Small-World Characteristics of Internet Topologies and Implications on Multicast Scaling

Shudong Jin
Department of Electrical Engineering and Computer Science, Case Western Reserve University
Cleveland, OH 44106
jins@cwru.edu

Azer Bestavros
Computer Science Department, Boston University
Boston, MA 02215
best@cs.bu.edu

Abstract

Recent work has shown that the physical connectivity of the Internet exhibits small-world behavior. Characterizing such behavior is important not only for generating realistic Internet topology, but also for the proper evaluation of large-scale content delivery techniques. Along this line, this paper tries to explain how the small-world behavior arises in the Internet topologies and how it impacts the performance of multicast techniques. First, we attribute the small-world behavior to two possible causes—namely the variability of vertex degree and the preference for local connections for vertices. We have found that both factors contribute with different relative degrees to the small-world behavior of the autonomous system (AS) level and router level Internet topologies. For the AS level topology, we observe that the high variability of vertex degree is sufficient to cause the small-world behavior, but for the router level topology, the preference for local connectivity plays a more important role. Second, we propose better models to generate small-world Internet topologies. Our models incorporate both causes of small-world behavior, and generate graphs closely resemble real Internet graphs. Third, using simulation we demonstrate the importance of our work by studying the scaling behavior of multicast techniques. We show that multicast tree size largely depends on the network topology. If topology generators capture only the variability of vertex degree, they are likely to underestimate the benefit of multicast techniques.

Keywords: Internet topology, small-world graphs, clustering, multicast

1 Introduction

Characterizing Internet topological properties, while interesting simply for the sake of discovery, is crucial for the evaluation of new protocols and design choices. Indeed, many significant innovations in the networking community in recent years have resulted from a more accurate understanding of the fundamental properties of that complex system. Not only it is important to characterize the Internet emergent properties, but also it is extremely valuable to explain *how* and *why* these properties emerge. Such an understanding would allow us to build models that could be used to generate synthetic artifacts (*e.g.*, large graphs) that resemble the real Internet well. Such synthetic artifacts are necessary for the simulation and proper evaluation of various network algorithms and protocols.

In this paper we focus on one aspect of Internet topology characterization that attracted much attention in recent years—namely the prevalence of the small-world phenomenon in the Internet routing graphs.

Smallworld graphs were mathematically formulated by Watts and Strogatz [38]. The structural properties of these graphs typically exhibit a short path length between any two vertices and strong clustering behavior. Using a careful analysis of real datasets, we study what cause the small-world phenomenon. We illustrate how and to what extent the small-world phenomenon is caused by the high variability of vertex degree, *i.e.*, the skewed vertex degree distribution; how and to what extent such a phenomenon is caused by the preference for local connectivity, *i.e.*, the property that the vertices tend to connect physically neighboring vertices. We show that both of these causes contribute with different relative degrees to the small-world behavior of the AS level and router level topologies. This finding provides a basis for more promising approaches to the development of more accurate Internet topology generators beyond degree-only models such as the Barabási-Albert model [5, 3]. We describe new models to generate Internet topologies that capture both the skewed vertex degree distributions and the preference for local connectivity, and show the small-world behavior in the synthetic graphs.

To demonstrate the usefulness of our topology characterization, we study the scaling behavior of multicast techniques–how multicast tree size increases with group size. The small-world behavior is relevant to multicast scaling. First a short path length implies usually two vertices are separated by few hops. Hence, there is a limited benefit of utilizing multicast to eliminate duplicate transmission. Second, on the other hand, the strong clustering implies that more likely local vertices share communication links to the other parts of the network. Hence, we will observe better (slower) multicast scaling. Previous studies, *e.g.* [31, 27], have analyzed multicast scaling in tree networks (or random networks) which are not small-world graphs. In this paper we simulate multicast techniques in different network topologies including Erdős-Rényi random graphs [16], random graphs with highly variable degree, our small-world graphs, and real Internet graphs. Our results show that using random degree-based topology generation models, it is likely to underestimate the efficacy of multicast techniques; but using our small-world models, the scaling behavior is more accurately portrayed.

The paper is organized as follows. Section 2 reviews recent work on topology characterization and multicast scaling. In Section 3, we describe the Internet routing graphs under study, and illustrate two possible causes of their small-world characteristics. We describe and evaluate our new models to generate small-world topologies in Section 4. In Section 5, we use simulation to study the scaling behavior of multicast techniques. Our findings and conclusions are summarized in Section 6.

2 Related Work

Recent work on Internet topology characterization has been related to one feature distinct from the early Erdős-Rényi random graph model [16]: highly variable vertex degrees. In such networks, vertices have a non-uniform probability of being connected to others, with some vertices having extremely large number of neighbors. The degree distributions were often observed to follow approximately a power-law [17].

Power-law networks are particularly emphasized by the work of Barabási and Albert [5, 3] who explored a promising class of models that yield strict power-law vertex degree distributions. In their model (called the B-A model), three generic mechanisms are defined: (1) Incremental growth, which follows from the observation that networks develop by adding new vertices or new edges. (2) Preferential connectivity, which relies on an observation that highly popular vertices are more likely to be connected again in the process of incremental growth, *i.e.* a so called "rich-get-richer" phenomenon. (3) Re-wiring, which removes some links randomly and re-wires them according to the preferential connectivity mechanism. The combined use of these mechanisms drives the evolution of the network to a steady state, in which the vertex degree distribution follows a power-law (so called a scale-free property).

There are debates on the ability of the B-A model to explain the evolution of the Internet. First, the mechanisms of the B-A model are found to be inconsistent with observations from the real Internet growth.

For example, preferential connectivity was shown to be stronger in the AS level graph growth and re-wiring was shown to be an insignificant factor [10]. Second, the strict power-law vertex degree distributions of the B-A model can not be confirmed [10]. Indeed, Internet object sizes are better captured by other distributions such as Weibull distribution [6]. This would imply that the high variability of vertex degrees in the AS level graphs may be the result of mechanisms [34] other than those in the B-A model. A more recent study [24] revealed that even the observed power-law degree distributions can be the result of sampling bias in traceroute-based measurement.

One interesting work [7] made some connections between power-law networks and small-world graphs [28, 38]. Small-world graphs exhibit connectivity properties that are between random and regular graphs (*e.g.*, regular lattices). Like regular graphs, they are highly clustered; yet like random graphs, they have typically short distances between arbitrary pairs of vertices. It has been shown that many networks have similar small-world property. In [7], the authors observed that most Internet topology generators capture power-law vertex degree distributions well, but usually do not do as well in capturing the clustering phenomenon exhibited in the Internet topologies. They proposed a variant of the B-A model by using a different preferential connectivity mechanism. Nevertheless, this variation of the B-A model is still degree-based. It obtains stronger clustering by generating a more skewed vertex degree distribution. No other causes of clustering and small-world behavior were identified. Moreover, their new model targets the AS level Internet topology, but not the router level topology.

In addition to the topology generators [26, 7] inspired by the B-A model, there are other generators that have been proposed and used to model the Internet. The Waxman model [39] extends the classical Erdős-Rényi model by randomly distributing vertices on a plane and creating edges by considering the distance between the vertices. The widely used GT-ITM Internet-specific topology generator [42] takes a hierarchical approach, capable of creating large graphs by composing smaller random graphs. Inet [21] assigns degrees to the vertices, following a power-law distribution, and then uses a linear preferential model to realize the assigned vertex degrees. The power-law random graph model in [2] generates degree distributions that strictly follow a power-law. A recent study [35] compared degree-based topology generators such as Inet and structural topology generators such as GT-ITM. Another study [33, 25] has focused on measurement techniques to infer network topologies more accurately and more completely.

2.1 Multicast Scaling and Topology

IP multicast was first proposed by Deering and Cheriton in [15], as a network level infrastructure that extends the Internet protocol in order to realize an efficient one-to-many communication service. Since then, much attention has been paid to algorithms for generating multicast trees and the actual implementation of these network layer multicast techniques. Shortest path tree is currently used in distance vector multicast routing protocol for Internet multicast traffic on the virtual multicast backbone (Mbone) network.

However, more than a decade after the protocols were developed, large parts of the Internet have yet to support native multicast. First, deployment of network layer multicast has not been adopted by many commercial ISPs. Second, even though many ISPs provide it, customers have yet to adopt and use native multicast. On the contrary, application level multicast approaches [12, 18, 9, 20, 30, 11, 32, 4] have been recently proposed as a viable alternative to IP multicast. In particular, end system multicast [12] has gained considerable attention.

Several recent studies have addressed the impact of topology on IP multicast using shortest path trees. In particular several studies considered how the size of a multicast tree increases with the size of the multicast group, *i.e.* the multicast scaling problem. Multicast scaling is important since it indicates the efficacy of utilizing multicast to reduce redundant transmission on the same links. The authors of [31] studied multicast scaling primarily under an assumption of k-ary tree topology. They provided a theoretical result which does not obey the Chuang-Sirbu law for IP multicast scaling [13] observed earlier. This law asserts that multicast

Table 1: Internet graphs used in this study.

Graphs	AS-2001	AS-2001+	Lucent	Scan+Lucent	Skitter	Skitter-CA
Number of vertices	11174	11461	112669	282672	258329	145067
Mean degree	4.19	5.71	3.21	3.15	3.39	2.44
Maximum degree	2389	2432	423	1973	412	359
Fit distribution	Power law ($\alpha \approx$ 1.22)	Power law ($\alpha \approx$ 1.22)	Weibull ($\beta \approx$ 0.49)	Power law ($\alpha \approx$ 1.7)	Weibull ($\beta \approx$ 0.48)	Weibull ($\beta \approx$ 0.48)

tree size increases as $n^{0.8}$, where n is the group size. This was generalized in [8] for more realistic multicast tree shapes. The authors of [27] considered a stricter approximation of the link cost reduction. Their results show that the Chuang-Sirbu law is not asymptotic for random graphs and k-ary trees. Recently, [1] also re-examined the analysis in [31] and provided precise asymptotic scaling behavior of tree size. This asymptotic term scales as $n(1 - \frac{\ln n}{\ln N})$ where n is the group size and N is the network size. More importantly, they showed that by replacing the k-ary complete tree topology by a self-similar tree, multicast tree size satisfies a power law. This finding strongly supports the claim that the essence of the problem lies in the modeling assumptions on the topology.

3 Characterizing Small-World Internet Topologies

This section describes the AS level and router level graphs studied in this paper, and provides evidence of their small-world behavior. We then examine the possible causes of the small-world phenomenon in Internet topologies.

3.1 Internet Graphs

Table 1 summarizes the AS level graphs and router level graphs used in this paper. The first two are AS level Internet graphs. One was obtained from the routing tables at route-views.oregon-ix.net. Since 1997, the routing tables have been collected once a day by the National Laboratory for Applied Network Research (NLANR) [29]. The graph we are using is dated May 26, 2001. Hereafter, we call it the “AS-2001” graph. It was found that the Oregon route-views is incomplete [10]. We obtained the second AS level graph from [36], which incorporates not only the Oregon route-views, but also the Looking glass data and the Routing registry data. This graph is also dated May 26, 2001. Hereafter, we call it the “AS-2001+” graph. Although AS-2001+ has only a few more vertices, it has much more edges (it has about 40% more edges. Therefore, it has higher average vertex degree).

We use four router level graphs. The first was obtained from traceroutes collected by the Internet Mapping project at Lucent Bell Laboratories around November 1999. Hereafter, this graph is called the “Lucent” graph. The second router level graph was obtained by merging the Lucent graph and the SCAN graph obtained around October/November 1999 using the Mercator software [19]. Hereafter, this graph is called the “Scan+Lucent” graph. In preprocessing both graphs, we discarded a few edges with undefined vertices. The percentage of these edges was negligible. These two graphs are available at [37]. These maps are relatively obsolete, and we also suspect the extremely high degrees of some nodes. Thus, we obtained two more recent maps from [14]. They were measured using the Skitter tool. We obtained a snapshot on October 15, 2002, recorded by 13 of the 19 Skitter monitors worldwide. We merged the corresponding 13 subgraphs and obtained the “Skitter” graph. In addition, we also use a subgraph recorded by the monitor located at Palo Alto, CA. This subgraph is called the “Skitter-CA” graph. The inclusion of this subgraph is intent to indicate that the observed small-world invariants are not due to the incompleteness of the graphs.

Figure 1 shows the complementary cumulative distribution function (CCDF) $1 - F(d)$ of vertex degrees for the six graphs under study. Here $F(d)$ is defined as the probability that a vertex has degree not higher than d. The CCDF quantifies the probability that a vertex has a degree d larger than a certain value. A common property of these graphs is their long tails. That is, vertex degree is highly variable. The level of variability appears to be different though.

Previous work [17] has shown that vertex degree distributions follow a power-law. This is best illustrated by the AS-2001 graph. With a power-law distribution, $1 - F(d) = cd^{-\alpha}$, where c and α are constants, and the log-log scale plot of CCDF is a straight line, as shown in Figure 1(a). Using a linear regression, we estimated that for the AS-2001 graph, the exponent α is close to 1.22. For the AS-2001+ graph in Figure 1(b), an observation is that the distribution follows a power-law less well [10].

For the router level graphs, we have also estimated their power-law exponents. However, as evident in Figure 1(c-f), it appears that the power-law distribution for the router level graphs does not fit our empirical datasets. This is particularly the case for the Lucent graph and for the Skitter graphs, for which the tails of the distributions drop faster than any power-law. In [6], Weibull distribution was found to provide a good fit to many Internet object size distributions. The Weibull distribution is one of the widely used lifetime distributions in reliability engineering. Its tail takes on the form $e^{-(x/\eta)^{\beta}}$, where η is the scale parameter and β is the shape parameter. Using rank regression on Y-axis, we estimated that the Weibull fits these empirical distributions better, though not perfectly.

Noticeably, these router level maps have some strange characteristics. The degree distributions indicate that many routers have have high degrees. We suspect the accuracy of these measurements. In particular, the maps could have included Internet Exchange Points (IXPs). Strictly, the router level maps should not contain such layer-2 nodes. Nevertheless, we believe, since the number of IXPs is small (according to NANOG, there are currently only over one hundred IXPs globally), they will not change the heavy-tailed nature of degree distribution.

3.2 Small-world Behavior

In an influential paper [38], Watts and Strogatz defined a range of graphs termed *small-world graphs*. Small-world graphs are highly clustered, like regular graphs such as lattices; yet they have typically short distances between arbitrary pairs of vertices, like random graphs. The structural properties of these graphs are quantified by two metrics: the characteristic path length L and the clustering coefficient C. As in [38], we define

Definition 1 *Characteristic path length L is the number of edges in the shortest path between two vertices, averaged over all pairs of vertices.*

Definition 2 *Clustering coefficient C is defined as follows. Consider a vertex v which has k_v neighbors. There are at most $k_v(k_v - 1)/2$ edges among these k_v neighbors. Let C_v denote the fraction of these edges that actually exist. C is the average of C_v over all vertices with degree at least two.*[1]

Following the original definition in [38], a small-world graph has two properties: (1) its L is not much larger than L_{random}, the characteristic path length of a random graph with the same number of vertices and edges, and (2) its C is much larger than C_{random}, the clustering coefficient of a random graph.[2] It is

[1] Since C_v is undefined when $k_v = 1$, this averaging excludes vertices with only one neighbor. If a graph has many vertices with degree one (especially when vertex degree is highly skewed), then all of them are ignored. We believe it is a limitation of the original metrics.

[2] In the original definition, it is not clear what is *not much larger* L value and what is *much larger* C value. We consider L not much larger if it is only a fraction larger than that of a random graph. We consider C much larger if it is orders of magnitude larger than that of a random graph.

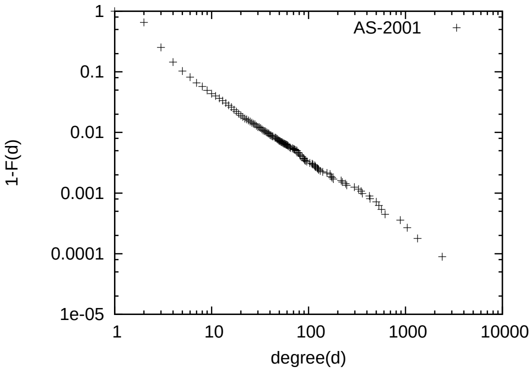

(a) AS-2001

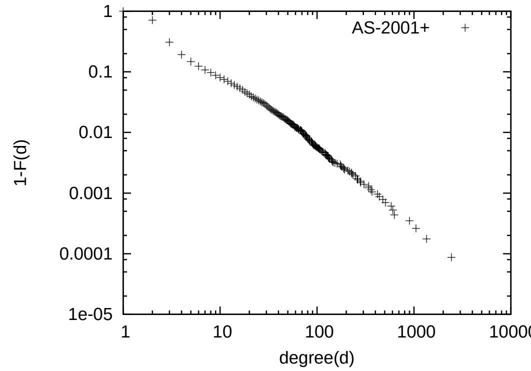

(b) AS-2001+

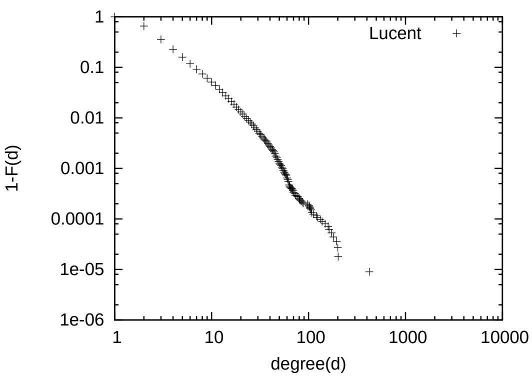

(c) Lucent

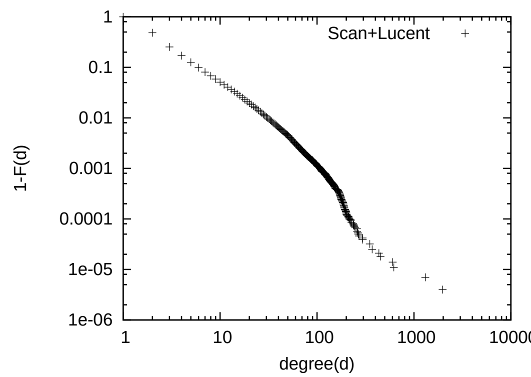

(d) Scan+Lucent

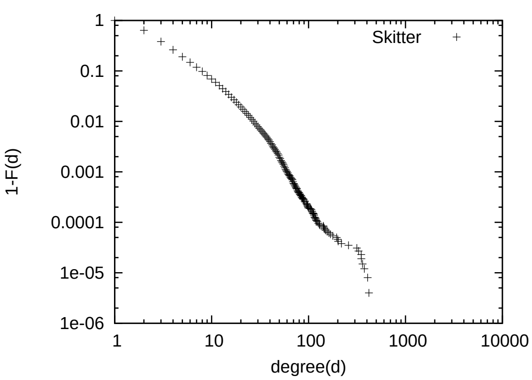

(e) Skitter

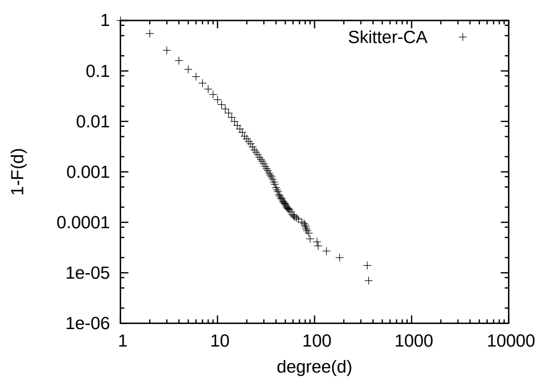

(f) Skitter-CA

Figure 1: Vertex degree distributions of the AS level and router level Internet graphs.

Table 2: Characteristic path length and clustering coefficient of the Internet graphs.

Graphs	AS-2001	AS-2001+	Lucent	Scan+Lucent	Skitter	Skitter-CA
L	3.637	3.535	10.02	8.803	11.36	16.26
C	0.432	0.4943	0.1001	0.0996	0.0233	0.0109
L_{random}	6.797	5.547	10.49	11.38	10.12	15.73
C_{random}	0.0003	0.0003	0.000022	0.000007	0.000014	0.000012

not difficult to see that for a connected random graph with N vertices and with an average degree of k, $L_{\text{random}} \sim \ln N / \ln k$ and $C_{\text{random}} \sim k/N$.

To show that the small-world phenomenon holds for both the AS level and router level Internet graphs, we computed their characteristic path lengths and clustering coefficients. Results are shown in Table 2. For comparison, we have also generated corresponding random graphs with approximately the same number of vertices and edges, and computed L_{random} and C_{random} by averaging over a number of random graph instances. The random graphs can be disconnected. In such cases, the largest connected components are used. Table 2 shows that the values of L for the Internet graphs are not much larger than L_{random}. Indeed, in most instances, L is smaller than L_{random}. Table 2 also shows that C is larger than C_{random} by 3-to-5 orders of magnitude (Note that this significant difference may not be the artifact due to the incompleteness of the Internet graphs under study). These two observations provide clear evidence of the small-world phenomenon in the Internet topologies.

In [7], AS level graphs were found to exhibit small-world behavior. Our findings complement their work by showing that the router level graphs are also small-world.

3.3 Cause One: High Variability of Vertex Degree

One question is whether the variability of vertex degree can introduce small-world behavior, given that all Internet graphs under study have highly variable degree.

To study the effects of high variability of vertex degree distributions, we generated random graphs whose vertex degree distribution follows a power law. For such a distribution, the CCDF is defined as $1 - F(d) = cd^{-\alpha}$, where c is a constant. Appendix A.1 explains how we use a random matching to generate graphs with highly variable degrees. We realized that there can be different random matching algorithms such as the one in [2]. However, this random matching is successful in realizing strict power-law degree distributions. In addition, if the resulting graph is disconnected, we choose its largest connected component. We do not use the power-law random graph model in [2], since it results in a multi-graph with duplicate edges and self-loops. Removing the duplicate edges and self-loops may result in graphs with vertex degree distributions of lower variability.

We generated graphs with about 10000 vertices and 100000 vertices, respectively. The average vertex degree is fixed at 4.2. The value of α varies over a wide range. The constant factor c is determined such that the average vertex degree is roughly equal to 4.2. We ensured that neither the number of vertices nor the average vertex degree of generated graphs departs from their targets by more than 2%. We computed the characteristic path length and clustering coefficient of these graphs, and plotted them against α as shown in Figure 2. Each point represents one graph instance.

Figure 2 indicates that smaller α values result in shorter characteristic path lengths and much larger clustering coefficients. Note that a smaller α means higher variability of vertex degree distribution. The presence of both short characteristic path length and high clustering coefficient is the signature of small-world graphs. Thus, we conclude that a highly skewed power-law vertex degree distribution is a possible

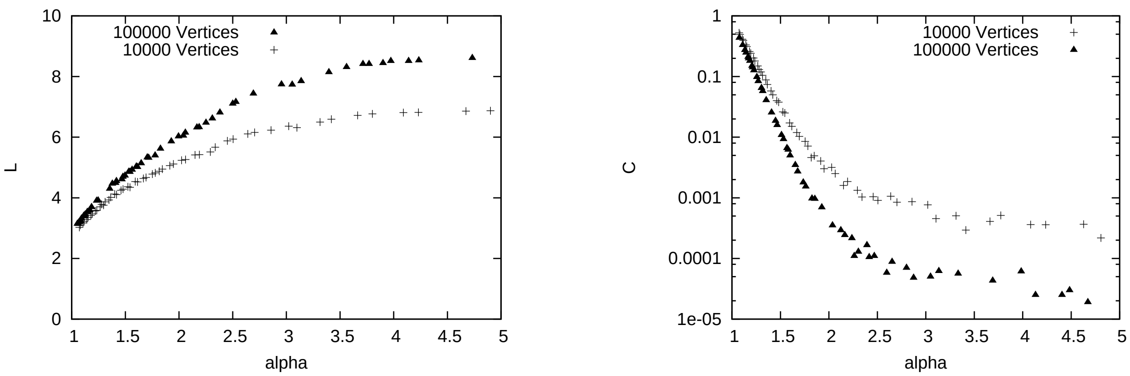

(a) Characteristic path length (b) Clustering coefficient

Figure 2: Power Law vertex degree distribution results in small-world behavior. Smaller power-law exponent (higher variability) results in shorter path length and stronger clustering.

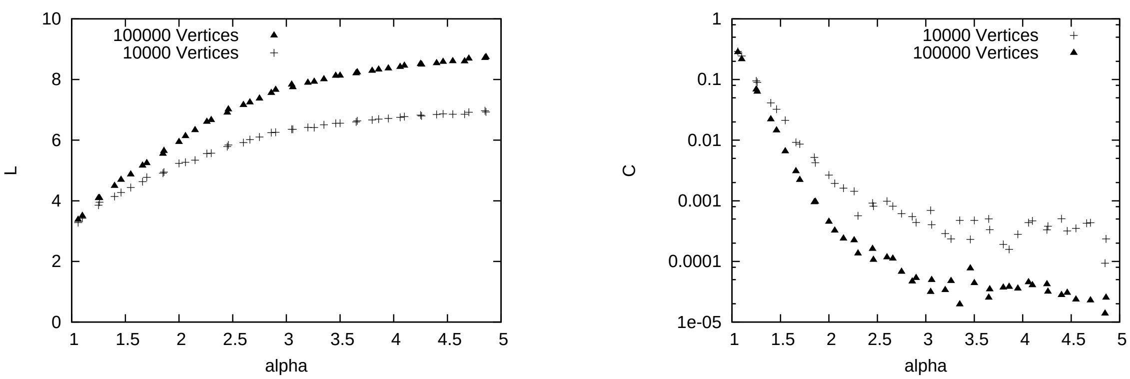

(a) Characteristic path length (b) Clustering coefficient

Figure 3: Power Law vertex degree distribution results in small-world behavior. Graphs were generated using the random graph model in [2].

cause of small-world behavior.

One question is, whether Figure 2 simply shows artifacts of our particular random matching algorithm. To answer this, we also used the random matching model in [2]. This model results in a multi-graph with duplicate edges and self-loops. Removing the duplicate edges and self-loops may result in graphs with vertex degree distributions of lower variability. We notice that when degree variability is not high, this model is still successful. When variability is very high (α close to unity), there are many duplicate edges and self-loops. We removed them and found usually the largest connected component still contains most vertices. We computed the characteristic path length and clustering coefficient of these graphs, and plotted them against α as shown in Figure 3. This figure shows the conclusions do not change when we use different random matching models.

Another question remains. It has been often observed that vertex degree distributions do not fit power-law distributions well [6, 10, 34]. Nevertheless, we found that as long as the vertex degree exhibits high variability, other distributions can also give rise to small-world behavior. To show this, we have generated random graphs whose vertex degree follows a Weibull distribution, still using the random matching model in Appendix A.1. The CCDF of Weibull distribution is $e^{-(x/\eta)^{\beta}}$, where η is the scale parameter and β is the

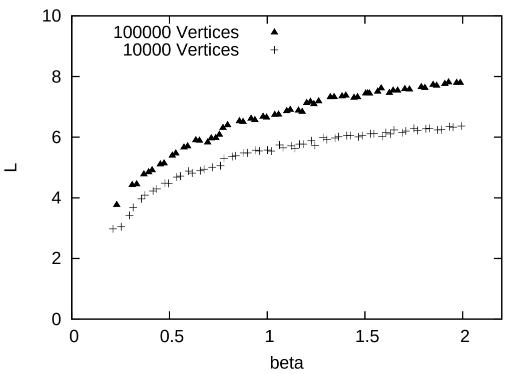

(a) Characteristic path length

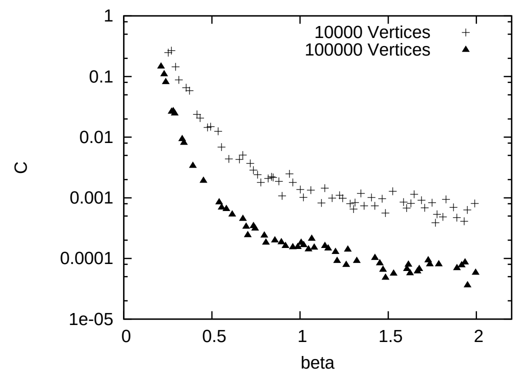

(b) Clustering coefficient

Figure 4: Weibull vertex degree distribution results in small-world behavior. Heavier-tailed vertex degree distribution also leads to shorter path length and stronger clustering.

shape parameter. The value of β varies from 0.2 to 2.0, and the value of η is determined such that the average vertex degree is still 4.2. For each generated graph, we computed its L value and C value. The results are plotted with varying β in Figure 4. Notice that the results we obtained here using a Weibull distribution are similar to those we obtained using a power-law distribution. Moreover, we observed that a smaller value of β for Weibull distribution (*i.e.*, heavier tails and higher variability of vertex degree) results in smaller L and in much larger C. Thus, we conjecture that, high variability of vertex degree distributions, whether it is the result of a power law or that of other distributions, can cause small-world behavior.

3.4 Cause Two: Preference for Local Connectivity

So far we have shown that the high variability of degree results in small-world behavior. But, are there other causes? To answer this question, we first conducted the following experiment. We generated graphs whose vertex degree distribution follows exactly the same distribution of the real Internet graphs. However, the edges were created randomly using a random matching. In this way, we preserved the high variability of vertex degree, but destroyed other topological properties that may exist in real Internet graphs. We call these synthetic graphs *randomized Internet graphs*.

Table 3: Characteristic path length and clustering coefficient of randomized Internet graphs.

Graphs	Randomized AS-2001	Randomized AS-2001+	Randomized Lucent	Randomized Scan+Lucent	Randomized Skitter	Randomized Skitter-CA
L	3.406	3.346	6.944	5.971	6.704	8.748
C	0.2688	0.2507	0.00028	0.00076	0.00016	0.00007

For each real Internet graph, we generated a number of randomized instances, and for each we computed their L and C values, averaged over all randomized instances. A randomized instance is often disconnected. However, in our experiment, its largest connected component contained more than 90% of the vertices in the corresponding real graphs. The results are reported in Table 3. Comparing these results with those of the real graphs in Table 2 (the first two rows), we can make the following observations. First, for the AS level graphs, the L values are very close to each other and the C values differ only by a small fraction (although the

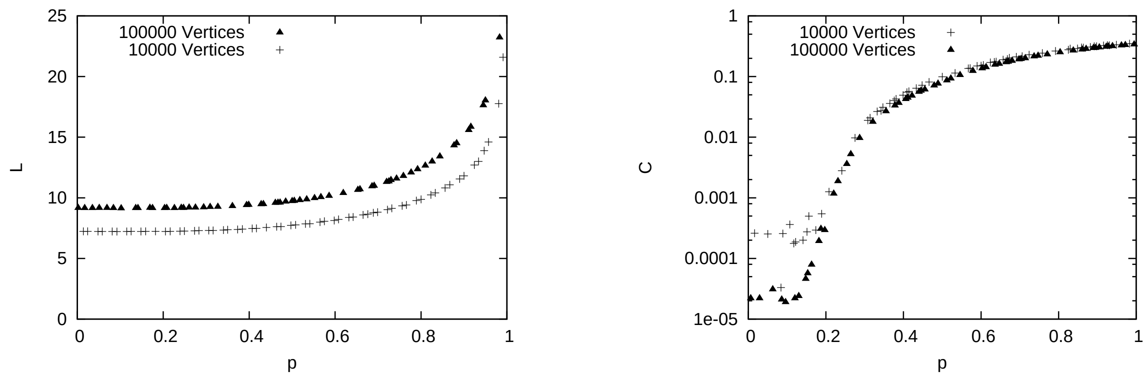

(a) Characteristic path length

(b) Clustering coeffi cient

Figure 5: Small-world behavior as the result of preference for local connectivity.

absolute differences are large). Second, for the router level graphs, the L values are considerably different, and the C values differ by several orders of magnitude.

Our conclusions from this experiment are the follows. First, clearly there are other causes that contribute to the small-world behavior of the Internet topologies. Those other causes lead to larger clustering coeffi cient and longer path length. Second, it appears that those other causes are more pronounced when the variability of vertex degree distributions is not extremely high. In other words, when the variability of vertex degree distributions is extremely high, which is the case for the AS level graphs, the effect of those other causes is overshadowed. When the variability of vertex degree distributions is moderate or low, which is the case for the router level graphs, the effect of those other causes is evident.

So, what are these other causes of small-world behavior in the Internet graphs? In attempting to answer this question, we do not intend to provide a complete and exclusive explanation of small-world phenomenon, but to identify one plausible explanation—namely, the preference for local connectivity in the Internet. Indeed, this possible explanation was inspired by the work of Watts and Strogatz [38], who found that if only a portion of the edges of a regular lattice are reconnected randomly, but most edges are intact, then the resulting graph would exhibit small-world behavior. By doing so, the clustering coeffi cient remains large due to local connectivity, but the characteristic path becomes closer to that of random graphs due to remote connectivity.

To illustrate this possible cause of small-world phenomenon, we generated a set of 10000 and 100000 vertices, which we randomly placed on a two-dimensional plane. We set the average vertex degree to 4.2. The vertex distribution is the same as that of a random graph with the same number of vertices and edges. Specifi cally, the distribution has an exponentially decaying tail, *i.e.* low variability. Connections between vertices were made as follows. For each vertex v with degree k_v, on average v is connected to its pk_v nearest neighbors. The other edges of v are random. Here, $0 \leq p \leq 1$ is the probability of local connectivity, which we call *local probability*. We varied p to generate many graph instances and computed their L values and C values. The results are plotted in Figure 5, where each point represents one graph instance.

The results in Figure 5 reveal that preference for local connectivity leads to small-world graphs. First, the characteristic path length increases slowly when p is small or moderate. Second, the clustering coeffi cient increases drastically when p is small or moderate. In overall, there is a wide intermediate regime for p that yields the characteristic signature of small-world graphs.

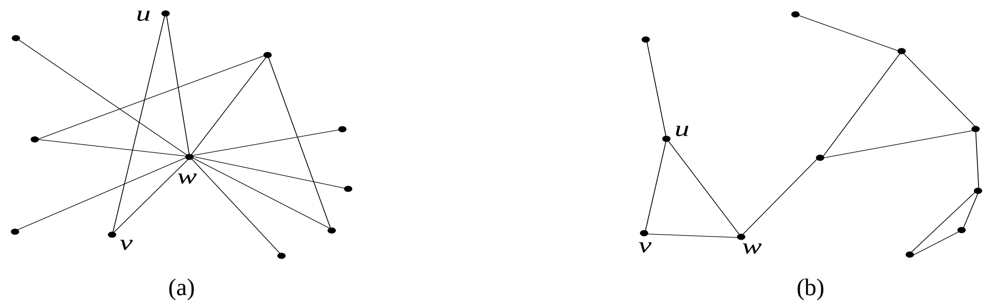

Figure 6: An illustration of strong clustering caused by (a) high variability of vertex degree, as vertices tend to have common neighbors (those with extremely high degrees) and hence tend to have large clustering coefficient, (b) preference for local connectivity, as nearby vertices tend to form triangles and hence tend to have large clustering coefficient.

3.5 Discussions

How does high variability of vertex degree distributions result in small-world behavior? With such high variability, it is likely that two interconnected vertices, say u and v, will have the same neighbor, say w. This occurs more frequently when w is a vertex with an extremely high degree, even though the edges are made random. An example is shown by Figure 6(a). It means that u, v, and w form a triangle. Such a pattern contributes directly to the computation of C_u, C_v, and C_w, and results in larger average clustering coefficient C, according to its definition. Intuitively, C grows with the variability of vertex degree. Also, notice that with highly variable vertex degree, the average distance between two vertices (L) is short. This is because the shortest path is usually through those extremely popular vertices. That is, highly popular vertices serve as good navigators in the graph.

How does preference for local connectivity result in small-world behavior? The answer to this question is straightforward. With a non-negligible probability of a local connection, if a vertex u is connected to v and w, then it is likely that v and w are also close to each other. See the example in Figure 6(b). As a result, there is a non-negligible probability that a triangle will form among these vertices, resulting in a higher clustering coefficient. Meanwhile, as long as there are still many long-range edges (short-cuts) in the graph, it is still easy to find a short path between two randomly chosen vertices. Those vertices with long-range edges serve as good navigators.

Both high variability of vertex degree distributions and the preference for local connectivity appear to be plausible causes of small-world behavior. In particular, the highly variable nature of vertex degree distributions is similar to the high-variable nature of many other Internet artifacts. Such distributions may be the result of some specific processes related to the evolution of these artifacts [5], or they may exist due to other reasons [34]. Preference for local connectivity may be explained for both the router level and AS level topologies as follows. At the router level, links are created by considering the distances since there are tighter physical constraints (there are certainly other relevant factors such as administrative considerations, which may cause hierarchy of the topology). At the AS level, although the physical distance in a space is less important, there can be another kind of locality, for example Internet Service Providers (ISPs) may form cliques due to their business relationship.

It is important to note that different causes contribute with different relative degrees to the small-world behavior of both AS level and router level maps. AS level connectivity and router level connectivity have certainly different characteristics. Our study shows that the the causes of their small-world behavior are also different. At AS level, the Internet map is largely policy-based. Preference for local connectivity makes less sense. On the other hand, at router level, physical distance is more important in determining the connectivity.

It is important to note that, the entire router level map is not a simple overlapping of individual ISP maps. It is more or less affected by policies. Our discussion has assumed a flat router level map, but the actual Internet is hierarchical, with different kinds of routers, *e.g.*, border gateway routers versus interior routers.

To summarize our findings in this section:

- Higher variability of vertex degree distributions leads to shorter characteristic path length and larger clustering coefficient. Extremely high variability gives rise to small-world behavior, which is true for the AS level topology.
- When the variability of vertex degree distributions is not high enough, it alone does not cause small-world behavior, which is true for the router level topology.
- Preference for local connectivity is a possible cause of small-world behavior. This cause appears more pronounced for the router level topology.

Our findings imply the failure of the Barabási-Albert model as an explanation of Internet growth and evolution [5]. Recently in [41], the same group recognized the limitations of their original model and combined distance dependence in a new hybrid model. Several other studies [10, 34, 6, 40] have also casted doubts on the adequacy of the Barabási-Albert model. However, none has examined the causes of small-world phenomenon as evidence. Specifically, the Barabási-Albert model targets power-law degree distributions. With only power-law degree distributions, the resulting graphs tend to have shorter characteristic path lengths and lower clustering coefficients. The comparisons between Table 2 and Table 3 provide the evidence. When the power-law exponent α departs much from unity, the Barabási-Albert model fails to generate small-world graphs at all. This result also calls for better models to generate more realistic Internet topologies.

4 Generating Small-world Internet Topologies

Based on our characterization, this section describes two models to generate small-world Internet topology with variable vertex degree. The models also capture the other cause of the small-world behavior, the preference for local connectivity.

4.1 The First Model

Figure 7 describes the first model. It is derived directly from our characterization. Given a sequence of vertex degrees and the probability of local connection, this model generates graphs exhibiting small-world behavior. To do so, it first creates local connections for each vertex, then the remaining edges are created using a random matching. Note, The sequence of degrees can be obtained from a distribution (*e.g.*, a power-law) or obtained from a real Internet graph. The resulting graph may be disconnected. In such a case, the largest connected component is used.

We have used the above model to generate graphs which exhibit a wide range of small-world behavior. For example, we generated graphs with 10000 vertices and with average degree 4.2. Vertex degree follows a power-law with exponent α. We varied parameters α from 1.05 to nearly 8.0 and p between 0 and 1. For each pair of α and p, we generated ten graphs, computed their average L value and C value, and plotted them in Figure 8. We observed that when α is small or p is moderate, both the characteristic path length and the clustering coefficient satisfy the requirements for small-world graphs.

We have also compared the synthetic graphs to actual Internet maps. We found that for the AS level graphs, if we pick α=1.22, no matter what is the value of p, average path length is very close to 3.5, the path length of AS level map. The difference is usually within 5%. However, when p is close to 0, the clustering coefficient is often less than 0.3, much smaller than the target value, 0.43-0.49. A moderate value of p works

Given a sequence of vertex degrees and the local probability p, generating a graph with N vertices as follows:

(1) Randomly place N vertices on a plane. A degree d_v is assigned to each vertex v, $1 \leq v \leq N$.

(2) Create local connections among the vertices. Connect each vertex v to its nearest pd_v neighbors. [a]

(3) Create remote connections among the vertices. Let d'_v be the number of edges already created for vertex v (the result of step (2)). Then, $(d_v - d'_v)$ more random edges are created for each vertex v. This is done by using the random matching model described in Appendix A.1.

[a] Precisely, pd_v is rounded down to $\lfloor pd_v \rfloor$ or rounded up to $\lceil pd_v \rceil$, in a probabilistic way. That is, if a real number r is closer to its ceiling, it is more likely to be rounded up. With probability $r - \lfloor r \rfloor$, r is rounded up. In this way, the expected value of the result is equal to r.

Figure 7: The first model for generating small-world graphs with variable degree.

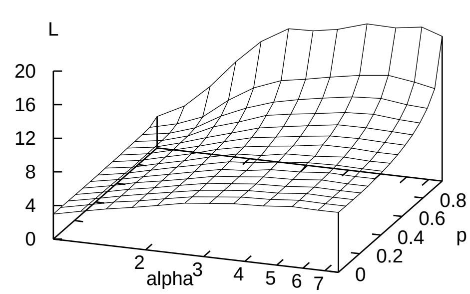

(a) Characteristic path length

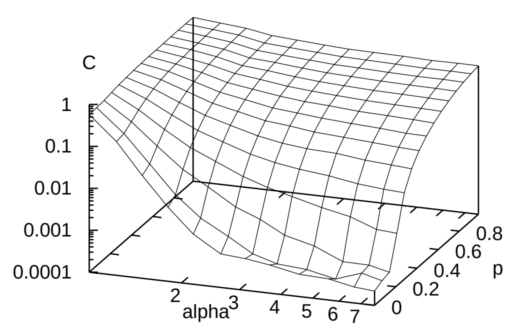

(b) Clustering coefficient

Figure 8: A range of small-world graphs generated by the first model.

Given a sequence of vertex degrees and the exponent of distance-dependence r, generate a graph with N vertices as follows:

(1) Randomly place N vertices on a plane. A degree d_v is assigned to each vertex v, $1 \leq v \leq N$.

(2) Create connections among the vertices. For each vertex v, in descending d_v order, repeat create connections (if not enough yet) as follows:

- Choose u with probability proportional to $\frac{d_u}{l_{u,v}^r}$, where $l_{u,v}$ is the Euclidean distance between u and v, such that (i) $u \neq v$ and (ii) there is no edge (u, v) yet. Create edge (u, v).

Figure 9: The second model for generating small-world graphs with variable degree.

well. For the router level graphs, the value of p is consistently important. When p is too close to 0, the average path length is always too large. Again, a moderate value of p works well.

4.2 The Second Model

The first model creates either local or remote connections. One issue is whether the probability of having an edge between two vertices can be explicitly expressed as a function of their distance. Previous studies have shown such distance-dependence exists [23, 41]. Notice that our model can be more realistic by explicitly defining such distance-dependent probabilities. Let us consider the following model: the probability of having an edge between two vertices u and v is proportional to $l_{u,v}^{-r}$, where $l_{u,v}$ is the distance between u and v and r is a non-negative constant. It was originally used by Kleinberg [22] to generalize the Watts-Strogatz model. The key is the choice of r. If r is larger, $l_{u,v}^{-r}$ decays faster and the edge tends to be local; but r can not be too large since otherwise the lack of long-range connections would lead to an excessively large L value. If r is smaller, $l_{u,v}^{-r}$ decays slowly and long-range connections become frequent; but r can not be too small since otherwise the graphs would be close to random (*i.e.*, the value of C would be excessively small). Using this model, Figure 9 describes an implementation.

Using this model, we also generated graphs which exhibit a wide range of small-world behavior. For example, we generated graphs with 10000 vertices and with average degree 4.2. Vertex degree follows a power-law with exponent α. We varied parameters α from 1.05 to nearly 8.0 and r between 0 and 4.5. For each pair of α and r, we generated ten graphs, computed their average L value and C value, and plotted them in Figure 10.

We observed that when α is small or r is moderate, both the characteristic path length and the clustering coefficient satisfy the requirements for small-world graphs. We also compared these synthetic graphs to the Internet maps. For example, for the AS level graphs, when $\alpha = 1.22$, no matter what is the value of r, the average path length is very close to the target value 3.5. From this experiment, we found the second model is as effective as the first model.

4.3 Synthetic Internet Graphs

The last two subsections have shown that by considering preference for local connectivity, there are different and easy ways to generate small-world topologies while still preserving the highly variable degree sequence. In this subsection, we try to generate synthetic Internet graphs closely resemble the real Internet graphs. We describe here our effort to use the second model to synthesize the Lucent router level graph.

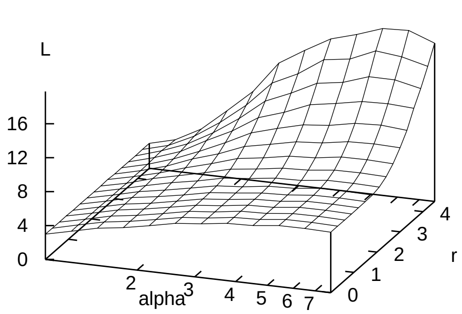

(a) Characteristic path length

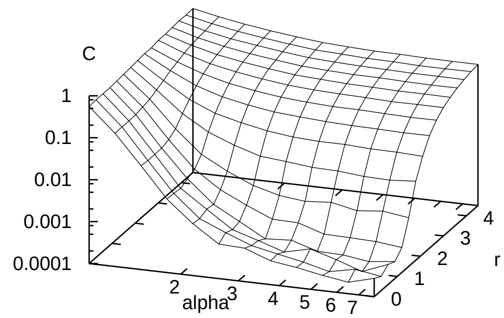

(b) Clustering coeffi cient

Figure 10: A range of small-world graphs generated by the second model.

We obtained the sequence of vertex degree from the real graph, and fed it into the graph generator. The other parameter r was set to 3.15. The output largest connected component has 111328 vertices and the mean degree is 3.24. We then computed its characteristic path length 9.814 and clustering coeffi cient 0.193. Comparing them to those of the real graph, we found that the path lengths are close but the clustering coeffi cients are still different. A possible explanation is as follows. The model assumes the probability of having an edge between two vertices u and v is proportional to $l_{u,v}^{-r}$. But the distance-dependence in the real graph may not follow a power-law well [23]. In other words, the power-law relationship may over-estimate local preference when $l_{u,v}$ is small. As a result, the model tends to produce higher clustering coeffi cient. In the next section, this graph is used in our simulation study of multicast scaling.

5 Implications on Multicast Scaling

In this section, we use simulations to study the scaling behavior of IP multicast and end-system multicast techniques, *i.e.*, how the multicast tree size increases with the group size. In our simulations, we used three topologies in comparison. The fi rst is the original Lucent router level graph in Table 1. The second is generated using the model described in Appendix A.1. We found the generated largest connected component (after eliminating duplicated edges and self-loops) has 110986 vertices and mean degree 3.18. Notice that this graph preserves the original vertex degree distribution in the Lucent graph only. The third topology, as we described before, is generated using our small-world model in Figure 9. It has 111328 vertices and the mean degree is 3.24. These three graphs have approximately the same size and mean degree. However, their small-world behaviors are different. Our small-world graph is more resemble the real Lucent graph.[3]

To study and scaling behavior of IP multicast with shortest path tree, we proceeded as follows. We randomly chose a vertex as the server and n other vertices as clients. Then a shortest path tree is constructed. The size of this tree was computed and compared to the average distance from these clients to the server. The ratio is the normalized multicast tree size and it reflects how the tree size increases with n. For each n we repeated above simulation many times and computed the average multicast tree size. By varying n we plotted the scaling behavior of IP multicast tree. Figure 11 compares the results using the Lucent graph and using the degree-only graph. Figure 12 compares the results using the Lucent graph and using our

[3]This paper focuses on the small-world aspect of the Internet topology. Thus we consider mainly two metrics, characteristic path length and clustering coeffi cient. We have also conducted experiments to compare the graphs using other metrics.

small-world graph.

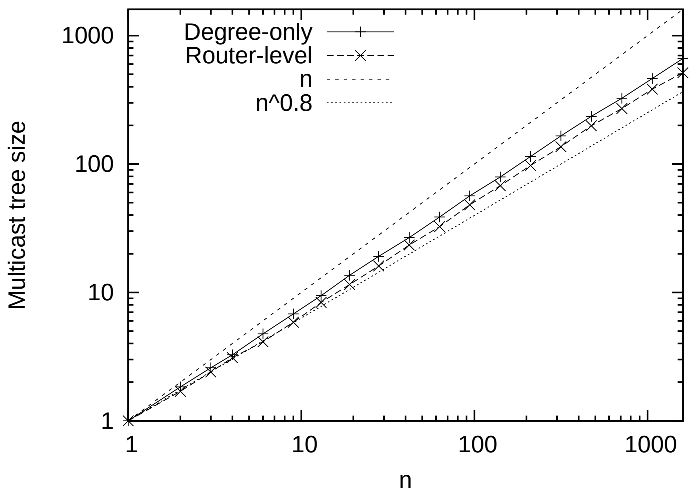

Figure 11: Comparison of the scaling behaviors of IP multicast using degree-only topology and real Internet topology.

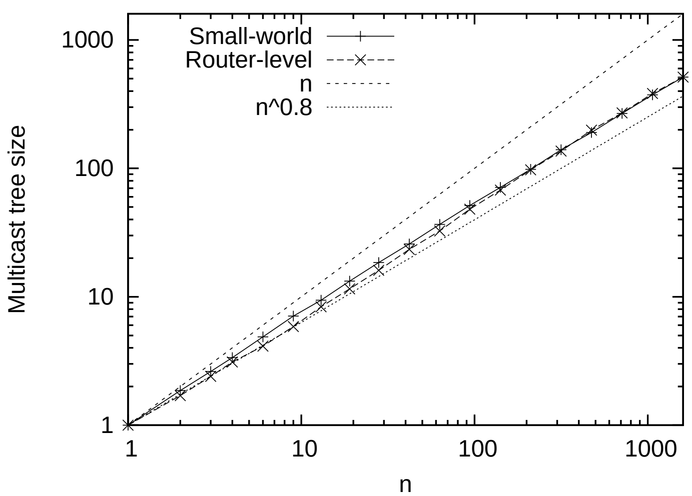

Figure 12: Comparison of the scaling behaviors of IP multicast using small-world topology and real Internet topology.

We observed that using the degree-only graph, IP multicast has worse scaling behavior than it using the real Lucent graph. On the contrary, using our small-world graph, the scaling behavior fits that using the real Lucent graph very well. The multicast tree sizes usually differ by 30-40% when different topology models are used. It thus suggests the importance of capturing the small-world behavior in the Internet topologies. Without doing so, flawed topology generation models may underestimate the benefit of IP multicast in reducing overall network cost.

We should also point out that accurate modeling of the Internet topology helps understand the so-called Chuang-Sirbu law [13] on IP multicast scaling. To better illustrate this point, we have also shown the log-scale plots of our simulation results in the figures. The Chuang-Sirbu asserts that multicast tree size increases as $n^{0.8}$. Since then several studies [31, 8, 27, 1] tried to explained it. From our simulation, we can see two points. First multicast tree size increases faster than $n^{0.8}$, *i.e.*, the exponent is larger. This suggests that the Chuang-Sirbu law is at best an approximation, and it overestimates the efficacy of IP multicast. Second we conjecture that small-world behavior (coupled with the variable degree nature in the Internet topologies) may draw the true scaling behavior line toward the Chuang-Sirbu law.

Next we studied the scaling behavior of the end-system multicast scheme originally described in [12]. There are one randomly-chosen sender and n randomly chosen receivers in the entire network. A complete virtual graph is constructed. This graph consists of $(n + 1)$ vertices corresponding to all participants, with a virtual link (with distance) between any pair of vertices. A minimum spanning tree is then constructed from this complete virtual graph. Note in the actual implementation of end-system multicast, a mesh network is constructed instead of complete virtual graph. Then a spanning tree is obtained from this mesh network. In our evaluation we didn't have any metrics (e.g., link bandwidth and latency) to facilitate mesh construction. We believe our method results in a smaller end-system multicast tree size. We expect this modification doesn't change our conclusion, since we are studying the scaling behavior.

We then computed the size of the minimum spanning tree (for each n we computed tree size many times and took the average). By varying n we plotted the scaling behavior of end-system multicast. Figure 13 compares the results using the Lucent graph and using the degree-only graph. Figure 14 compares the results using the Lucent graph and using our small-world graph.

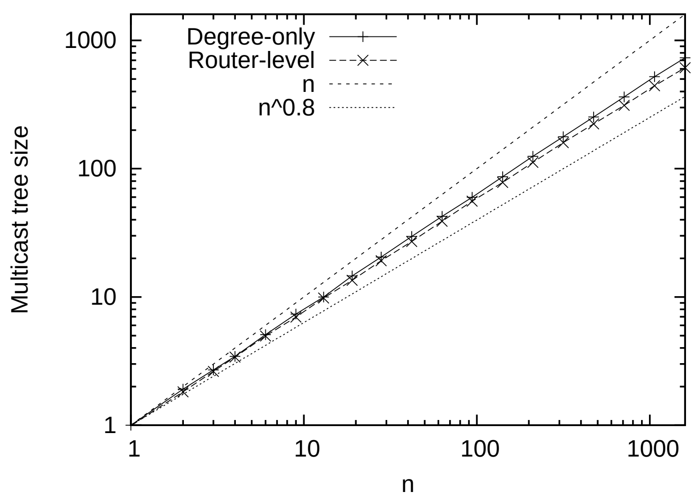

Figure 13: Comparison of the scaling behaviors of end-system multicast using degree-only topology and real Internet topology.

We observed that using the degree-only graph, end-system multicast also has worse scaling behavior than it using the real Lucent graph. It might be the same reason: the degree-only graph does not capture the strong clustering in the real Lucent graph. On the contrary, using our small-world graph, the scaling behavior is slightly better than that using the real Lucent graph. This result has also surprised us a little. There are two potential explanations. First our small-world graph has slightly higher clustering coefficient than that of the actual router level graph. Second the actual router level graph may have topological properties other than clustering behavior. Such properties are not captured by our small-world models.

To summarize, these experiments suggest the importance of capturing the small-world behavior in the Internet topologies. Otherwise, it is likely to underestimate the benefit of using multicast.

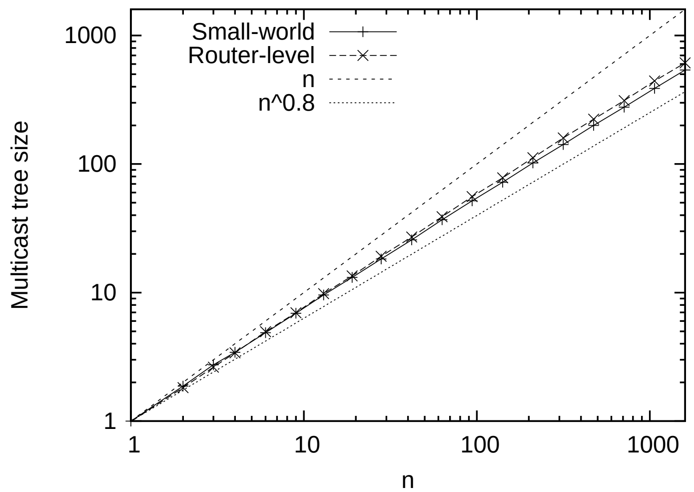

Figure 14: Comparison of the scaling behaviors of end-system multicast using small-world topology and real Internet topology.

6 Conclusions and Future Work

This paper has focused on the small-world aspect of the Internet topology. We illustrated the possible causes of such behavior in the physical connectivity of the Internet, and demonstrated its importance in evaluating the scaling behavior of multicast techniques. Our main findings and conclusions are:

- Two possible causes of small-world behavior in the Internet are (i) high variability of vertex degree and (ii) preference for local connectivity. Our analysis revealed that extremely high variability may give rise to the small-world behavior in AS level topology; but when the variability of vertex degree distributions is moderate, clustering in small-world graphs is mainly caused by local connectivity, which is true for router level topology.

- If Internet topology generators target vertex degree distributions only, they may generate less realistic topologies. We proposed more promising models to generate network topologies that captures both the vertex degree distributions and the preference for local connectivity. By doing so, it is easy to generate small-world synthetic graphs. Indeed, there are different ways to achieve the goal.

- Multicast tree size depends on the small-world behavior of the Internet topology. If Internet topology generators target the variable vertex degrees only, then it is likely to underestimate the benefit of multicast techniques.

We are currently extending our small-world topology model in several directions. First, we hope to capture a more realistic distribution of vertices in the plane. For example, Internet objects are distributed with highly skewed geographical density [19, 23]. Second, we hope to model the incremental growth process of the Internet that could give rise to the scale-free small-world behavior.

A.1 Random Graphs with Variable Vertex Degrees

We describe how we generate a random graph, given a set of N vertices with possibly highly variable degrees d_v, $1 \leq v \leq N$. Our technique modifies the original model in [2] and creates a random matching of the vertex degrees as follows.

Given a set of N vertices with possibly highly variable degrees d_i, $1 \leq i \leq N$.

(1) Form a set S containing d_v distinct copies of each vertex v.

(2) For each v in S, in descending d_v order:

– Choose u randomly from S such that (i) $u \neq v$ and (ii) there is no edge (u, v) yet. Create edge (u, v) and let $S \leftarrow S \backslash \{u, v\}$.

There are two differences between our model and that in [2]. First, we ensure that no duplicate edges or self-loops are produced. Second, we start with vertices of higher degrees, since otherwise it is possible to fail in finding distinct edges to satisfy their degree when set S becomes too small.

References

[1] C. Adjih, L. Georgiadis, P. Jacquet, and W. Szpankowski. Multicast tree structure and the power law. In *Proceedings of ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2002.

[2] W. Aiello, F. R. K. Chung, and L. Lu. A random graph model for massive graphs. In *Proceedings of ACM Symposium on Theory of Computing (STOC)*, pages 171–180, 2000.

[3] R. Albert and A. Barabási. Topology of evolving networks: Local events and universality. *Physical Review Letters*, 85:5234–5237, 2000.

[4] S. Banerjee, B. Bhattacharjee, and C. Kommareddy. Scalable application layer multicast. In *Proceedings of ACM SIGCOMM*, August 2002.

[5] A. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.

[6] A. Broido and kc claffy. Internet topology: Connectivity of IP graphs. In *Proceedings of SPIE ITCOM*, 2001.

[7] T. Bu and D. Towsley. On distinguishing between Internet power law topology generators. In *Proceedings of IEEE INFOCOM*, 2002.

[8] R. Chalmers and K. Almeroth. Modeling the branching characteristics and efficiency gains in global multicast trees. In *Proceedings of IEEE INFOCOM*, 2001.

[9] Y. Chawathe. *Scattercast: An architecture for Internet broadcast distribution as an infrastructure service*. PhD thesis, University of California, Berkeley, 2000.

[10] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. The origin of power laws in Internet topologies revisited. In *Proceedings of IEEE INFOCOM*, 2002.

[11] Y.-H. Chu, S. G. Rao, S. Seshan, and H. Zhang. Enabling conferencing applications on the Internet using an overlay multicast architecture. In *Proceedings of ACM SIGCOMM*, August 2001.

[12] Y.-H. Chu, S. G. Rao, and H. Zhang. A case for end system multicast. In *Proceedings of ACM SIGMETRICS*, June 2000.

[13] J. Chuang and M. Sirbu. Pricing multicast communications: A cost based approach. In *Proceedings of Internet Society INET*, 1998.

[14] Cooperative Association for Internet Data Analysis. The CAIDA Website. http://www.caida.org.

[15] S. Deering and D. Cheriton. Multicast routing in datagram internetworks and extended lans. In *Proceedings of ACM SIGCOMM*, 1988.

[16] P. Erdős and A. Rényi. The evolution of random graphs. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, 5:17–61, 1960.

[17] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the Internet topology. In *Proceedings of ACM SIGCOMM*, 1999.

[18] P. Francis. Yoid: Extending the multicast Internet architecture, September 1999. White paper http://www.aciri.org/yoid/.

[19] R. Govindan and H. Tangmunarunkit. Heuristics for Internet map discovery. In *Proceedings of IEEE INFOCOM*, March 2000.

[20] J. Jannotti, D. Gifford, K. Johnson, M. Kaashoek, and J. O'Toole. Overcast: Reliable multicasting with an overlay network. In *Proceedings of USENIX Symposium on Operating System Design and Implementation (OSDI)*, October 2000.

[21] C. Jin, Q. Chen, and S. Jamin. Inet: Internet topology generator. Technical Report CSE-TR-433-00, EECS Department, University of Michigan, 2000.

[22] J. Kleinberg. The small-world phenomenon: An algorithm perspective. In *Proceedings of ACM Symposium on Theory of Computing (STOC)*, pages 163–170, May 2000.

[23] A. Lakhina, J. Byers, M. Crovella, and I. Matta. On the geographic location of Internet resources. In *Proceedings of ACM SIGCOMM Internet Measurement Workshop (IMW)*, 2002.

[24] A. Lakhina, J. Byers, M. Crovella, and P. Xie. Sampling biases in IP topology measurements. In *Proceedings of IEEE INFOCOM*, April 2003.

[25] R. Mahajan, N. Spring, D. Wetherall, and T. Anderson. Inferring link weights using end-to-end measurements. In *Proceedings of ACM SIGCOMM Internet Measurement Workshop (IMW)*, November 2002.

[26] A. Medina, I. Matta, and J. Byers. On the origin of power laws in Internet topologies. *ACM SIGCOMM Computer Communication Review*, April 2000.

[27] P. V. Mieghem, G. Hooghiemstra, and R. van der Hofstad. On the efficiency of multicast. *IEEE/ACM Transactions on Networking*, 9(6):719–732, 2001.

[28] S. Milgram. The small world problem. *Psychology Today*, 2:60–67, 1967.

[29] National Lab for Applied Network Research. Global ISP interconnectivity by AS number. http://moat.nlanr.net/AS/.

[30] D. Pendarakis, S. Shi, D. Verma, and M. Waldvogel. ALMI: An application level multicast infrastructure. In *Proceedings of USENIX Symposium on Internet Technologies and Systems (USITS)*, March 2001.

[31] G. Phillips, S. Shenker, and H. Tangmunarunkit. Scaling of multicast trees: Comments on the Chuang-Sirbu scaling law. In *Proceedings of ACM SIGCOMM*, 1999.

[32] S. Ratnasamy, M. Handley, R. Karp, and S. Shenker. Application-level multicast using content-addressable networks. In *Proceedings of International Workshop on Networked Group Communication (NGC)*, November 2001.

[33] N. Spring, R. Mahajan, and D. Wetherall. Measuring isp topologies with rocketfuel. In *Proceedings of ACM SIGCOMM*, August 2002.

[34] H. Tangmunarunkit, J. Doyle, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. Does AS size determine degree in AS topology? *ACM SIGCOMM Computer Communication Review*, October 2001.

[35] H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. Network topology generators: Degree based vs. structural. In *Proceedings of ACM SIGCOMM*, 2002.

[36] Topology Project at University of Michigan. AS graph data sets. http://topology.eecs.umich.edu/data.html.

[37] USC Information Sciences Institute. Internet maps. http://www.isi.edu/div7/scan/mercator/maps.html.

[38] D. Watts and S. Strogatz. Collective dynamics of small-world networks. *Nature*, 363:202–204, 1998.

[39] B. M. Waxman. Routing of multipoint connections. *IEEE Journal on Selected Areas in Communications*, 6(9):1617–1622, 1988.

[40] W. Willinger, R. Govindan, S. Jamin, V. Paxson, and S. Shenker. Scaling phenomena in the Internet: Critically examining criticality. In *Proceedings of the National Academy of Sciences*, 2001.

[41] S. Yook, H. Jeong, and A. Barabási. Modeling the Internet's large-scale topology. *Proceedings of the National Academy of Sciences*, 2002.

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