# Small-World Internet Topologies

## *Possible Causes and Implications on Scalability of End-System Multicast*

Shudong Jin Azer Bestavros
Computer Science Department
Boston University
Boston, MA 02215
{jins,best}@cs.bu.edu

**T**echnical Report
BUCS-TR-2002-004

*January 2002*

*Abstract*—**Recent work has shown the prevalence of small-world phenomena [28] in many networks. Small-world graphs exhibit a high degree of clustering, yet have typically short path lengths between arbitrary vertices. Internet AS-level graphs have been shown to exhibit small-world behaviors [9]. In this paper, we show that *both* Internet AS-level *and* router-level graphs exhibit small-world behavior. We attribute such behavior to two possible causes–namely the high variability of vertex degree distributions (which were found to follow approximately a power law [15]) and the preference of vertices to have local connections. We show that *both* factors contribute with different relative degrees to the small-world behavior of AS-level and router-level topologies. Our findings underscore the inefficacy of the Barabasi-Albert model [6] in explaining the growth process of the Internet, and provide a basis for more promising approaches to the development of Internet topology generators. We present such a generator and show the resemblance of the synthetic graphs it generates to real Internet AS-level and router-level graphs. Using these graphs, we have examined how small-world behaviors affect the scalability of end-system multicast. Our findings indicate that lower variability of vertex degree and stronger preference for local connectivity in small-world graphs results in slower network neighborhood expansion, and in longer average path length between two arbitrary vertices, which in turn results in better scaling of end system multicast.**

## I. Introduction

Because of its phenomenal growth in size, scope, and complexity—as well as its increasingly central role in society—the Internet has become an important object of study and evaluation. Moreover, as possibly the most complex and largest artifact of human engineering that was not deliberately designed, the Internet must be approached very much like a natural or physical phenomena, whose complex emergent properties cannot be understood by simple composition of well-understood behaviors. It is for these reasons that the last few years have witnessed a surge in research that attempts to empirically identify invariants about the Internet static (e.g., topological) as well as dynamic (e.g., traffic) characteristics.

This work was partially supported by NSF research grants ANI-9986397 and ANI-0095988.

Characterizing Internet properties and behaviors, 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 the complex system that is the Internet. Not only is the characterization of Internet emergent properties important, but also explaining *how* and *why* these properties emerge is extremely valuable for many reasons. First, such an understanding would allow us to build models that could be used to generate synthetic artifacts (e.g., large graphs, traffic traces and datasets) that resemble those in the "real" Internet. Such synthetic artifacts are necessary for the evaluation of proposed protocols. Second, understanding the processes that are in play in the current Internet may give us a handle on predicting the behaviors of the Internet as it grows even larger, and more importantly, it may allow us to devise mechanisms that could alter unwanted emergent behaviors.

In this paper we focus on one aspect of Internet topology characterization, which has attracted significant attention in recent years—namely the prevalence of small-world phenomena in Internet routing maps. In Section II, we review recent findings along these lines . In Section III, using a careful analysis of real datasets, we show that small-world behaviors are not only prevalent in AS-level routing maps, but also they are prevalent in router-level maps. In Section IV, we attribute such behaviors to two possible causes—namely the well-documented, high variability of vertex degree distributions and the preference of vertices to have local connections. We show that *both* of these causes contribute with different relative degrees to the small-world behavior of AS-level and router-level topologies. This finding underscores the inefficacy of existing evocative models of Internet growth. Also, it provides a basis for more promising approaches to the development of more accurate Internet topology generators, which we discuss and evaluate in Section V.

As we postulated earlier, one of the important reasons for characterizing Internet properties—and of explaining the processes that contribute to the emergence of such properties—is the study of their impact on the effectiveness of networking protocols. To that end, in Section VI, we examine how small-world behaviors affect the scalability of end system multicast. Our findings indicate that the second cause that we believe is responsible for small-world behavior in Internet maps (namely, preference for local connectivity) results in better scalability of end-system multicast. Our findings and conclusions, summarized in Section VII, are supported both analytically and empirically.

## II. Related Work

*Features of Internet Topologies:* Recent work on Internet topology characterization has focused on two features that are distinct from early random graph models [14].

The first topological feature of recent interest, "small-world" behaviors [22] (popularly called six degrees of separation [17]), was mathematically formalized by Watts and Strogatz [28]. 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. Examples include actor collaboration networks, power grids, the World Wide Web links [1], [7], and autonomous system (AS) graphs of the Internet [9].

The second topological feature of recent interest is the skewed degree distributions of network vertices. This feature is present in paper citation databases [26], actor collaboration networks, Web links [5], and the physical connectivity of the Internet [15]. In such networks, vertices have a non-uniform probability of being connected to others, with some vertices having extremely large numbers of neighbors (e.g., popular actors and popular Web pages in the examples above). The degree distributions were often observed to follow approximately a power law.

*The BA Model:* Power-law networks are particularly emphasized by the work of Barabasi and Albert [6], [4] who explored a promising class of models that yield strict power-law degree distributions. In their model (the BA model), three generic mechanisms are defined: (1) Incremental growth, which follows from the observation that networks develop by adding new vertices or new connections. (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, 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 topology to a steady-state, in which the vertex degree distribution follows a power-law (so called scale-free).

While the BA model has been applied to Internet topology generation [21], [9], there have been some debate on its ability to explain the evolution of real networks. First, the mechanisms of the BA model are found to be inconsistent with observations from real Internet growth. For example, preferential connectivity was shown to be stronger in AS graph growth and re-wiring was shown to be an insignificant factor [11]. The shortcoming of preferential connectivity is not surprising as other findings revealed that the number of links to a Web page is not correlated to its age [2]. Second, the strict power-law degree distributions resulting from the BA model cannot be confirmed [11]. Indeed, Internet object sizes may be better captured by other distributions such as Weibull distribution [8]. This would imply that the high variability of vertex degrees in AS graphs may be the result of mechanisms [27] other than those in the BA model.

*Internet Topology Generators:* In addition to topology generators inspired by the BA model [21], [9], there are other generators that have been proposed and used to model Internet topologies. The Waxman model [29] extends the classical Erdos-Renyi model by randomly distributing vertices on a plane and creating edges by considering the distance between the vertices. The pioneering GT-ITM Internet-specific topology generator [31] uses this approach, resulting in topologies that do not exhibit power-law degree distributions. Inet [19] assigns degrees to the vertices, following a power-law distribution, and then uses a linear preferential model to realize the assigned vertex degrees. The random graph model in [3] generates degree distributions that strictly follow a power-law. However, the resulting graph is in fact 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. The model in [20] generates small-world graphs by extending the original Watts-Strogatz model to two-dimensional lattices. The probability of having an edge between two vertices is a function of the distance between the two vertices, for example inverse-square distribution. This model does not consider power-law degree distributions.

*Impact of Topology Features on Network Protocols:* Several recent studies have addressed the impact of topology on network protocols. Phillips $et\ al.$ [24] showed how the size of a multicast tree increases in tandem with the size of the multicast group, primarily under the assumption that the network neighborhood size (the number of vertices within a certain distance) increases exponentially. They provided a more rigid result which *roughly* obeys the Chuang-Sirbu[13] law for IP multicast scaling. This law stated that multicast tree size increases as $m^{0.8}$, where $m$ is the number of receivers. Chalmers and Almeroth [10] considered more realistic and different shapes of multicast tree. Radoslavov $et\ al.$ [25] considered additional topological properties—expansion, resilience, and distortion—to characterize real and synthetically-generated networks. They studied how these topological properties impact several multicast design questions, including end system multicast.

## III. Evidence

This section describes the AS-level and router-level graphs used in this paper, and provides evidence of their small-world behavior.

### A. Internet Graphs

To present evidence of small-world behaviors in Internet topologies, we use the AS-level graphs and the router-level graphs, summarized in Table I. The AS-level graphs were 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) [23]. For the purpose of our study, it suffices to use two graphs, dated September 19, 2000, and September 19, 2001, respectively. A preprocessing program (available to the public) allows the compilation of lists of AS interconnection pairs from the original routing tables. We carefully eliminated many self-edges appearing in the lists. Notice that the 2001 graph is noticeably larger than the earlier one, reflecting the growth of Internet AS community. Hereafter, these two graphs are called "AS2000 graph" and "AS2001 graph", respectively.

We use two router-level graphs available at [18]. 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 [16]. 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.

Figure 1 shows the complementary cumulative distribution function (CCDF) $1 - F(d)$ of vertex degrees for the four graphs under consideration. The CCDF quantifies the probability that a vertex has a degree larger than a certain value. A common property of these graphs is that vertex degrees exhibit high variability. The level of variability appears to be different though.

Previous work [15] showed that vertex degree distributions follow a power-law. This is confirmed by the AS graphs. With a power-law distribution, $1 - F(d) = cd^{-\alpha}$, the log-log scale plot of CCDF is a straight line, as shown in Figure 1(a-b). Using a linear regression, we estimated that for both AS graphs, the exponent $\alpha$ is close to $1.22$.

For the router-level graphs, we also estimated their power-law exponents. The values in Table I were obtained for the tail ($d > 10$). However, as evident in Figure 1(c-d), it appears that the power-law distribution for router-level graphs does not perfectly fit our empirical dataset. This is particularly the case for the Lucent graph, for which the tail of the distribution drops faster than any power law. In [8] the Weibull distribution was found to provide a good fit to many Internet object size distributions. The Weibull distribution is one of the widely used life-time distributions in reliability engineering. Its tail takes on the form $e^{-(x/\eta)^\beta}$, where $\eta$ is the scale parameter and $\beta$ is the shape parameter. Using rank regression on y-axis, we estimated that the Weibull fit to the tail (when $d > 10$) has $\beta = 0.41$. However, we also noticed that the fit really depends on where the tail starts. For example, the Weibull fit to the tail (when $d > 40$) has $\beta = 0.32$. A smaller $\beta$ value means a heavier tail. Our conclusion is that Weibull distribution does not capture well the heavier tail of the empirical data for the vertex degree distribution of router-level graphs.

### B. Small-World Phenomena

In an influential paper [28], Watts and Strogatz defined a range of graphs termed "small-world graphs". Small-world graphs are highly clustered, like regular graphs (e.g., lattices), yet 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 [28], we define $L$ as the number of edges in the shortest path between two vertices, averaged over all pairs of vertices. Also, we define the clustering coefficient $C$ as follows. Consider a vertex $v$ which has $k_v$ neighbors. Clearly, 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. We define $C$ to be the average of $C_v$ over all vertices $v$ with degree larger than one.[1]

By definition, 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. It is not difficult to see that for a random graph with $N$ vertices and with an average degree of $k$, $L_{random} \sim \ln(N)/\ln(k)$ and $C_{random} \sim k/n$.

To show that the small-world phenomenon holds for both AS-level and router-level Internet graphs, we computed their characteristic path lengths and clustering coefficients. Results are shown in Table II. For comparison purposes, we have also generated corresponding random graphs with approximately the same number of vertices and edges, and computed $L_{random}$ and $C_{random}$. Table II shows that the values of $L$ for the Internet graphs are not much larger than $L_{random}$. Indeed, in some instances, $L$ is smaller than $L_{random}$. Table II also shows that $C$ is larger than $C_{random}$ by 3-to-5 orders of magnitudes. These two observations provide clear evidence of the presence of small-world phenomenon in Internet topologies.

## IV. Possible Causes

In this section, we examine the possible causes of small-world phenomena in Internet topologies. First, we show that extremely high variability of vertex degree distributions results in a short

[1] Since $C_v$ is undefined when $k_v = 1$, this averaging used to calculate $C$ excludes vertices with only one neighbor.

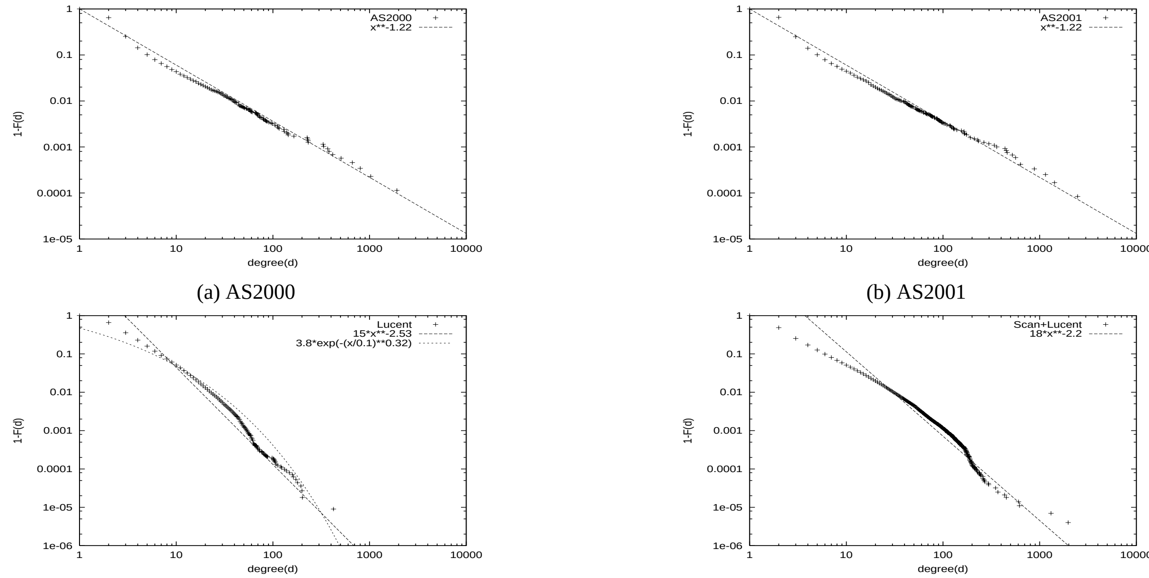


(a) AS2000 (b) AS2001

(c) Lucent (d) Scan+Lucent

Fig. 1. Vertex degree distributions of Internet AS graphs and router-level graphs.

characteristic path lengths and in a high clustering coefficients. Next, we show that such variability alone *cannot* account for all of the small-world behavior in Internet maps. We postulate that preference for local connectivity is another possible contributor.

### *A. High Variability of Vertex Degree*

To study the effects of high variability of vertex degree distributions, we generate random graphs whose vertex degree distributions 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.A-A explains how we generate random graphs with highly-variable degrees.

We generated graphs with about 10000 vertices and 100000 vertices, respectively. The average vertex degree is fixed at $4.2$. The value of $\alpha$ 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 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 $\alpha$ as shown in Figure 2. Each point represents one graph.

Figure 2 indicates that smaller $\alpha$ values result in shorter characteristic path lengths and much larger clustering coefficients. Note that a smaller $\alpha$ 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 skewed power-law vertex degree distribution is a possible cause of small-world behavior.

It has been often observed that vertex degree distributions do not fit power-law distributions well [8], [11], [27]. Nevertheless, we found that as long as the vertex degree exhibits high variability, other distributions can also lead to small-world behavior. To show this, we have generated random graphs whose vertex degrees follows a Weibull distribution. The CCDF of Weibull distribution is $e^{-(x/\eta)^{\beta}}$, where $\eta$ is the scale parameter and $\beta$ is the shape parameter. The value of $\beta$ varies from $0.2$ to $2.0$, and the value of $\eta$ 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 $\beta$ in Figure 3. Notice that the results we obtain here using Weibull distribution are similar to those we obtained using a power-law distribution. Moreover, we observe that smaller values of $\beta$ for Weibull distribution (i.e., heavier tails and higher variability of vertex degree) result in smaller $L$ and in much larger $C$. Thus, we conjecture that, the 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.

### *B. Preference for Local Connectivity*

So far, we have shown that the high variability of vertex degree distributions results in small-world behavior. But, are there other causes? In this section, we show that the answer to this question is affirmative.

Specifically, to answer the above question, we generate graphs whose vertex degree distribution follows *exactly* the same distribution of the real Internet graphs. However, the edges were created randomly. In this way, we *preserve* the high variability of vertex degree, but *destroy* other topological properties that may exist in real Internet graphs. We call these synthetic graphs *"randomized" Internet graphs*. Appendix A.A-A describes the algorithm we used to generate these randomized graphs.

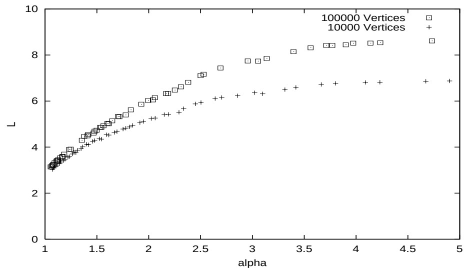


(a) Characteristic path length

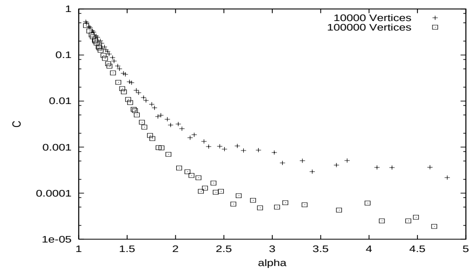


(b) Clustering coefficient

Fig. 2. Power Law degree distribution results in small-world behavior.

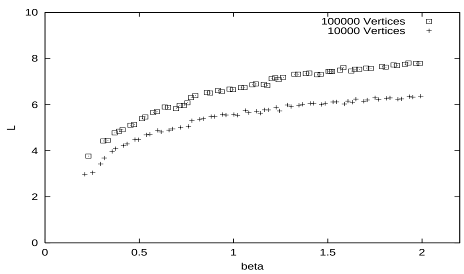


(a) Characteristic path length

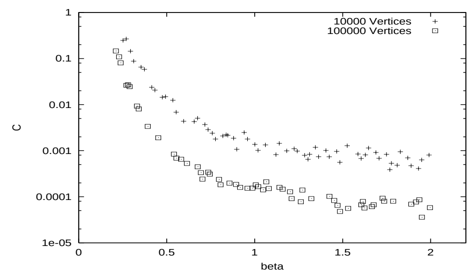


(b) Clustering coefficient

Fig. 3. Weibull degree distribution results in small-world behavior.

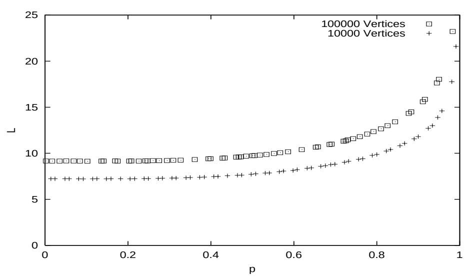


(a) Characteristic path length

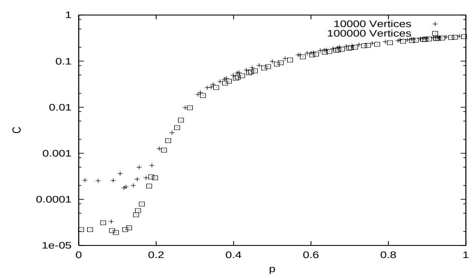


(b) Clustering coefficient

Fig. 4. Preference for local connectivity results in small-world behavior.

For each "real" Internet graph, we generate a number of randomized instances, and for each we compute their $L$ and $C$ values, averaged over all randomized instances. The results we obtained are tabulated in Table III. Comparing these results with those of the real graphs in Table II, we can make the following observations. First, for AS-level graphs, the $L$ values are very close to each other and the $C$ values differ only by a small factor (although absolute differences are large). Second, for 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 as follows. First, clearly there are other causes that contribute to the small-world behavior of Internet topologies. Second, it appears that these "other causes" are more pronounced when the variability of vertex degree distributions is moderate but not extreme. In other words, when the variability of vertex degree distributions is extremely high, which is the case for AS-level graphs, the effect of these other causes is overshadowed. When the variability of vertex degree distributions is only moderate, which is the case for router-level graphs, the effect of these other causes is evident.

So, what are those other causes of small-world behavior in Internet graphs? In attempting to answer this question, we do not intend to provide a complete and exclusive explanation of small-world phenomena, but to identify one plausible explanation—namely, *the preference for local connectivity in the Internet.* Indeed, this possible explanation was originally suggested by Watts and Strogatz [28], who found that if only a portion of the edges of a regular lattice are reconnected randomly, the resulting graph would exhibit small-world behavior. By doing so, the clustering coefficient remains very high due to local connectivity, but the characteristic path becomes closer to that of random graphs due to long-range (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 with the same distribution as that of a random graph with the same number of vertices and edges. Specifically, the distribution has an exponentially-decayed 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 connections of $v$ are random. Here, $0 \leq p \leq 1$ is the probability of local connectivity, which we call the "local probability". We varied $p$ to generate many graph instances and computed their $L$ values and $C$ values. The results are plotted in Figure 4, where each point represents one graph instance.

The results in Figure 4 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 coefficient increases drastically when $p$ is small or moderate. Overall, there is a wide regime for $p$ that yields characteristic signatures of small-world graphs.

### C. Discussion

*How does the high variability of vertex degree distributions result in small-world behaviors?* 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 large degree. 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 overall average clustering coefficient $C$. Thus, $C$ grows with the variability of vertex degree. Also, notice that with highly-variable vertex degrees, 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 through 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. As a result, there is a non-negligible probability that a triangle will form among these vertices, resulting in a higher clustering coefficient. Meanwhile, since there are still many long-range connections, it is easy to find a short path between two randomly-chosen vertices. It is those vertices with long-range connections that serve as good navigators.

From above discussions, both high variability of vertex degree distributions and the preference for local connectivity appear to be possible causes of small-world behavior. Moreover, both of these causes are plausible. 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 [6], or they may exist due to other reasons [27]. Preference for local connectivity may be explained for both router-level and AS-level topologies as follows. At the router-level, links are created by considering physical distances. At the AS-level, ISPs may form cliques in which the logical distance is shorter.

Notice that our findings imply the failure of the BA model as an explanation of Internet growth and evolution [6], as well as the adequacy of topology generators based thereupon [21], [9]. Although several previous studies [11], [27], [8], [2], [30] have casted doubts on the adequacy of the BA model, none has examined the causes of small-world phenomena as evidence. Specifically, the BA 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 II and Table III provide the evidence. When the power-law exponent departs much from unity, the BA model fails to generate small-world graphs (as observed for router-level graphs).

## V. Synthesis

In this section, we describe an Internet topology model, and show that it generates a range of small-world graphs. We also use this generator to create synthetic graphs with both characteristic path lengths and clustering coefficients closely matching those of real AS-level and router-level Internet graphs.

### A. Topology Generation

We use the following model to generate a graph, given the power-law exponent $\alpha$ of vertex degree distribution, and the local probability $p$.

(1) Randomly place $N$ vertices on a plane. A degree $d_v$ is assigned to each vertex $v$, $1 \leq v \leq N$, such that $d_v$ follows the power-law.

(2) Create local connections among the vertices. Connect each vertex $v$ to its nearest $pd_v$ neighbors. To be precise, $pd_v$ is rounded down to $\lfloor pd_v \rfloor$ or rounded up to $\lceil pd_v \rceil$, in a probabilistic way[2].

(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 graph model described in Appendix A.A-A.

We have used this model to generate graphs which exhibits a wide range of small-world behaviors. For example, we generated graphs with 10000 vertices and with average degree 4.2. We varied parameters $\alpha$ from 1.05 to nearly 8.0 and varied $p$ between 0 and 1. For each pair of $\alpha$ and $p$, we generated ten graphs, computed their average $L$ value and $C$ value, and plotted them in Figure 5. We observe that when $\alpha$ is small or $p$ is moderate but not too high, both the characteristic path length and the clustering coefficient satisfy the requirements for small-world graphs.

[2] If a real number $r$ is closer to its ceil, 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 $r$


(a) Characteristic path length

(b) Clustering coefficient

Fig. 5. A range of small-world graphs are generated with different choices of $\alpha$ and $p$.

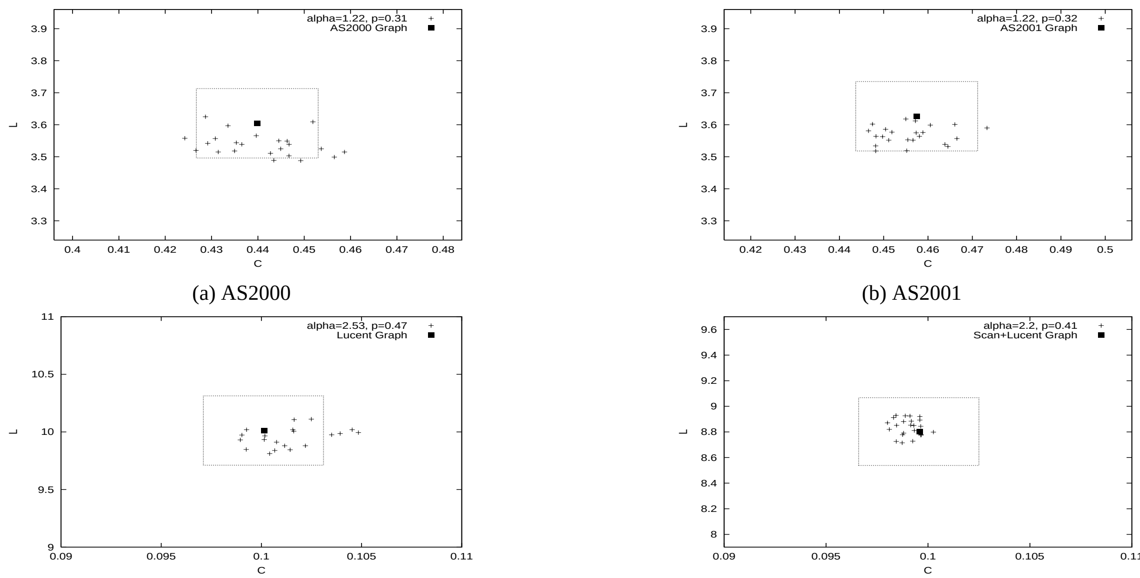


(a) AS2000

(b) AS2001

(c) Lucent

(b) Scan+Lucent

Fig. 6. Synthetic graphs have both characteristic path length and clustering coefficient very close to those of the real Internet graphs.

## *B. Synthetic Internet Graphs*

We have also used the above model to synthesize Internet graphs which are similar to the AS-level and router-level graphs we have empirically characterized. To do so, we first estimate the power-law exponent of vertex degree distribution of the real Internet graph. Next, we find a value of $p$ that yields a synthetic graph with $L$ and $C$ values that are close to those of the real Internet graph. For example, for the AS2001 graph, we have $\alpha \approx 1.22$ from Table I and a corresponding $p \approx 0.31$. For the Lucent graph, we have $\alpha \approx 2.53$ from Table I and a corresponding $p \approx 0.47$. We used these parameter settings to create a number of synthetic graphs. The resulting $L$ values and $C$ values are reported in Figure 6, where one point represents one synthetic graph. The values of $L$ and $C$ for the real Internet graph we aim to synthetically replicate are shown by a solid point in the center of the plot. The box in the plot shows a range where both the $L$ and $C$ values are within $3\%$ of the target. For the AS-level graph, the $L$ value of synthetic graph is slightly smaller, but the difference is often less than $3\%$. For the Lucent graph, occasionally the $C$ value of the synthetic graph was found to be slightly larger, but the difference is still insignificant. For the Scan+Lucent graph, the fit is the best.

To summarize, by choosing appropriate settings for $\alpha$ and $p$, our algorithm generates graphs that very much resemble the real AS-level and router-level graphs in terms of their small-world characteristics.

## *C. Discussion*

Three related issues are discussed in this subsection. The first issue is whether our model is amenable to incremental growth, i.e., adding new vertices or edges. When a new vertex $v$ is added into a graph, it is attached to an existing vertex $u$ according to some preferential model. This attachment is either remote, in which case $v$ is allocated at a random point on the plane, or local, in which case $v$ is allocated within a certain distance of $u$. When a new edge is added, an existing vertex $v$ is chosen according to some preferential model. The edge is either remote, in which case another vertex is also chosen according to the preferential model, or local in which case another vertex is chosen within a certain distance from $v$.

The second issue is whether the probability of having an edge between two vertices can be a function of their distance? 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 $d_{u,v}^{-r}$, where $d_{u,v}$ is the distance between $u$ and $v$ and $r$ is a positive constant. It was originally used by Kleinberg [20] to generalize the Watts-Strogatz model. The key is the choice of $r$. If $r$ is larger, $d_{u,v}^{-r}$ decays faster and the edge tends to be local; but $r$ cannot be too large since otherwise the lack of long-range connections would lead to an excessively large $L$ value. If $r$ is smaller, $d_{u,v}^{-r}$ decays slowly and long-range connections become frequent; but $r$ cannot be too small since otherwise the graphs would be close to random (i.e., the value of $C$ would be excessively small). In [20], it was found that $r = 2$ (inverse-square distribution) is the optimal choice for two-dimensional lattices.

The third issue is how to generate a graph exhibiting skewed geographical density. For example, Govindan and Tangmunarunkit [16] identified the geographical location of the routers in the Scan+Lucent map. They found different patterns in the density map. To augment our model with such functionality, a natural approach is distributing the vertices on a plane non-uniformly.

## VI. Implications on End-System Multicast

In this section, we study the scalability of end system multicast in light of the small-world behavior of Internet topologies. We show that in small-world graphs, neighborhood expansion exhibits various patterns. Using a neighborhood expansion function, we compute a theoretical bound on the scalability of end system multicast. These results are validated by simulations using synthetic graphs and real Internet graphs.

### A. Neighborhood Expansion in Small-world Graphs

The neighborhood expansion function $E(d)$ for a graph is defined as the average fraction of vertices reachable in $d$ hops, starting from an arbitrary vertex.

To characterize $E(d)$ for graphs exhibiting small-world behavior resulting from highly-variable vertex degree distributions *only*, we generated synthetic graphs of varying power-law exponent ($1.11 \le \alpha \le 10.0$), while keeping the local probability $p$ set to zero. The generated graphs have approximately 30000 vertices and an average degree of 4.0. Their $L$ values and $C$ values were computed and reported in Table IV. We also computed their neighborhood expansion function, which are plotted in Figure 7. We computed $E(d)$ by finding the fraction of vertices reachable in $d$ hops from every vertex, and then calculating $E(d)$ to be the average over all vertices.

Figure 7 shows that $E(d)$ varies significantly with $\alpha$. A higher variability of vertex degree distributions results in fast neighborhood expansion. Also, when $\alpha$ is large (e.g., $\alpha = 10.0$), $E(d)$ appears to increase *exponentially* until it approaches saturation, which simply means that the entire graph has been covered. The exponential nature of $E(d)$ is clearly shown on the semi-log plot in Figure 7(b), where a straight line means an exponential increase. Notice, however, when $\alpha$ is small, e.g., $\alpha = 1.11$, $1.43$, or $2.0$, a power-law seems to provide a better fit for $E(d)$.

To characterize $E(d)$ for graphs exhibiting small-world behavior resulting primarily from preference for local connectivity, we generated synthetic graphs of varying local probability $p$, while fixing $\alpha = 5.0$. Thus, the variability of the vertex degree is low, but the graphs exhibit small-world behavior due to preference for local connectivity. Table V gives the $L$ and $C$ values of these graphs. As before, we compute their neighborhood expansion functions, which are shown in Figure 8. An interesting observation from this figure is that the growth of $E(d)$ gradually shifts from exponential (when $p = 0$) to approximately power-law (when $p = 1$). This can be explained by noting that when $p$ is close to zero, there is strong randomness in the graphs, resulting in exponential growth of the neighborhood expansion function. On the other hand, when $p$ is close to unity, all connections are made locally, resulting in fairly slow neighborhood expansion. Note that in regular two-dimensional grids or meshes, $E(d)$ grows quadratically. In our case, since vertex degree still exhibits some variability, even when $p = 1$, $E(d)$ still increases faster than $d^2$.

To summarize, from both experiments we have conducted, it is evident that the neighborhood expansion function is largely affected by small-world behavior. How does such such network neighborhood expansion affect the scalability of end system multicast is exploited next.

### B. Theoretical Results

We consider the following end system multicast scheme, originally described in [12]. There are one sender and $n$ 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 the complete virtual graph. The scalability of end system multicast using such a tree reduces to how the tree size increases with $n$. Our main theoretical results characterizing the asymptotic growth of the tree size as a function of $n$ follow. (Note: By asymptotic, we mean that $n \gg 1$ but $n \ll N$, where $N$ is the number of vertices in the network.)

*Exponential Neighborhood Expansion Implications:* If the neighborhood expansion function is exponential and homogeneous for all vertices, then the tree size increases asymptotically at least as fast as

$$n\left(1 - \frac{\ln(n)}{\ln(N)}\right) \tag{1}$$

*Power-Law Neighborhood Expansion Implications:* If the network neighborhood expansion function is a power-law with ex-

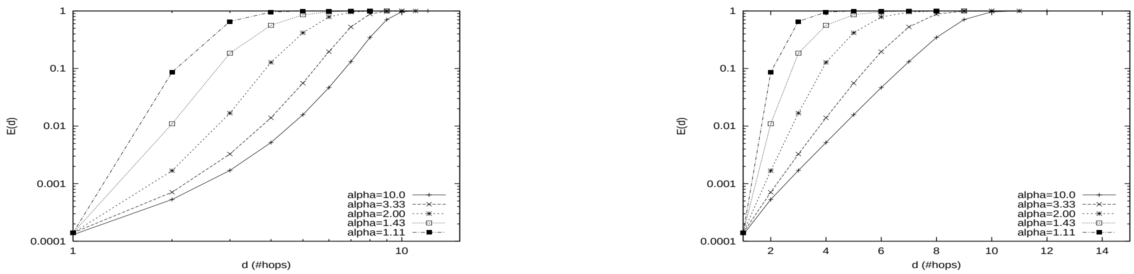


(a) Log-log scale (b) Linear-log scale

Fig. 7. Neighborhood expansion varying with $\alpha$.

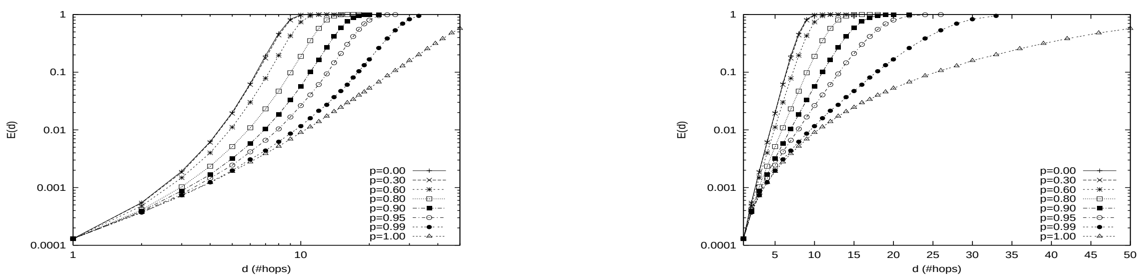


(a) Log-log scale (b) Linear-log scale

Fig. 8. Neighborhood expansion varying with $p$.

ponent $H$ and is homogeneous for all vertices, then the tree size increases asymptotically at least as fast as

$$n^{1-1/H} \tag{2}$$

The above results are obtained by reducing the problem to another problem: if there are $n$ independent copies of an object in a network, what is the expected distance to the nearest copy? Notice that the expected size of the minimum spanning tree should not be smaller than $n$ *times this expected distance* since $n$ edges in the virtual complete graph appear in the minimum spanning tree. In Appendix A.A-B, we show that, asymptotically, if the neighborhood expansion function is exponential, then this expected distance is proportional to $(1 - \frac{\ln(n)}{\ln(N)})$. And, if the network neighborhood expansion function is a power-law with exponent $H$, then this expected distance is proportional to $n^{-1/H}$.

These results suggest that with exponential neighborhood expansion, the scalability of end-system multicast *does not* depend on the small-world behavior ($L$ or $C$), given the constant $N$. However, with power-law expansion, the scalability of end system multicast *does* depend on the $L$ value of small-world graphs[3].

The above results only show *asymptotic* behavior. In practice, network size and multicast group size are bounded. Also, these results hold only under the ideal assumption of *homogeneous neighborhood expansion*. In practice, neighborhood sizes of different vertices may vary. Thus, the question is whether these theoretical results apply to real Internet topologies. We address this question next by presenting simulation results over synthetic as well as real Internet topologies that exhibit small-world properties.

### C. Simulation Results

In the first set of experiments, we simulated an end-system multicast system over synthetic graphs generated using different values of $\alpha$, with the values of $L$ and $C$ given in Table IV and the neighborhood expansion functions shown by Figure 7. Given a graph, we randomly choose $n$ receivers and one sender, generate the complete virtual graph, find a minimum spanning tree, and calculate its size. Using different random seeds, we repeat the above process and compute an average tree size. This experiment is repeated by varying $n$. Figure 9 shows the results.

Figure 9 shows that smaller values of $\alpha$ (i.e., larger $L$ and smaller $C$) result in a better scaling behavior for end-system multicast. The results also show a large deviation from the bound of equation (1), suggesting that this asymptotic behavior (under an exponential neighborhood expansion) does not hold here. On the other hand, the asymptotic bound $n^{1-1/H}$ of equation (2) seems to apply here. Several reasons can account for this. First, the network neighborhood expansion for the graphs we considered may well follow a power-law, especially when $\alpha$ is small. Second, the graph size (30000 vertices) may have been too small to allow exhibition of asymptotic behaviors, especially for large values of $n$. Third, the graphs we used have a higher vertex degree variability, which results in a heterogeneous network expansion function, thus violating the homogeneity as-

[3]Why? Let $D$ be the diameter of the network and assume power-law expansion has $E(x) = (x/D)^H$. We can compute the characteristic path length $L \approx \int_1^D x dE(x) \approx \frac{H}{H+1} D$. Note, $D^H = N$, so $L \approx \frac{H}{H+1} N^{1/H}$. In other words, $L$ is related to $H$, given the constant $N$.

sumption of network neighborhood expansion, which is necessary for the asymptotic bound given in equation (1) to hold.

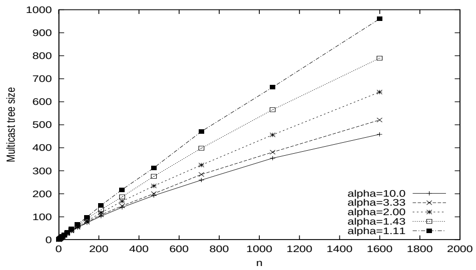


Fig. 9. Scaling of end system multicast varying with $\alpha$.

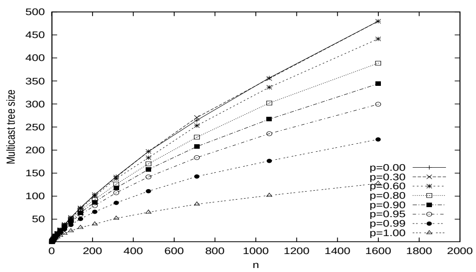


Fig. 10. Scaling of end system multicast varying with $p$.

In the second set of experiments, we simulated end system multicast using synthetic graphs generated using varying values of $p$. The $L$ and $C$ values for these graphs are given in Table V, and their neighborhood expansion functions are shown by Figure 8. Figure 10 shows the simulation results we obtained. It shows that larger values of $p$ (i.e., larger $L$ and larger $C$) result in better scaling behavior. By comparing these results to the $n^{1-1/H}$ bound of equation (2), where $H$ is estimated from Figure 8, we find that our simulation results are slightly higher. This is expected since $n^{1-1/H}$ is not a tight bound.

To summarize, it appears from these two sets of experiments that the value of $L$ of small-world graphs dictates the scaling behavior of end system multicast. For a network with given size, the larger the $L$, the better the scalability of end system multicast. On the other hand, the $C$ value appears to be less important as evident from both our analytical and experimental results.

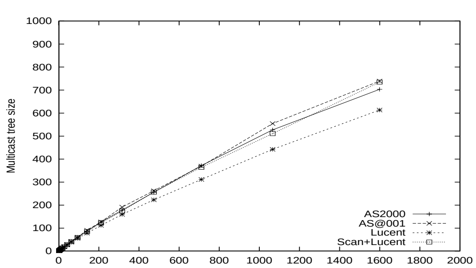


Fig. 11. Scaling of end system multicast in Internet graphs.

In our last set of experiments, we simulated end-system multicast using real Internet graphs. Results of these experiments are shown in Figure 11. To better understand these results, we also show the neighborhood expansion function of these graphs in Figure 12. For both the AS2000 and AS2001 graphs, $\alpha = 1.22$ is small, resulting in fast neighborhood expansion and smaller $L$ values. These small $L$ values lead to faster increases in multicast tree size. However, since these two graphs are smaller (8742 and 11927 vertices, respectively), it is easier to reach saturation, at which point the asymptotic behavior does not hold. For the Lucent and Scan+Lucent graphs, we observed that the former has a slower neighborhood expansion and a larger $L$ value, even though its size (112669) is smaller than the size (282672) of the Scan+Lucent graph. As a result, end system multicast displays better scalability in the Lucent graph. Notice that these two graphs have approximately equal $C$ values, but their end system multicast scaling behavior is quite different. This confirms our previous conclusion that the $C$ value of small-world graph is less important.

## VII. Conclusion

Accurate characterization of the emergent topological properties of the Internet and better understanding of the underlying processes that yield these characteristics are crucial for proper evaluation of network protocols and systems. In that vein, recent work [9] has shown the prevalence of small-world behaviors [28] in AS-level routing maps. In this paper, we extend these findings to router-level maps as well. More importantly, we attribute small-world behaviors in Internet maps to two possible causes—namely the high variability of vertex degree distributions and the preference of vertices to have local connections. We quantified the relative contributions of both of these causes to the observed small-world behavior of AS-level and router-level topologies. In addition to establishing the inefficacy of existing models [6] that attempt to explain the growth process of the Internet, we describe a more promising approach to the generation of synthetic Internet topologies and show the resemblance of the synthetic graphs it generates to real Internet AS-level and router-level graphs.

To underline the importance of being able to quantify the relative strengths of both causes of small-world behavior, we considered the implications of small-world behavior on the scalability of end system multicast. Our analytical and experimental findings indicate that the small-world behavior in Internet graphs affects the scalability of end-system multicast, with the characteristic path lengths of such graphs being the dominant factor.

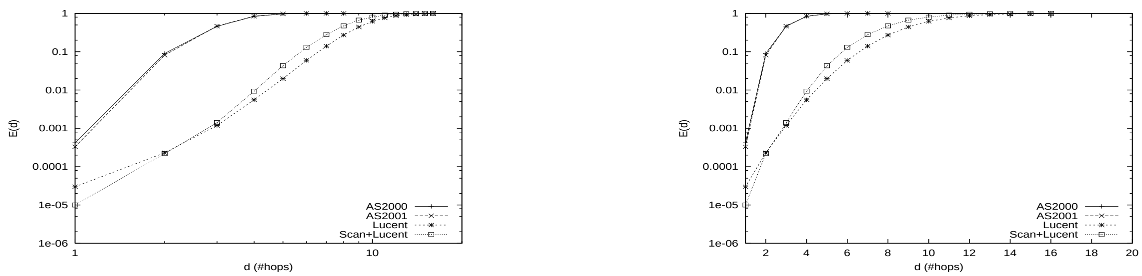


(a) Log-log scale

(b) Linear-log scale

Fig. 12. Neighborhood expansion in Internet graphs.

## REFERENCES


[1] L. Adamic. The small world Web. In *Proceedings of ECDL*, pages 443–452, 1999.

[2] L. Adamic. Technical comments: Power-law distribution of the World-Wide Web. Science, 287, 2000.

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

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

[5] R. Albert, H. Jeong, and A. Barabasi. Diameter of the World-Wide Web. *Nature*, 401:130–131, 1999.

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

[7] A. Broder, R. Kumar, F. Maghoul, P. Raghavan, S. Rajagopalan, R. Stata, A. Tomkins, and J. Wiener. Graphs structures in the Web. *Computer Networks*, 33:309–320, June 2000.

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

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

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

[11] 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 INFOCOM*, 2002.

[12] Y.-H. Chu, S. G. Rao, and H. Zhang. A case for end system multicast. In *Proceedings of SIGMETRICS*, pages 1–12, 2000.

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

[14] P. Erdos and A. R'enyi. The evolution of random graphs. *Pupl. Math. Inst. Hungar. Acad. Sci.*, 7:17–61, 1960.

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

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

[17] J. Guare. *Six Degrees of Separation: A Play*. Vintage Books, New York, 1990.

[18] I. S. Institute. Internet maps. http://www.isi.edu/div7/scan/mercator/maps.html.

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

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

[21] A. Medina, I. Matta, and J. Byers. On the origin of power laws in internet topologies. *ACM SIGCOMM computer communication review*, April 2000.

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

[23] National Laboratory for Applied Network Research. Global ISP interconnectivity by as number. http://moat.nlanr.net/as/.

[24] G. Phillips, S. Shenker, and H. Tangmunarunkit. Scaling of multicast trees: Comments on the Chuang-Sirbu scaling law. In *Proceedings of SIGCOMM*, pages 41–51, 1999.

[25] P. Radoslavov, H. Tangmunarunkit, H. Yu, R. Govindan, S. Shenker, and D. Estrin. On characterizing network topologies and analyzing their impact on protocol design. Technical Report 00-731, Dept. of CS, University of Southern California, February 2000.

[26] S. Redner. How popular is your paper? An empirical study of the citation distribution. *Eur. Phys. J. B*, 4:131–134, 1998.

[27] 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.

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

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

[30] 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.

[31] 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.

# APPENDIX

### A. 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 $d_i$, $1 \leq i \leq N$. Our technique modifies the original model in [3] and generates random graphs as follows.

*(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 [3]. 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.

### B. End-System Multicast Scalability Bounds

We show the expected distance from the nearest copy, given that there are $n$ independent copies in the network, under a network neighborhood expansion function that is (1) exponential or (2) power-law.

First consider the general case, $E(x)$ is defined on $a < x < b$ in continuous form. Let us consider the probability that an arbitrary copy is not within distance $x$. By definition, this probability is $1 - E(x)$. Since there are $n$ independent copies, the probability that none of them is within distance $x$ is $(1 - E(x))^n$. Let $F(x) = 1 - (1 - E(x))^n$. The probability density function $f(x)$, i.e., the probability that the nearest copy is at distance $x$, is computed as $\frac{dF(x)}{dx}$. Let $g(n)$ denote the expected distance of the nearest copy, which is computed as follows

$$\begin{aligned} g(n) &= \int_a^b x f(x) dx \\ &= \int_a^b x dF(x) \qquad \text{Let } y = F(x) \\ &= \int_0^1 F^{-1}(y) dy, \end{aligned} \tag{3}$$

where $F^{-1}(.)$ is the inverse of $F(.)$. We then consider exponential expansion and power-law expansion separately.

#### B.1 Exponential Neighborhood Expansion

Let neighborhood expansion function $E(x) = k^{x-D}$, $0 \leq x \leq D$, where $D$ is called the diameter of the network. $F(x) = 1 - (1 - E(x))^n = 1 - (1 - k^{x-D})^n$. Inverse function $F^{-1}(y) = D + \frac{\ln(1-(1-y)^{1/n})}{\ln(k)}$. From equation (3), $g(n)$ is computed as follows

$$\begin{aligned} g(n) &= \int_0^1 F^{-1}(y) dy \\ &= D + \frac{1}{\ln(k)} \int_0^1 \ln(1 - (1-y)^{1/n}) dy \\ &= D + \frac{1}{\ln(k)} \int_0^1 \ln(1 - y^{1/n}) dy. \end{aligned}$$

Interestingly, $-\int_0^1 \ln(1 - y^{1/n}) dy = \sum_{i=1}^n \frac{1}{i}$, the harmonic number which is equal to $\ln(n) + 0.5772156 \cdots + \frac{1}{2n}$. This term is asymptotically close to $\ln(n)$. Therefore, $g(n) \approx D - \frac{\ln(n)}{\ln(k)}$. Notice, $\frac{\ln(N)}{\ln(k)} = D$. Therefore, $g(n) \approx D(1 - \frac{\ln(n)}{\ln(N)})$, which grows asymptotically as $1 - \frac{\ln(n)}{\ln(N)}$.

The exact computation of term $-\int_0^1 \ln(1 - y^{1/n}) dy$ is as follows. Notice $\ln(1 + x)$ has series expansion $x - \frac{x^2}{2} + \frac{x^3}{3} - \cdots$. Therefore,

$$\begin{aligned} -\int_0^1 \ln(1 - y^{1/n}) dy &= \sum_{i=1}^{\infty} \int_0^1 \frac{y^{i/n}}{i} dy \\ &= \sum_{i=1}^{\infty} \frac{n}{(n+i)i} \\ &= \sum_{i=1}^{n} \frac{1}{i} \end{aligned}$$

#### B.2 Power Law Neighborhood Expansion

Let neighborhood expansion function $E(x) = (x/D)^H$, $1 \leq x \leq D$, where $D$ is called the diameter of the network. $F(x) = 1 - (1 - E(x))^n = 1 - (1 - (x/D)^H)^n$. Inverse function $F^{-1}(y) = D(1 - (1-y)^{1/n})^{1/H}$. From equation (3), $g(n)$ is computed as follows

$$\begin{aligned} g(n) &= \int_0^1 F^{-1}(y) dy \\ &= D \int_0^1 (1 - (1-y)^{1/n})^{1/H} dy \\ &= D \int_0^1 (1 - y^{1/n})^{1/H} dy. \end{aligned}$$

The computation of $g(n)$ is complicated (more details are in the next paragraph). It is easier to numerically solve $\frac{g(n)}{D}$ for different values of $n$. We have done so using different choices of $H$. The results are plotted in Figure 13. Interestingly, the results show that $\frac{g(n)}{D}$ match $n^{-1/H}$ very well for various choices of $H$ and not too small $n$. Notice that, the numerical lines are parallel to the corresponding $n^{-1/H}$ lines. It means their slight difference does not change the scaling behavior.

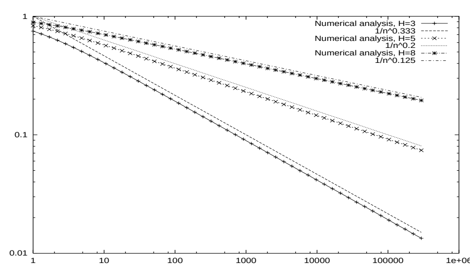


Fig. 13. Numerical results match $n^{-1/H}$ well.

The computation of $\frac{g(n)}{D}$ is the follows

$$\begin{aligned} \frac{g(n)}{D} &= \int_0^1 (1 - y^{1/n})^{1/H} dy \qquad \text{Let } x = 1 - y^{1/n} \\ &= n \int_0^1 x^{1/H} (1-x)^{n-1} dx \\ &= nB(1 + \frac{1}{H}, n) \\ &= n \frac{\Gamma(1 + \frac{1}{H}) \Gamma(n)}{\Gamma(n + 1 + \frac{1}{H})}, \end{aligned}$$

where $\Gamma(a)$ is the Gamma function defined as $\int_0^\infty x^{a-1} e^{-x} dx$, and $B(a, b)$ is the Beta function equal to $\frac{\Gamma(a)\Gamma(b)}{\Gamma(a+b)}$. It appears, due to the factorial nature of Gamma function, $\frac{\Gamma(n)}{\Gamma(n+1+\frac{1}{H})}$ is roughly proportional to $n^{-1-1/H}$ when $n$ is large. In addition, $\Gamma(1 + 1/H)$ is a constant, corresponding to the displacement between the numerical lines $\frac{g(n)}{D}$ and $n^{-1/H}$ lines in Figure 13.

TABLE I

INTERNET GRAPHS USED IN THIS PAPER.

| Graphs | AS2000 | AS2001 | Lucent | Scan+Lucent |
|---|---|---|---|---|
| Number of vertices | 8742 | 11927 | 112669 | 282672 |
| Average degree | 4.06 | 4.19 | 3.21 | 3.15 |
| Maximum degree | 1918 | 2467 | 423 | 1973 |
| Fit distributions | Power law ($\alpha \approx 1.22$) | Power law ($\alpha \approx 1.22$) | Power law ($\alpha \approx 2.53$), Weibull ($\beta \approx 0.32$) | Power law ($\alpha \approx 2.20$) |

TABLE II

CHARACTERISTIC PATH LENGTH AND CLUSTERING COEFFICIENT OF THE GRAPHS.

| Graphs | AS2000 | AS2001 | Lucent | Scan+Lucent |
|---|---|---|---|---|
| $L$ | 3.655 | 3.627 | 10.02 | 8.803 |
| $C$ | 0.4399 | 0.4578 | 0.1001 | 0.0996 |
| $L_{random}$ | 6.721 | 6.797 | 10.49 | 11.38 |
| $C_{random}$ | 0.00035 | 0.00025 | 0.000022 | 0.000007 |

TABLE III

CHARACTERISTIC PATH LENGTH AND CLUSTERING COEFFICIENT OF RANDOMIZED GRAPHS.

| Graphs | Randomized AS2000 | Randomized AS2001 | Randomized Lucent | Randomized Scan+Lucent |
|---|---|---|---|---|
| $L$ | 3.464 | 3.411 | 6.944 | 5.971 |
| $C$ | 0.2417 | 0.2653 | 0.00028 | 0.00076 |

TABLE IV

CHARACTERISTIC PATH LENGTH AND CLUSTERING COEFFICIENT OF GENERATED GRAPHS VARYING WITH $\alpha$.

| Graphs | $\alpha = 10.0$ | $\alpha = 3.33$ | $\alpha = 2.00$ | $\alpha = 1.43$ | $\alpha = 1.11$ |
|---|---|---|---|---|---|
| $L$ | 8.775 | 7.338 | 5.724 | 4.453 | 3.278 |
| $C$ | 0.00018 | 0.0010 | 0.0119 | 0.0312 | 0.3602 |

TABLE V

CHARACTERISTIC PATH LENGTH AND CLUSTERING COEFFICIENT OF GENERATED GRAPHS VARYING WITH $p$.

| Graphs | $p = 0.0$ | $p = 0.3$ | $p = 0.6$ | $p = 0.8$ | $p = 0.9$ | $p = 0.95$ | $p = 0.99$ | $p = 1.00$ |
|---|---|---|---|---|---|---|---|---|
| $L$ | 8.422 | 8.482 | 9.571 | 11.98 | 14.72 | 17.94 | 25.48 | 45.82 |
| $C$ | 0.00010 | 0.01868 | 0.1668 | 0.2702 | 0.3168 | 0.3327 | 0.3480 | 0.3525 |