Understanding and Modeling the Internet Topology: Economics and Evolution Perspective

Xiaoming Wang, *Student Member, IEEE*, and Dmitri Loguinov, *Senior Member, IEEE*

***Abstract*—In this paper, we seek to understand the intrinsic reasons for the well-known phenomenon of heavy-tailed degree in the Internet AS graph and argue that in contrast to traditional models based on preferential attachment and centralized optimization, the Pareto degree of the Internet can be explained by the evolution of wealth associated with each ISP. The proposed topology model utilizes a simple multiplicative stochastic process that determines each ISP's wealth at different points in time and several "maintenance" rules that keep the degree of each node proportional to its wealth. Actual link formation is determined in a decentralized fashion based on random walks, where each ISP individually decides when and how to increase its degree. Simulations show that the proposed model, which we call *Wealth-based Internet Topology* (WIT), produces scale-free random graphs with tunable exponent α and high clustering coefficients (between 0.35 and 0.5) that stay invariant as the size of the graph increases. This evolution closely mimics that of the Internet observed since 1997.**

***Index Terms*—Autonomous systems, clustering coefficient, degree distribution, Internet topology, random walk, wealth evolution.**

I. Introduction

RECENT studies show that real-life large-scale networks not only exhibit power-law degree distributions, but are highly clustered. Thus, a significant effort has recently focused on developing graph generators that are capable of constructing random networks with power-law degree distributions [1]–[3], [7], [12], [15], [23], [24], [33], [48], [49], [52], [59] and high clustering [7], [25]. Among the previous approaches, preferential attachment [7], [18], [60] and optimized-based construction [15] have become the two major paradigms for explaining the Internet topology. The former theory relies on the principle that each joining node attaches its links to existing nodes with a probability proportional to their current degree, *without any regard for the existing link structure.* The latter theory models node join as an optimization problem and argues that each joining ISP aims to solve a certain tradeoff between the benefit of improved connectivity and the cost of adding new links. This optimization is run over the existing graph and takes the current connectivity into account.

As we discuss next, the existing evolution theories exhibit several limitations in the context of the Internet AS-graph. While acceptable in certain cases (such as social networks), preferential attachment [7], [18], [60] is usually too restrictive to realistically model the Internet graph as it bases link formation *solely* on the degrees of existing nodes[1] and places too much weight on ISP "popularity." From the practical perspective, it is clear that such complex factors as geographic location, technical feasibility, business strategy, existing connectivity, and various economic considerations contribute to the evolution of each network rather than the mere size of other networks. Optimization-based topology models [15] are viable alternatives to preferential attachment that capture more diverse factors related to ISP peering; however, the lack of *mutuality* (i.e., a joining provider cannot attach to an ISP that does not wish to peer with it) and absence of *economic basis* for link formation (e.g., a joining network operator would not attach to an ISP close to bankruptcy, regardless of how well-connected the latter one is) make them potentially unrealistic as well.

In addition, both preferential attachment and optimization-based construction depend on the *global* knowledge of the system and always create random graphs using centralized information. While this is certainly not a problem during simulations (i.e., most generators are centralized), we argue that any theory that relies on global knowledge inherently fails to *explain* how the Internet could have reached its current stage given the fact that no single ISP has complete information about the AS graph. In preferential attachment, it is hard to conceive that new ISPs will test the probability p_i of connecting to each existing ISP i and then select the peering point exactly according to the ratio of degree d_i to the global sum $\sum_{k=1}^{n} d_k$, where n is the number of nodes in the system. As for the theory of optimization-based tradeoff, the algorithm requires complete information about the *structure* of the graph (not just the degree of each ISP) and burdens each new node with an optimization process with complexity $\Theta(n^2)$, which is hardly possible in simulations.

Recent studies [21] have revealed that there are missing links in the data sets based on BGP routing tables and thus the graph induced from these sources is not representative of the true Internet structure. While many measurement efforts are under way to address this problem of completeness, no consensus is reached yet about how many links are still missing from the existing measurement data and how far the current view of the Internet deviates from its true view. Therefore, a graph model must not only represent the current view, but also be compatible with future discoveries of new features of the Internet, e.g.,

[1]We use terms node, user, and ISP interchangeably.

Manuscript received March 30, 2006; revised October 03, 2007; October 15, 2008; and March 20, 2009; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor D. Raz. First published September 29, 2009; current version published February 18, 2010. This work was supported by the National Science Foundation (NSF) under Grants CCR-0306246, ANI-0312461, CNS-0434940, CNS-0519442, and CNS- 0720571. An earlier version of this paper was presented at the IEEE INFOCOM, Barcelona, Spain, 2006.

The authors are with the Department of Computer Science, Texas A&M University, College Station, TX 77843 USA (e-mail: xmwang@cs.tamu.edu; dmitri@cs.tamu.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TNET.2009.2024145

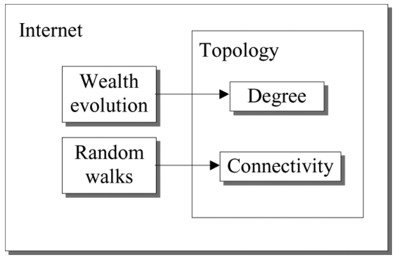

Fig. 1. Components of the wealth-based evolution model.

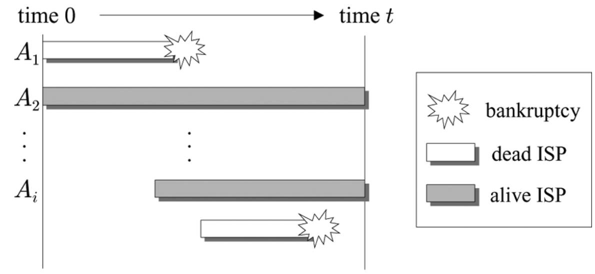

Fig. 2. Birth-death wealth evolution.

nonpower-law degrees and more locality (higher clustering). This requirement calls for a more natural and flexible solution than the previous models.

In this paper, we overcome the above limitations and complement the previous efforts by proposing a different theory for the structure of the Internet that relies on: 1) principles of economic evolution that govern the degree of each ISP; and 2) distributed random walks that determine the actual attachment decisions. While the main focus of this paper is to *understand* the evolution of the Internet, we also provide specific algorithms that can be used to create new graphs and test them against those observed in the Internet over the last decade. In addition, we also discuss possible ways to extend our algorithms to support a wide variety of topologies.

A. Degree

The structure of the proposed model is shown in Fig. 1, where the construction of the graph is driven by two paradigms—wealth evolution and random walks. As shown in the figure, the former is responsible for the degree distribution, while the latter for the formation of actual links. The main principle of the proposed model is that the *degree* of an ISP is a consequence of complex forces that can be macroscopically modeled by the wealth[2] of the ISP and not by the metrics found in the topology itself. This characteristic of the Internet makes it fundamentally different from other real-life graphs such as neural networks [3], [55], actor collaborations [3], scientific citations [43], [46], and numerous networks observed in physics [2], [8], [25], which also exhibit power-law degree distributions, but lack the financial orientation of the Internet.

Since individual and company wealth in many free-market societies is governed by Pareto distributions [41], we argue that the heavy-tailed degree of Internet ISPs is *a result of the particular structure of their wealth rather than anything else*. To understand this correlation, notice that it makes little sense to build topologies in which small local ISPs are modeled with extremely large degree, well-established backbone providers are assigned a handful of peering points, and the structure of individual companies evolves *only based on the degree of other ISPs*. Causality between company wealth and degree can be explained by many factors such as cost of link maintenance that makes higher degree more expensive, customer pressure that forces networks with many subscribers to be better connected, and the various QoS objectives that necessitate more peering points to deliver better service and extract more revenue from transit traffic; however, the exact specifics of this relationship are not essential and may be hidden under the umbrella of a simple economic model discussed below.

To capture the dynamics of open-market competition between the ISPs, our model assigns certain wealth $w_i(t)$ to each ISP i and acts on behalf of the ISP to keep its degree $d_i(t)$ proportional to its wealth at time t. Individual ISP wealth $w_i(t)$ is governed by one of the simplest wealth evolution models that relies on random multiplicative increases/decreases in response to the stock market and various random economic decisions of the company. To account for bankruptcy that is prevalent among new startups, each ISP is removed from the system when its wealth drops below a certain threshold w_b needed to operate and provide service to its customers. This framework is illustrated in Fig. 2, where ISPs $A_1, A_2, \ldots$ sequentially join the system and compete in the Internet market using their individual assets $w_i(t)$ that include their initial funding, customer revenues, stock market gains, etc. As shown in the figure, at any time t, the system is composed of a random number of ISPs that are still alive (i.e., those that avoided bankruptcy), whose distribution of wealth determines the degree structure of the AS-level graph.

B. Links

For the construction of actual links, it is rather clear that the Internet evolves in a *distributed* fashion where the ISPs might not be aware of any global characteristics of the network. To reflect the distributed nature of real attachment decisions, our model allows each ISP that plans to expand to perform random walks along the existing graph until it finds a neighbor that is willing to accept its peering request and satisfy its financial requirements (e.g., offer the right customer base, necessary economic model, and reasonable peering conditions). While the actual attachment decisions in real life are not "haphazard," we argue that the event that a given ISP satisfies all of the above criteria for attachment may be modeled at some high level as *purely random*.

The final note is that our theory of random walks (as opposed to other means of finding neighbors) may be viewed as a by-product of the Internet market being a large social network, where many companies and individuals discover new acquaintances through existing links (i.e., business or personal relationships) rather than by approaching complete "strangers." This also allows our model to preserve locality, i.e., geographically close ISPs are more likely to peer, which is a well-known phenomenon in [4], [6], [19], [24], [28], [31], [34], [54], and [57].

We combine the above two methods (degree evolution and random walks) into a set of algorithms we call *Wealth-based Internet Topology* (WIT). Simulations show that (WIT) succeeds

[2]Company wealth is an abstract concept that includes its revenue, customers, income, property and stock value, equipment, bandwidth, etc.

in producing power-law degree distributions with a flexible exponent α (including $\alpha = 1.2$ observed in the Internet) and is able to achieve levels of clustering close to those in the Internet (i.e., 0.45). More importantly, we find that the clustering coefficient of (WIT) matches that of the Internet during the *entire evolution* of the graph (i.e., as the size of the system increases) rather than for a single value of n as usually examined in prior work. We also show that WIT preserves the evolving nature of other graph-theoretic metrics in the Internet (e.g., assortativity, average path length, and the second smallest eignevalue) better than previous approaches.

C. Possible Extensions

It should be noted that WIT can be extended to support a wider range of degree distributions and clustering coefficients than shown in this paper (see Section VI for more details). The former objective can be achieved by using an *interaction-based* or *exchange-based* wealth evolution process [20], which can produce a variety of degree distributions besides the traditional power-law. The latter objective can be implemented by adjusting the length of random walks for finding neighbors, which leads to different levels of locality and thus clustering. Another possible extension is to construct hierarchical Internet topologies under the framework of (WIT). This can be performed by explicitly labeling nodes as ISPs and non-ISPs during join, limiting neighbor attachment to only ISPs with enough wealth, and creating inter-ISP peering or customer-provider links based on a purely random event or the ratio of wealth among the connecting nodes. This flexibility allows one to cover a relatively general class of topologies, potentially including the true Internet structure that is at present time partially hidden from RouteViews observations.

The remainder of the paper is organized as follows. We first review the background and related work in Section II. We then present our wealth evolution model in Section III and discuss the details of the topology construction algorithm in Section IV. Finally, we compare our model with existing methods in Section V, discuss its possible extensions in Section VI, and conclude the paper in Section VII.

II. Background and Related Work

In this section, we overview a small subset of related work and mention several well-known models that we study later in the paper.

A. Internet Topology and Power-Law Degree Distribution

Faloutsos *et al.*. [16] show that the Internet AS-level topology exhibits a power-law degree distribution, or the so-called "scale-free" phenomenon

$$P(d_i > x) = (x/\beta)^{-\alpha} \quad (1)$$

where d_i is the degree of node i, β is the scale parameter, and α is the shape parameter of the power-law distribution. Note that many similar observations [7], [9] are obtained from the archived snapshots of Border Gateway Protocol (BGP) routing tables collected by the Oregon RouteViews Server [45]. While sometimes it is argued that this data set does not reflect the whole view of the Internet, it has been reported in [9], [22] that graph properties such as the degree distribution are robust even with certain incompleteness in the data set of the Oregon Route-Views Server. In fact, it is shown in [9] that the number of links is the only difference between the graph inferred from the information collected by the Oregon Server and the one complemented with other sources (e.g., the Looking Glass tool and Internet Routing Registry database) and that the power-law degree distribution holds for both graphs. Note that similar observation can be found in [22] regarding the size of the largest connected component.

To model the scale-free property in the Internet, many efforts have been brought forward to design topology generators that produce power-law degrees. Some of them construct random graphs incrementally and others do not allow the growth of the network. We call the former algorithms *evolving* and the latter *nonevolving*. We next review two major classes of evolving models, i.e., preferential-attachment and optimization tradeoffs, and follow it up with a discussion of nonevolving methods.

B. Preferential Attachment

The most common scale-free models used today are based on the theory of "preferential attachment" which is proposed by Barabási *et al.*. [3] and implemented in their topology model *Barabási-Albert* (BA). At each discrete time step, BA adds a new node x to the graph, which is then randomly linked to $m \geq 1$ existing nodes using the preferential-attachment function

$$p_i(t) = \frac{d_i(t)}{\sum_{k=1}^{n(t)} d_k(t)} \quad (2)$$

where $p_i(t)$ is the probability that node i is selected for link formation at time t, $d_i(t)$ is its degree at time t, and $n(t)$ is the number of nodes in the graph at time t. This version of preferential attachment always produces graphs with shape parameter $\alpha \approx 2$. To relax the constraint on α, a method known as *Albert-Barabási* (AB) [2] adds the operations of link rewiring and growth suspension (i.e., the graph evolves without adding new nodes).

Bu *et al.*. [7] utilize shift-parameter $\lambda \in [-\infty, 1]$ in their model, which they call *Generalized Linear Preference* (GLP), and modify (2) to

$$p_i(t) = \frac{d_i(t) - \lambda}{\sum_{k=1}^{n(t)} (d_k(t) - \lambda)}. \quad (3)$$

Through the use of (3), GLP achieves arbitrary values of $\alpha = 2 - \lambda \in [1, \infty)$ and high levels of clustering. Similar methods are proposed by Simon [5], [48], [49] and Krapivsky *et al.*. [26], [27].

Other mechanisms in this category include *BRITE* [34] and *Inet* [23]; however, throughout this paper, we only study BA, AB, and GLP since their performance can be used to infer that of the other models.

C. Optimization-Based Models

Another major class in generating power-law degree distributions was first proposed by Carlson *et al.* [8] and later studied by Fabrikant *et al.* [15] in the context of the Internet. In these models generally called *Highly Optimized Tolerance*

(HOT) each new node selects the attachment point based on the minimization of two objectives: the geographical length of the peering link and the average number of hops to other nodes in the graph. In particular, a new node i attaches to node k that minimizes the following:

$$k = \arg\min_{j<i}\{\theta d_{ij} + h_j\} \quad (4)$$

where d_{ij} is the Euclidean length of link (i,j), h_j is the average distance from j to other nodes in the graph, and θ is a parameter tuning the relative significance of the two objectives d_{ij} and h_j. Chang *et al.*. [9], [10] further explore optimization-based construction methods by allowing each AS to have multiple geographical locations, called *Points of Presences* (PoPs), where each new node i computes (4) by replacing d_{ij} with the minimum distance to all PoPs of node j.

D. Nonevolving Power-Law Generators

In this category, we mention several generators that do not grow (evolve) the network over time. One of the simplest power-law graph construction models is called *Given Expected Degree* (GED) [12], [35], [36], [38]. GED is an extension of the classical Erdös-Rényi graph model $G(n,p)$ [14] in which edge-existence probability p is adjusted on a per-link basis to produce a heavy-tailed degree distribution. Specifically, a sequence of *weights* $\{w_i\}$ is first generated according to a Pareto distribution and then each edge (i,j) is created with independent probability

$$p_{ij} = \min\left(\frac{w_i w_j}{D}, 1\right) \quad (5)$$

where $D = \sum_{k=1}^{n} w_k$. The $\min$ function is necessary since product $w_i w_j$ may exceed D, especially in sequences drawn from power-law distributions with shape parameter $\alpha < 2$.

A similar graph construction method called *Power-Law Random Graph* (PLRG) [1] replicates each node i exactly w_i times and then places random edges between the replicated nodes with equal probability. Thus, nodes with larger initial weight w_i receive proportionally more edges than nodes with smaller weight.

Additional nonevolving generators include random geometric graphs [24] and rewired small-world (Watts) networks that exhibit a heavy-tailed degree distribution [42], [55].

III. Wealth Model

We present our wealth evolution model in this section and show how it fits into our topology generator in Section IV.

A. Wealth Evolution

According to the theory proposed in this paper, the Internet can be modeled as an economic entity, where ISPs dynamically join and leave the system based on random events. Denote by $w_i(t)$ the wealth of ISP i at time t. When a new ISP joins the system at time t_i, it comes with a certain amount of initial wealth w_0, which accounts for the startup capital obtained from venture capitalists, banks, private investors, and other sources.[3] During the lifetime of an ISP, it invests its wealth in business activities, retrieves financial return, and suffers losses, all of which allows its wealth $w_i(t)$ to randomly evolve over time.

In what follows, we adapt the idea from the work of [29], [30], [51] and describe wealth $w_i(t)$ using a random process. Notice that the amount of investment return is usually considered to be proportional to the wealth of a company [29], [30], [51]. Thus, we start with a basic *unscaled* model in which the individual investment-return cycle is a multiplicative stochastic process

$$u_i(t) = \lambda_i(t) u_i(t-1), \text{ for } t > t_i \quad (6)$$

where $u_i(t)$ is the unscaled wealth of user i at time t, $\lambda_i(t)$ is a random variable drawn from some distribution describing the randomness of the investment-return market cycle, and t_i is the join time of node i. We assume that $\lambda_i(t)$ is a stationary process that is independent among the ISPs.

Note that in (6), we do not constrain the distribution of $\lambda_i(t)$ for generality of the model; however, such generality may result in a collapse or explosion of system wealth. Specifically, if $E[\lambda_i(t)] > 1$, the average wealth will grow to infinity as the system evolves. On the other hand, if $E[\lambda_i(t)] < 1$, the average wealth will diminish to zero. To keep system wealth constant, we counteract any possible inflation of wealth by scaling (6) and taking the result to be the *real* wealth of each ISP

$$w_i(t) = \frac{u_i(t)}{\rho(t)} \quad (7)$$

where $w_i(t)$ is the scaled wealth of user i at time t and $\rho(t)$ is a random process that we determine next.

As before, define $n(t)$ to be the number of ISPs in the system and $\bar{u}(t)$ to be the average unscaled system wealth at time t

$$\bar{u}(t) = \frac{1}{n(t)} \sum_{i=1}^{n(t)} u_i(t). \quad (8)$$

Define $\rho(t)$ to be

$$\rho(t) \equiv \frac{\bar{u}(t)}{w_0}. \quad (9)$$

Then combining (7)–(9), we establish that the average scaled wealth of the system remains constant and equals w_0

$$\bar{w}(t) \equiv \frac{1}{n(t)} \sum_{i=1}^{n(t)} w_i(t) = w_0. \quad (10)$$

Note that it may appear that (7) requires global knowledge. However, this is not the case. Stochastic process $\rho(t)$ can be viewed as inflation adjustment inherent in economic systems that is automatically reflected in the cost of goods and services (i.e., money spent in leasing links and the corresponding maintenance). To show that the system can be implemented in a distributed fashion, we use the expectation of the right-hand side of (9) instead of $\rho(t)$ in (7). It is easy to verify that the redefined $\rho(t)$ has the following form:

$$\rho(t) \equiv \frac{E[\bar{u}(t)]}{w_0} = \left(E[\lambda_i(t)]\right)^t. \quad (11)$$

Therefore, we can rewrite (7) as

$$w_i(t) = \frac{u_i(t)}{\left(E[\lambda_i(t)]\right)^t} = \frac{\lambda_i(t)}{E[\lambda_i(t)]} w_i(t-1) \quad (12)$$

[3]Our model uses fixed w_0; however, a simple extension to random startup funding is possible as well. Simulations show that such an extension produces almost identical results.

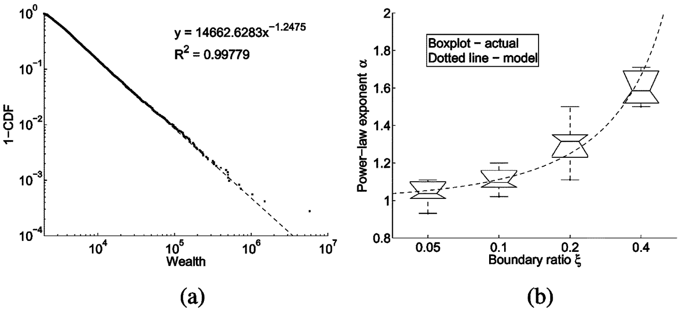

Fig. 3. Wealth distribution at time $t = 200$ under static join (10 000 joining ISPs). (a) $\xi = 0.2$. (b) Effect of boundary ratio ξ.

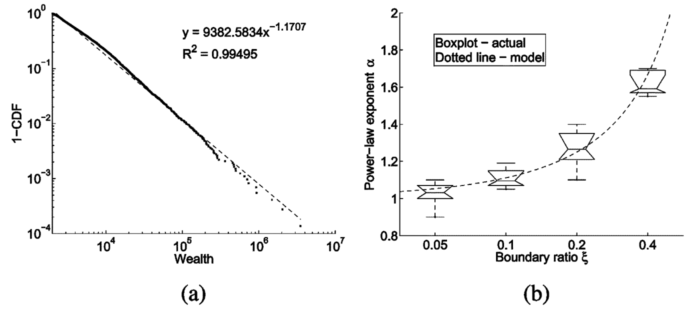

Fig. 4. Wealth distribution at time $t = 200$ under dynamic join ($\mu = 50$). (a) $\xi = 0.2$. (b) Effect of boundary ratio ξ.

with the initial condition $w_i(0) = w_0$. It is worth noting that $w_i(t)$ in (12) only depends on $\lambda_i(t)$ and $w_i(t-1)$, which can be computed locally. In addition, it is straightforward to verify that the average scaled wealth $\bar{w}(t)$ of the system defined by (10) using the new definition of $\rho(t)$ in (11) is asymptotically constant and equal to w_0

$$\lim_{n(t)\to\infty} \bar{w}(t) = w_0. \tag{13}$$

Given the model of ISP wealth represented by (12), we next examine the conditions of bankruptcy and obtain the power-law exponent of $w_i(t)$ as a function of the bankruptcy boundary. As discussed in the introduction, we impose a lower boundary $w_b = \xi w_0$ on each ISP's wealth such that no one in the system is poorer than this baseline. For simplicity of discussion in the rest of the paper, we use metric $\xi \in (0,1)$ as the ratio of the bankruptcy boundary w_b to the initial wealth w_0 of each joining ISP. Armed with (10) and bankruptcy definitions above, one can observe that the system defined by(12) is the same to the multiplicative stochastic process with a reflective barrier, which is well studied in economics. Specifically, the wealth distribution produced by(12) follows a power law with exponent α determined by the lower boundary w_b scaled by the average social wealth [29], [30]. In addition, [51] points out that α's sole dependence on w_b holds only if the average social wealth $\bar{w}(t)$ is fixed. Since this condition asymptotically holds in (13), we immediately obtain the following result.

Theorem 1: For sufficiently large t, the wealth evolution process $w_i(t)$ characterized by (12) achieves a power-law distribution with exponent

$$\alpha \approx \frac{1}{1-\xi} \tag{14}$$

where $0 < \xi < 1$ is the ratio of the lower boundary w_b of personal wealth to the initial wealth w_0.

B. Simulations

We confirm (14) under two different ISP join schemes. The first method, which we call *static*, forces all individual ISPs to start their wealth evolution processes at the same time. We plot in Fig. 3(a) the resulting distribution of $w_i(t)$ and observe that it is indeed power-law with an exceptionally good fit. We next vary ξ in simulations and generate 1000 instances of the evolution process to examine the correlation between ξ and α. Fig. 3(b) presents the resulting boxplot[4] distribution of actual power-law exponents α and shows that their average value agrees with model (14) very well.

The second method, which we call *dynamic*, allows ISPs to join the system according to some arrival process of rate μ. We experimented with several join processes and observed no impact on the corresponding wealth distribution. Without loss of generality, the rest of the paper uses Poisson arrivals of rate μ ISPs per time unit and keeps the average number of ISPs that join the network by time t equal to μt. Simulations shown in Fig. 4 demonstrate that dynamic join also produces power-law wealth distributions and that the corresponding expected power-law exponent $E[\alpha]$ follows (14) very accurately.

C. Discussion

In addition to multiplicative stochastic processes, prior efforts also use an interaction-based mechanism to model a wealth process. We refer readers to [20] for an extensive review of this direction. In such a model, two randomly selected nodes i and j *exchange* their wealth: at time t, node i loses some amount of its wealth to node j and the sum of their wealth remains constant before and after such an interaction, i.e., $w_i(t+1)+w_j(t+1) = w_i(t) + w_j(t)$ for all t. This scenario of wealth exchange arises when two ISPs form a contractual partnership, in which one provides services to the other and as a result gets paid over the service period. According to [20], the wealth exchange model can achieve a wide variety of wealth distributions ranging from exponential to Gamma, which could be useful when the target distribution is not power-law. However, due to the page limit, we leave the discussion of the wealth exchange model to future work.

With the result in (14), we are ready to present our topology model and show how the power-law wealth distribution affects topology generation in the next section.

IV. Topology Model

We start this section by introducing the role of wealth in our topology model, then present link-construction algorithms, and finally study the degree distribution and clustering properties of the resulting generator.

[4]The boxplot produces a box and whisker plot for each variable (e.g., power-law exponent or clustering coefficient). The box has lines at the lower quartile, median, and upper quartile values. The whiskers are lines extending from each end of the box to show the extent of the rest of the data. Outliers are data with values beyond the ends of the whiskers.

A. From Wealth to Topology

This subsection describes how the wealth model interacts with the topology generator and how it determines the degree evolution of the Internet.

In the hypothetical Internet market modeled in this paper, the connectivity of an ISP to the rest of the network decides its usefulness to the customers. In order to maximize its revenue, an ISP tends to build as many links to other ISPs as possible. However, such link expansion is always limited by the wealth of the ISP since each link incurs a certain amount of expense, which represents the cost of purchasing routers and other equipment, leasing bandwidth, and maintenance. Therefore, an ISP builds additional links when it has spare wealth and similarly removes links when its wealth cannot sustain the expense of existing links.

In this regard, we model the degree of an ISP as being proportionally dependent on its wealth and closely related to random fluctuations in $w_i(t)$. In reality, it is possible that ISP connectivity and wealth might be dependent through a more complex closed-loop system; however, the proposed model is a simple abstraction of this unmeasurable process, in which wealth evolves independently of ISP's topological properties (e.g., only based on the stock market, current business plan, customer satisfaction).

Suppose that the link between nodes i and j induces certain cost C. Denote by $z_i(t) = Cd_i(t)$ the expense induced by all links of node i at time t. Then, whenever the wealth of an ISP i drops below its current link expense

$$w_i(t) - z_i(t) < 0 \tag{15}$$

the ISP must remove some of the existing links to reduce expense $z_i(t)$ below its wealth $w_i(t)$. On the other hand, if the wealth of ISP i allows more links than it currently has

$$w_i(t) - z_i(t) > C \tag{16}$$

new connections are built until $z_i(t)$ reaches its wealth limit.

A direct result of the above mechanism for link adjustment is the linear mapping between link expense and wealth

$$z_i(t) \approx w_i(t) \tag{17}$$

which leads to the following result:

$$d_i(t) \propto w_i(t). \tag{18}$$

Recalling that the economic system presented in (6)–(9) produces a power-law wealth distribution, the next theorem follows immediately.

Theorem 2: For large enough t, the degree distribution of random graphs constructed under (15) and (16) is power-law with exponent

$$\alpha \approx \frac{1}{1-\xi} \tag{19}$$

where constant ξ is the wealth boundary ratio.

Proof: Notice that the degree $d_i(t)$ of node i is proportional to its wealth $w_i(t)$, that is

$$d_i(t) = cw_i(t) \tag{20}$$

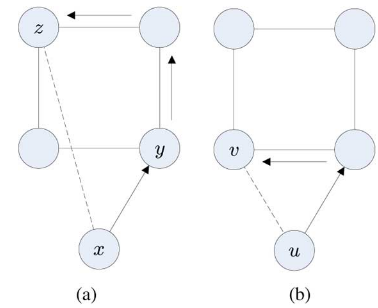

Fig. 5. Random walks in topology generation. (a) Initial attachment. (b) Link addition.

where c is constant for all t and i. Furthermore, the wealth $w_i(t)$ of node i has a power-law distribution with exponent given by (14) and constant c in (20) only affects the corresponding scale parameter. Therefore, we have that $d_i(t)$ also follows a power-law distribution with the same exponent as that in (14). ■

Notice that conditions (15)–(16) may result in oscillatory link behavior, which is not the case in reality. To this end, we provide a dampening threshold T to relax (15)–(16) to read

$$w_i(t) - z_i(t) < -T \tag{21}$$

and

$$w_i(t) - z_i(t) > C + T. \tag{22}$$

By carefully choosing dampening threshold T, we are able to suppress the oscillations in link adjustment while allowing (19) to hold (see below for simulations).

B. Topology Construction

At each step t, our generator first updates the wealth of all existing nodes according to (12) and then evolves the topology by adding new nodes and adjusts links of the existing nodes based on changes in their wealth. In this section, we introduce a basic model that uses uniform labeling for all links. In Section VI, we discuss a possible extension that exhibits an explicitly hierarchical structure with peer-peer and customer-provider links. The detailed algorithms of link addition and deletion are given next.

As mentioned in the Introduction, link addition in our topology generator depends on a simple set of rules based on random walks. When a new node x decides to enter an existing graph, it uniformly and randomly selects one existing node y in the network as the point of entry. Once the node is introduced into the graph, it decides to "explore" its new location by performing a random walk of l steps. Once the initial walk stops at a node z, the new node x establishes a link to z, which is illustrated as the dotted line in Fig. 5(a). The above represents a selection process in which x searches for the first "acceptable" peering ISP. In the second phase of the join, node x starts from z and performs $m - 1$ additional random walks to find the remaining $m - 1$ neighbors, where m is determined by the initial wealth w_0 and the link price C. The decision to attach is determined solely by the final node where the walk stops and represents the random process where ISPs seek business partners with matching interests.

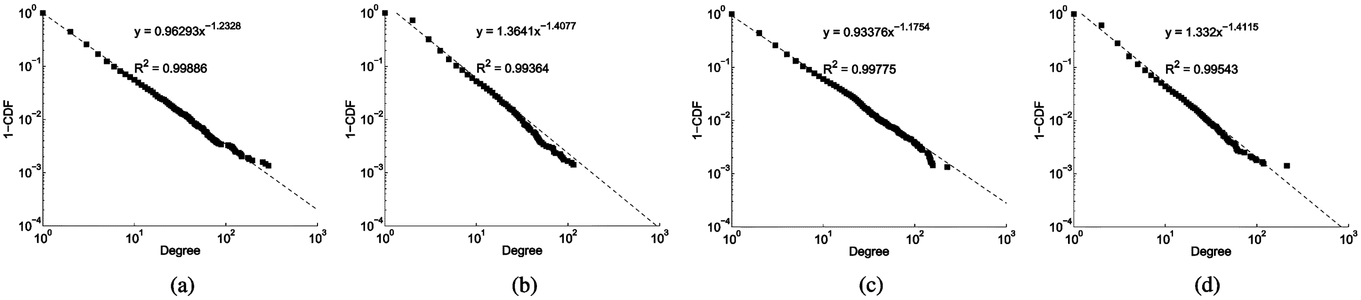

Fig. 6. Degree distributions of (WIT) graphs ($t = 200, \mu = 50, c = 0.4, \tau = 0.1$). (a) $\xi = 0.2$ and $l = 3$. (b) $\xi = 0.4$ and $l = 7$. (c) $\xi = 0.2$ and $l = 15$. (d) $\xi = 0.4$ and $l = 100$.

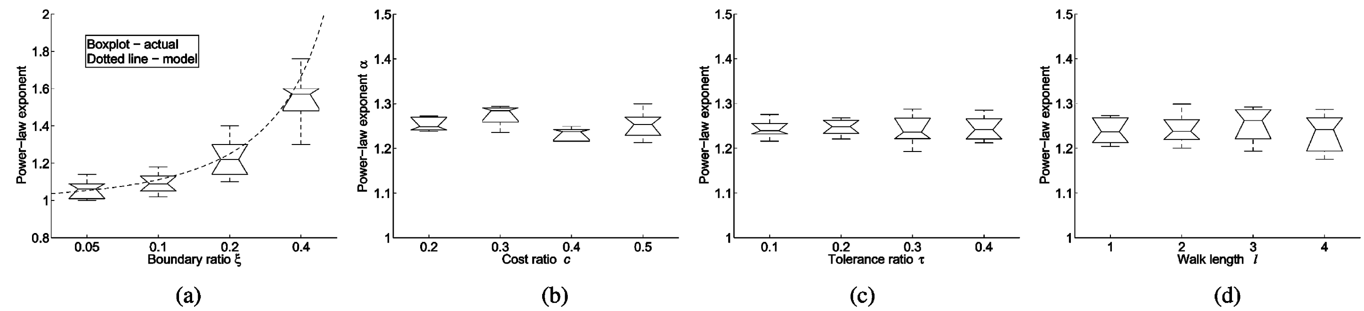

Fig. 7. Shape parameter α under $\xi = 0.25, c = 0.4, \tau = 0.1, l = 1$ ($t = 200, \mu = 50$). (a) Effect of boundary ratio ξ. (b) Effect of cost ratio c. (c) Effect of dampening ratio τ. (d) Effect of walk length l.

When an existing node u needs to build a new link, the procedure of finding a new neighbor is the same as in the initial attachment except that random walks start from u as shown in Fig. 5(b). On the other hand, when a node is forced to eliminate some of the existing links, it uniformly and randomly chooses a neighbor from its peering list and terminates the corresponding connection.

Finally, our topology model starts at time zero with a fully connected core network of size m_0, which represents the initial stage of the Internet. Since simulations show that m_0 does not affect the properties of constructed graphs, we use the commonly suggested value $m_0 = 3$ from prior methods BA, AB, and GLP. If a user fails to find a suitable attachment point after 20 random walks, it preserves its current wealth $w_i(t)$ until the next cycle $i + 1$.

We refer to the set of algorithms described above as *Wealth-based Internet Topology* (WIT) and next study its degree distribution and clustering properties.

C. Degree Distribution

Before we begin, we first verify the quality of power-law distributions produced by (WIT). Fig. 6 plots four examples of degree distributions obtained in simulations using several different values of l and ξ. This figure combined with additional results indicates that WIT's degree exhibits very clear power-law tails that hold remarkably well for both short and long walks l.

Our analysis of (WIT) in this subsection focuses on how degree exponent α is affected by four parameters: lower boundary w_b, link cost C, dampening threshold T, and walk length l. For convenience, we normalize the first three metrics by w_0 to obtain *lower-boundary ratio* $\xi = w_b/w_0$ (as defined in the previous section), *cost ratio* $c = C/w_0$, and *dampening ratio* $\tau = T/w_0$.

To avoid confusion as to which parameter is responsible for which graph property, we study the effect of these factors separately by changing only one of them and keeping the other three fixed. We generate 1000 random graphs for each point in the figures and show the distribution of shape parameter α in the boxplots of Fig. 7. The results in Fig. 7(a) demonstrate that α increases as a function of ξ and follows model (19) very accurately. Additionally, Fig. 7(b)–(d) indicate that α is not sensitive to cost ratio c, dampening ratio τ, or walk length l, which also agrees with Theorem 2 very well.

Numerous additional simulations with different parameters show similar results (omitted for brevity) and conclusively establish that boundary ratio ξ is the *only* parameter that affects shape α of degree distribution, which explains our global view of the model in Fig. 1. By tuning ξ, one can achieve arbitrary power-law exponents $\alpha \in [1, 2]$, and, as we show in the next subsection, tuning walk length l allows (WIT) to achieve flexible clustering $\gamma \in [0.008, 0.64]$.

D. Clustering

The clustering coefficient of a graph measures how frequently neighbors of a node connect to each other. Define T_i to be the number of triangles incident to node i. The clustering coefficient γ_i of node i is [56]

$$\gamma_i = \frac{T_i}{d_i(d_i - 1)/2} \tag{23}$$

where d_i is the degree of node i. Then for a graph $G(V, E)$, its clustering coefficient is given by [56]

$$\gamma(G) = \frac{1}{|V| - |V^{(1)}|} \sum_{i \in V - V^{(1)}} \gamma_i \tag{24}$$

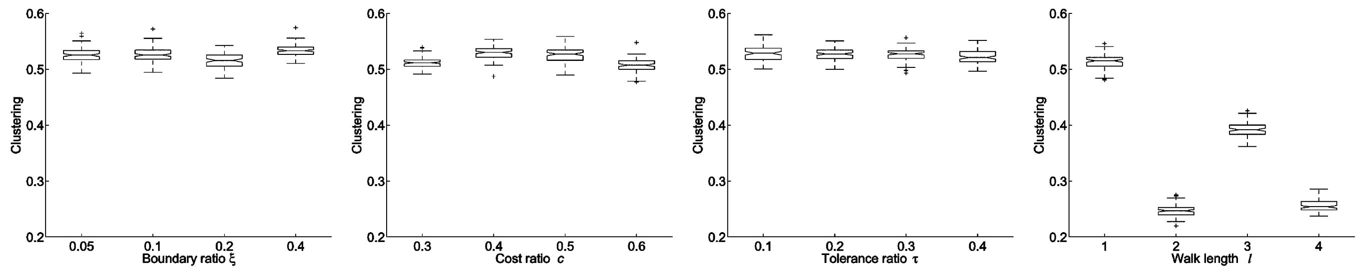

(a) (b) (c) (d)

Fig. 8. Clustering coefficient under $\xi = 0.2, c = 0.5, \tau = 0.1, l = 1$ ($t = 200, \mu = 50$). (a) Effect of boundary ratio ξ. (b) Effect of cost ratio c. (c) Effect of dampening ratio τ. (d) Effect of walk length l.

where $V^{(1)}$ is the set of degree-one nodes in G.

In what follows, we study clustering of (WIT) and show that it only depends on walk length l and not other parameters of the generator. Again, we conduct four sets of simulations as in the previous subsection and vary one of the four parameters in each set while keeping the other ones fixed. The boxplot of Fig. 8 shows clustering coefficients of (WIT) graphs generated under these conditions. In Fig. 8(a)–(c), the average clustering coefficient stays around 0.52 and does not exhibit much correlation with the change in the corresponding parameters. On the other hand, Fig. 8(d) shows that $\gamma(G)$ responds to walk length l, which we analyze in more detail next.

We start our discussion of how random walks determine local connectivity and clustering under the assumption that walk lengths are large, i.e., $l \gg 1$. Recall that random walks on graphs represent the evolution of a stationary Markov chain. For a stationary chain the probability for a walk to terminate at node i is simply [58]

$$p_i(t) = \frac{d_i(t)}{\sum_{k=1}^{n} d_k(t)} \tag{25}$$

which is exactly the same as preferential probability (2). Therefore, we immediately obtain the following result.

Theorem 3: For walks longer than the mixing time of the corresponding Markov chain, (WIT) 's clustering reduces to that of preferential attachment.

To validate Theorem 3, we implement a variant of our model that deploys preferential attachment instead of random walks in link construction. We refer to this variant as WIT-PA and compare it to pure (WIT) in simulations. By setting $l = 1{,}000$, we generate 1000 (WIT) graphs of different sizes n and plot their average clustering coefficients along with those of WIT-PA in Fig. 9(a). The figure shows that $E[\gamma(G)]$ of the two models is almost identical.

Analysis of short (i.e., significantly smaller than the mixing time of the chain) walks in (WIT) becomes nontrivial since the stationary distribution (25) does not hold. For $l = 1$, (WIT) always produces a lattice of triangles, where each node with degree $d_i \geq 2$ has $d_i - 1$ triangles. Therefore, without considering link deletion and rewiring, the clustering of node i is given by $2/d_i$, where d_i is the degree of node i. This immediately leads to

$$E[\gamma(G)] = E\left[2/d_i\right] = \int_{\beta}^{\infty} \frac{2}{x} dF(x) \tag{26}$$

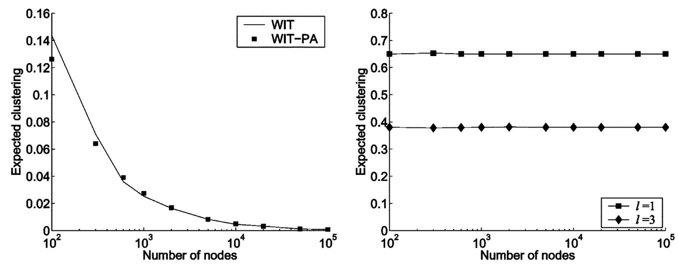

(a) (b)

Fig. 9. Clustering in (WIT) ($\xi = 0.2$, $c = 0.45$, $\tau = 0.3$). (a) $E[\gamma(G)]$ for $l = 1000$. (b) $E[\gamma(G)]$ for $l = 1$ for $l = 3$.

where $F(x) = 1 - (\beta/x)^{\alpha}$ is the CDF of the power-law degree distribution. It follows from (26) that

$$E[\gamma(G)] = \frac{2\alpha}{(1+\alpha)\beta}$$

which is a constant independent of the graph size n. Combining this with $\alpha = 1.2$ and $\beta = 1.45$ (obtained from simulations with $\xi = 0.2$, $c = 0.45$, $\tau = 0.3$, $l = 1$), we get $E[\gamma(G)] = 0.75$. The actual expected clustering of (WIT) for $l = 1$ is slightly lower and equals 0.64 as shown in Fig. 9(b), which can be explained by link deletion and rewiring not included in the above analysis.

For $l = 2$, new nodes produce mostly quadrangles instead of triangles and thus construct a poorly clustered graph, while for $l = 3$, (WIT) builds a mixture of triangles and pentagons, and exhibits lower clustering than with $l = 1$, but much higher than with $l = 2$. Fig. 9(b) plots (WIT) 's average clustering coefficient for $l = 3$ and shows that it also stays constant as the graph evolves.

Alternating behavior in clustering between odd and even walk lengths is obvious for short walks and disappears when l becomes long enough. In Fig. 10(a), we show that (WIT) 's clustering coefficient starts from 0.64 with $l = 1$, drops to 0.008 with $l = 2$, then oscillates with a decreasing amplitude, and finally converges to 0.038 as walk length l reaches 40. For larger n, Fig. 10(b) shows that when the system contains more "randomness" (i.e., 2000 nodes join the system), the clustering coefficient converges to its asymptotic value much quicker than in Fig. 10 (a).

In reality, it is not hard to conceive that a new node in the Internet prefers short walks instead of long ones when deciding on link attachment. This in practice means that (WIT) equipped

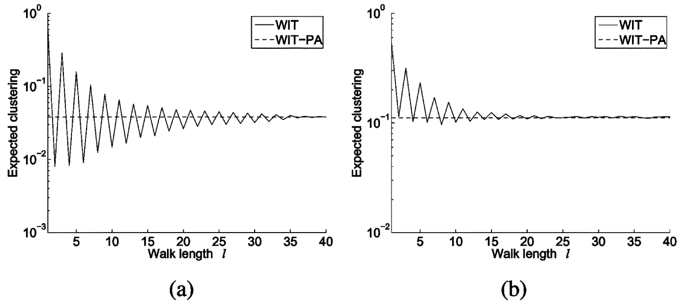

Fig. 10. Effect of walk length on the clustering coefficient ($\xi = 0.2, c = 0.45, \tau = 0.3$). (a) $E[\gamma(G)]$ for $n = 500$. (b) $E[\gamma(G)]$ for $n = 2000$.

with short walks builds graphs with both constant and high clustering. Other preferential attachment-based algorithms [7], [18], [60] on the other hand, only capture the behavior of long walks, produce small and decreasing clustering [see Fig. 9(a)], and, thus, cannot fully explain the observed structure of the Internet as we show next.

V. Evolutionary Comparison of Topology Models

With the topology model developed in the previous sections, we are now ready to answer the question of whether our model can track the evolution of the Internet. We start by understanding Internet's graph-theoretic properties as functions of time.

A. Metrics of Interest

In addition to the degree distribution and the clustering coefficient, the *characteristic path length* and *average degree* of a given graph are usually used in the evaluation of topology generators [3], [7]. Considering that analyzing degree correlation of neighboring nodes and spectrum have recently become two important methods in graph characterization [11]–[13], [37], [53], we include two additional metrics in our comparison: *assortativity coefficient* r and the *second smallest eigenvalue* λ_2 of the normalized Laplacian matrix of each graph. A combination of these six metrics is usually sufficient to distinguish between the existing random graph models.

Denote by $h(x,y)$ the hop distance between nodes x and y and by $h(x)$ the average distance from x to the rest of the graph. Recall that characteristic path length L is defined as the median of average distances $h(x)$ over all nodes in the graph, i.e., $L = \text{median}_{x \in V}\{h(x)\}$ [56]. For the degree-degree correlation, we first define q_{xy} to be the *empirical joint degree distribution* of graph $G(V,E)$

$$q_{xy} = \frac{|S_{xy}|}{|E|} \tag{27}$$

where set $S_{xy} \subseteq E$ consists of edges whose end-nodes have degrees x and y. Then, the assortativity r of a given graph is [13], [53]

$$r = \frac{E[d] \sum_{x,y} xy \left(E[d]q_{xy} - xyq_xq_y\right)}{E[d]E[d^3] - E^2[d^2]} \tag{28}$$

where $q_x = P(d = x)$ is the PMF of degree and $E[d^k]$ is its corresponding k-th moment. Note that metric r reflects how the degree of neighboring nodes is correlated with each other [53] negative values of r imply that high-degree nodes tend to establish links to low-degree nodes; on the contrary, graphs with positive r have many links connecting nodes with similar degrees; and $r = 0$ implies that the degree of neighbors is not correlated.

For the spectral analysis, we adopt the normalized Laplacian matrix as defined in [11], [53]

$$\mathcal{L}(G)(i,j) = \begin{cases} 1, & \text{if } i = j \text{ and } d_i \neq 0 \\ -(d_id_j)^{-1/2}, & \text{if } (i,j) \in E \\ 0, & \text{otherwise} \end{cases} \tag{29}$$

where d_i is the degree of node i. Denote by $\{\lambda_i\}$ the eigenvalues of $\mathcal{L}(G)$ and label them in an increasing order, i.e., $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_{n-1} \leq \lambda_n$, where n is the graph size. Note that $\lambda_1 = 0$ and $\lambda_2 > 0$ for any connected graph [11], [37]. Further recall that metric λ_2 of a graph is closely related to its connectivity, graph expansion, bisection width, and diameter; more specifically, small λ_2 implies that the graph is poorly connected and vulnerable to node/edge failures [11], [37]. Thus, to better understand the resilience of constructed topologies, we focus on λ_2 instead of the entire spectrum in our graph-evolution study.

With the help of these metrics, we next explore how the Internet has evolved in the last 8 years.

B. Observing Internet Evolution

The Internet topology is usually built using: 1) BGP routing tables provided by the Oregon RouteViews project [45] and RIPE [44]; or 2) traceroute data obtained from Skitter [50], DIMES [47], or iPlanet [32]. Both measurement approaches have their own drawbacks and their topologies are generally believed to be incomplete. This paper does not aim to solve the topology completeness issue, but rather evaluate the performance of our method against one of the existing datasets. Throughout the rest of the paper, we use BGP tables from Oregon and leave additional comparisons for future work.

We collected AS snapshots from BGP routing tables [45] between November 1997 and January 2006, extracted the corresponding AS topologies, and computed for each graph its average degree, power-law exponent α, clustering coefficient γ, characteristic path length L, assortativity r, and the second smallest Laplacian eigenvalue λ_2. To examine the dynamic behavior of the Internet, we plot in Fig. 11 these graph properties against the size of the Internet. Fig. 11(a) shows that the degree distribution exhibits a constant power-law exponent, which is rather stable in $[1.15, 1.23]$. Besides the scale-free degree distribution, the Internet is almost invariant in its average degree and characteristic path length, which stay around 4 and 3.7 in Fig. 11(b)–(c), respectively. More interestingly, Fig. 11(d) indicates that the clustering coefficient of the Internet is not only high as reported in [7], [25], but also fairly constant between 0.35 and 0.47. In Fig. 11(e), Internet's assortativity is approximately -0.19, which indicates that ISPs tend to connect to networks with dissimilar degree. The second smallest eigenvalues of the Internet in Fig. 11(f) are contained within $[0.02, 0.04]$ with a linear decaying trend.

C. Comparison of Topology Models

It is possible for a topology generator to construct graphs that match the structural metrics of the Internet at a given time point (i.e., size n). However, as the Internet evolves and its size increases, graph properties of the generator may deviate

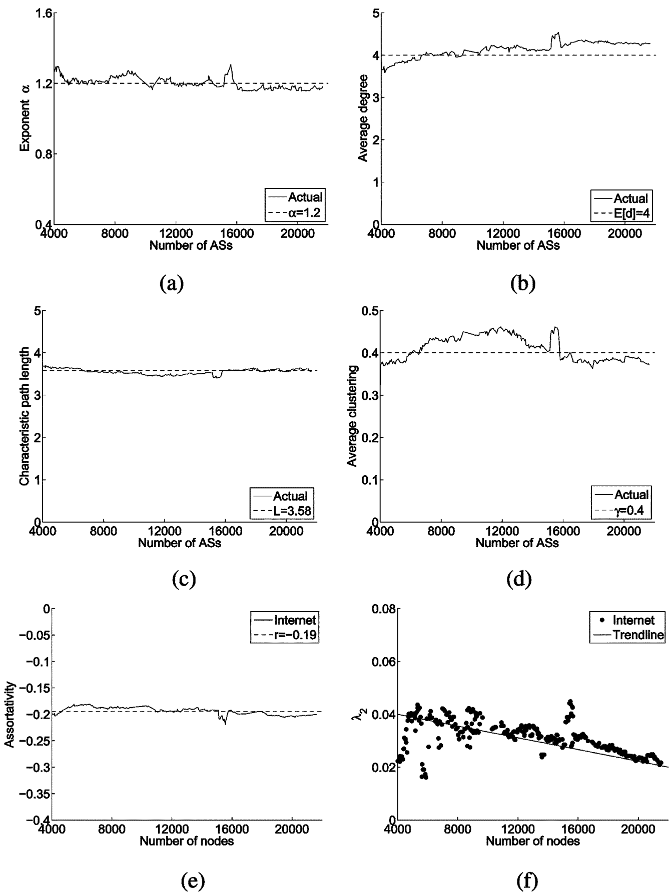

Fig. 11. Evolutionary view of the Internet topology observed from the Oregon RouteViews project [45] between November 1997 and January 2006. (a) Degree exponent α. (b) Average degree. (c) Characteristic path length L. (d) Clustering coefficient $\gamma(G)$. (e) Assortativity r. (f) Second smallest eigenvalue λ_2.

from those in the Internet. Therefore, it is important to compare existing topology models from an *evolutionary* point of view, which tracks the corresponding graph metrics over the entire construction process. Fig. 11 shows that even though the size of the Internet keeps increasing over time, the six graph-theoretic properties remain more or less invariant to the growth. The main question we aim to address in this section is *whether this invariance is captured by the existing generators?* We answer this question by examining how several existing models behave during the graph construction process and its evolution. Note that our *dynamic* comparison complements previous *static* efforts [7], [25] since it performs analysis from a completely different perspective.

In particular, we compare (WIT) to several classical topology generators. For preferential attachment, we use BA, AB, and GLP. For optimization-based algorithms, we study HOT and allow each new node to link to $m \geq 2$ peering points. For nonevolving models, we modify GED to support incremental construction, where each new node joins the system and builds links to existing nodes with the probability described by (5). We refer to our version of GED as *Evolving GED* (EGED).

In our simulations of these generators, we attempt to ensure that the average degree and exponent α are the same as in the Internet (i.e., 4 and 1.2, respectively). Considering that BA and HOT always produce fixed $\alpha \approx 2$, we allow this rule to be violated in these two methods, but guarantee full conformance of the two metrics in the remaining models examined in this paper. Further note that AB and EGED usually produce disconnected graphs, where the metric of characteristic path length L cannot be applied. However, it has been proved in [1] that for a power-law graph with exponent $1 < \alpha < 2$, there exists a giant component of size $\Theta(n)$ and all smaller components are of size $O(1)$, almost surely. Since $\alpha = 1.2$ in our comparison study, we only examine graph properties of the largest connected component of graphs built by AB and EGED. The difference in using the subgraph induced by the largest connected component instead of the whole graph can be neglected for sufficiently large graph size n.

In our first simulation, we let each generator build a random graph with 22 000 nodes. During graph construction, we record snapshots of the partial graphs at different time epochs and compute their clustering coefficients, characteristic path lengths, assortativity coefficients, and the second smallest eigenvalues. We omit the snapshots of small graphs to be consistent with the size of the Internet whose structure before 1997 is not currently available.

As we show below, oscillation in the clustering coefficient exists in all studied generators. To augment the information provided by a single instance of each stochastic process, we also show the expected clustering coefficients in all studied methods. For each generator, we create 1000 random graph evolutions and average the clustering coefficient at each time t. All figures discussed below plot instant clustering as solid lines and their expected values as dotted curves. In what follows, we compare these simulation results against the Internet topology observed from the Oregon RouteViews project [45].

Fig. 12(a) shows that BA exhibits very small clustering coefficients, which decay towards zero as the graph grows in size. This is clearly not representative of the situation in the Internet in Fig. 11(d). The characteristic path length of BA grows from 4.6 to 5.1 as shown in Fig. 12(b). Compared to the Internet where L is *fixed* around 3.7, BA tends to push new nodes away from the center of the graph and, thus, produces *increasingly* larger characteristic path lengths as n grows. The curve of BA's assortativity in Fig. 12(c) tends to 0, which indicates that the correlation of neighbor degree vanishes as the graph evolves in its size, which also does not match the evolution of this metric in Fig. 11(e). In Fig. 12(d), BA's λ_2 shows a similar decreasing trend as that in Fig. 11(f), but with a larger intercept, which indicates that BA builds graphs that are less congested than the Internet.

AB improves over BA in terms of clustering as shown in Fig. 13(a); however, its $\gamma(G)$ also decreases as the graph size increases. Fig. 13(b) indicates that AB's characteristic path length is not only high, but also increasing as a function of n, which is similar to that in BA. Fig. 13(c) shows that AB's assortativity r is constant at zero, which is the result of rewiring. In Fig. 13(d), AB's λ_2 is stable around 0.2, which is 4 times larger than that of the Internet.

Among the three preferential attachment methods, GLP shows the highest clustering in Fig. 14(a) reaching as high as 0.37 for $n = 4000$. However, as n increases to 22 000, $\gamma(G)$ drops to 0.32, which is a common drawback of all examined preferential attachment methods, whose clustering coefficient

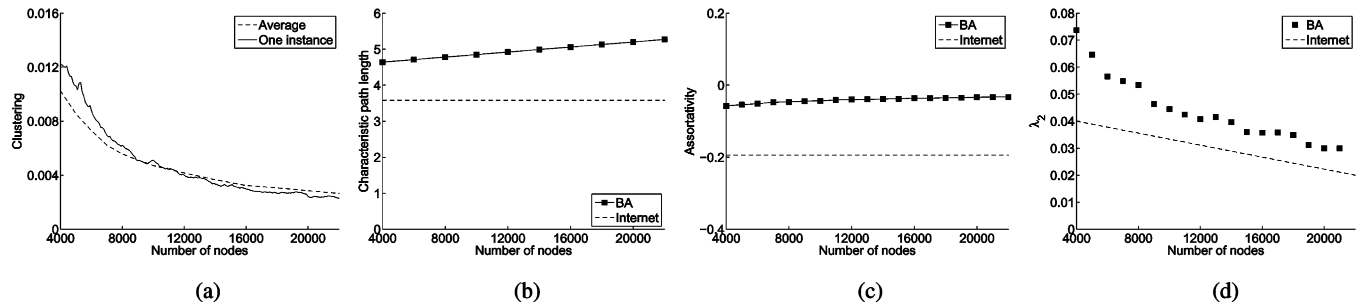

Fig. 12. Evolution of BA compared to the Internet topology observed from the Oregon RouteViews project [45]. (a) Clustering $\gamma(G)$. (b) Characteristic path length L. (c) Assortativity r. (d) Second smallest eigenvalue λ_2.

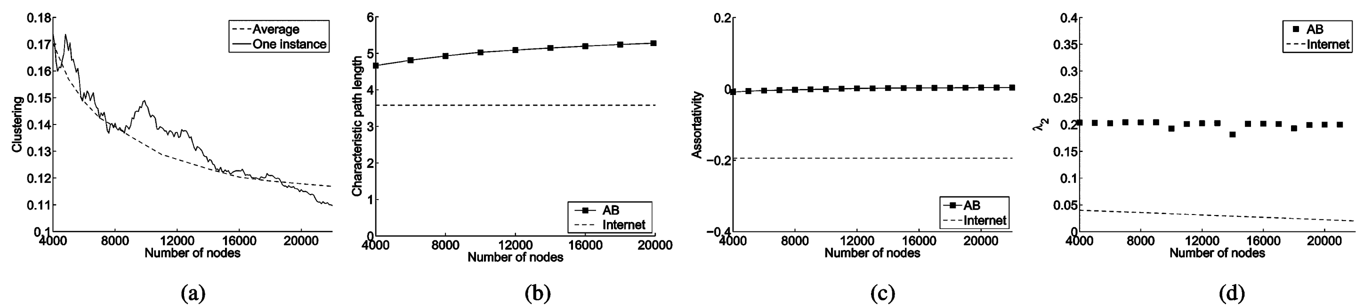

Fig. 13. Evolution of AB compared to the Internet topology observed from the Oregon RouteViews project [45]. (a) Clustering $\gamma(G)$. (b) Characteristic path length L. (c) Assortativity r. (d) Second smallest eigenvalue λ_2.

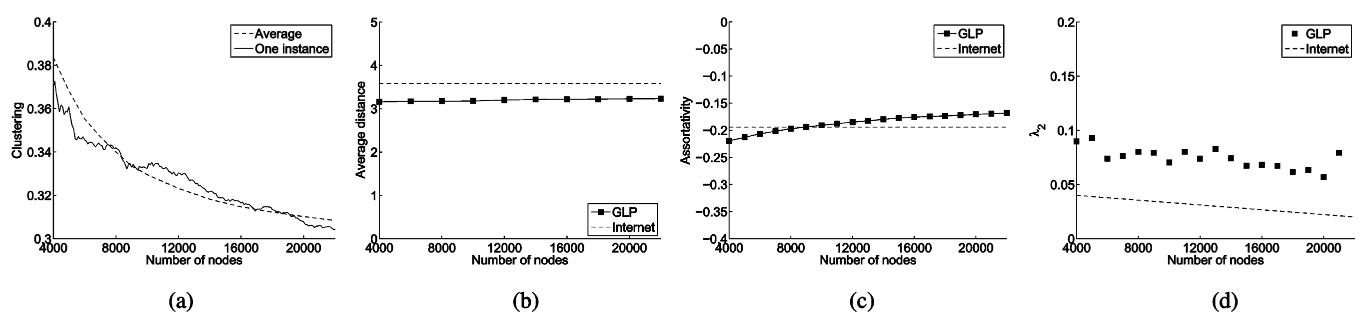

Fig. 14. Evolution of GLP compared to the Internet topology observed from the Oregon RouteViews project [45]. (a) Clustering $\gamma(G)$. (b) Characteristic path length L. (c) Assortativity r. (d) Second smallest eigenvalue λ_2.

decreases as $n \to \infty$. From Fig. 14(b), we observe that the characteristic path length of GLP is close to that of the Internet and its corresponding time average is 3.2. As shown in Fig. 14(c), GLP's assortativity is similar to that of the Internet in terms of its time average; however, its trend deviates from what is observed in the Internet. In Fig. 14(d), the second smallest eigenvalue of GLP graphs oscillates around 0.074, which is twice as large as Internet's λ_2.

In the category of nonevolving methods, EGED also demonstrates decaying clustering in Fig. 15(a) and keeps its $\gamma(G)$ significantly smaller than that of the Internet. Its characteristic path length starts from 2.9, but then "overshoots" to 4.0 when n reaches 22 000 as shown in Fig. 15(b). In Fig. 15(c), the assortativity of EGED is fairly stable around -0.14, which is slightly larger than the time average found in the AS graphs. The corresponding second smallest eigenvalue in Fig. 15(d) indicates significant variance and has a larger time average than that of the Internet.

Interestingly, HOT exhibits in Fig. 16(a) very high clustering that oscillates around 0.45. However, its characteristic path length is much higher than in the Internet and increases from 6.1 to 8.1 almost as a linear function of n as shown in Fig. 16(b). The preference of HOT for geographically short links leads to a graph that spreads out over the entire coordinate plane and thus results in a significantly larger characteristic path length than needed to model the Internet. This is also reflected in the evolution of r and λ_2 plotted in Fig. 16(c)–(d): assortativity r is stable around -0.27, which is smaller than observed in the Internet; metric λ_2 is almost negligible compared to that in Fig. 11(f), which implies a substantially more congested structure compared to the AS topology.

Similar to HOT, (WIT) displays high clustering during the entire graph evolution as shown in Fig. 17(a). The average clustering starts from 0.39 for $n = 4000$ and converges to 0.43 for $n = 22000$. Instant clustering oscillates around the dotted line in Fig. 17(a) and at certain points reaches as high as 0.45, which closely mimics the random fluctuation of γ in the Internet. In addition to producing flexible α, (WIT) is different from HOT in terms of its characteristic path length, assortativity, and eigenvalues. In Fig. 17(b)–(c), (WIT) 's metrics L and r are initially

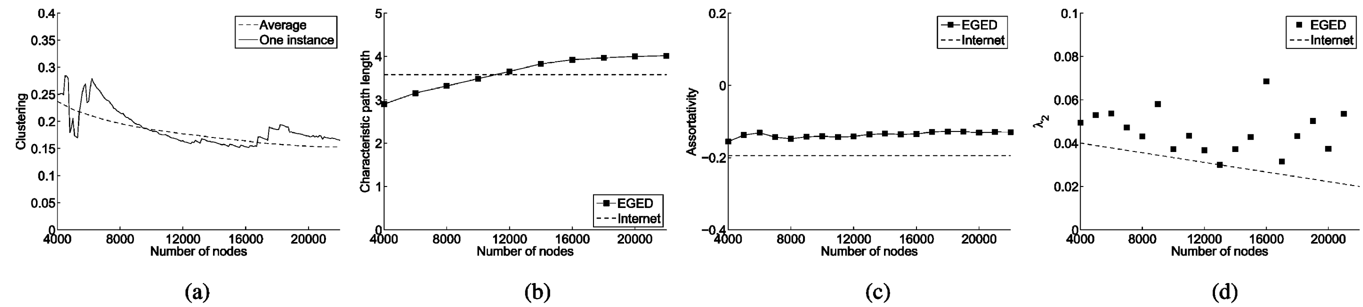

Fig. 15. Evolution of EGED compared to the Internet topology observed from the Oregon RouteViews project [45]. (a) Clustering $\gamma(G)$. (b) Characteristic path length L. (c) Assortativity r. (d) Second smallest eigenvalue λ_2.

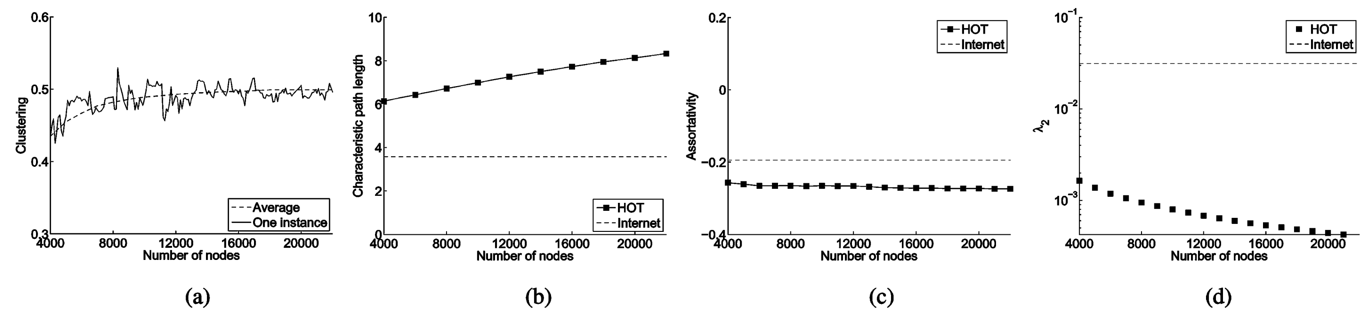

Fig. 16. Evolution of HOT compared to the Internet topology observed from the Oregon RouteViews project [45]. (a) Clustering $\gamma(G)$. (b) Characteristic path length L. (c) Assortativity r. (d) Second smallest eigenvalue λ_2.

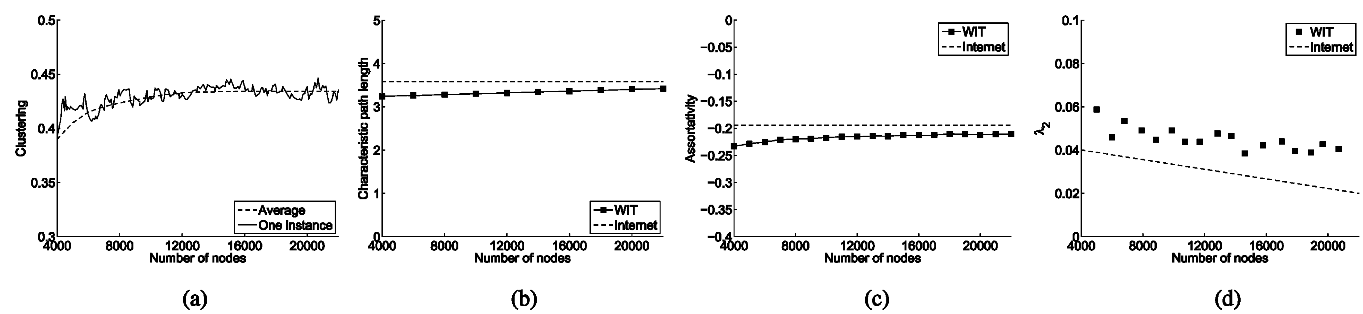

Fig. 17. Evolution of (WIT) compared to the Internet topology observed from the Oregon RouteViews project [45]. (a) Clustering $\gamma(G)$. (b) Characteristic path length L. (c) Assortativity r. (d) Second smallest eigenvalue λ_2.

small, but eventually converge from below to 3.7 and -0.19, respectively, observed in the Internet. In Fig. 17(d), (WIT) 's λ_2 is oscillating but its trend is slowly decreasing as in Fig. 11(f), showing that connectivity properties of (WIT) graphs tend to resemble those of the Internet.

Based on a combination of several metrics examined in this section, one must conclude that graphs constructed by (WIT) are the closest among the compared models to the Internet's evolutionary structure observed by the Oregon RouteViews project [45]. We also believe that (WIT) is a more realistic framework than some of the existing methods as it relies on distributed construction rules and allows each ISP to independently select its peering points based on internal factors such as its wealth, customer base, and QoS requirements that do not depend on the parameters or decisions of other ISPs.

VI. Discussion

We now make several observations and discussion points about the methods proposed in this paper. We also examine possible extensions of this framework to support construction of graphs with nonpower-law degree, different levels of clustering, and hierarchical structure.

A. Why Random Walks?

It may be argued that ISPs do not literally perform random walks on the existing graph and are likely to make connectivity decisions based on a plethora of factors such as technical feasibility, strategic performance considerations, business/market incentives, and price sensitivity; however, this paper merely suggests that their complex decisions (which cannot be easily formalized or even measured in practice) could be replaced with tractable theoretical approximations based on random walks. As with any model, one must weigh the tradeoff between the feasibility of understanding the actual underlying process and replacing it with a high-level approximation. Since the former is impossible in our case, the latter is the only possible solution.

It should also be noted that (WIT) 's short random walks serve the purpose of preserving locality (in the sense that geographically close ISPs are preferred in neighbor selection), which resonates with the well-known fact that the formation process of the Internet topology exhibits geographical preference [4], [6], [19], [24], [28], [31], [34], [54], [57].

B. Possible Extensions

We next briefly mention two possible extensions to (WIT).

1) Degree Distribution and Clustering: For now, we have seen how (WIT) produces graphs with a power-law degree distribution and high clustering that remains invariant as graph size evolves, which are the two key properties of the observed Internet structure. However, as reported by recent studies [21], certain links/nodes are missing from the data sets collected by existing measurement efforts, which implies that graphs inferred from these data sources may not be representative of the true Internet structure, e.g., the degree distribution and/or clustering coefficient of the underlying graph may be different. Therefore, our goal is this section is to explain how (WIT) can be adapted to future Internet topologies that may potentially exhibit drastically different properties from those currently observed in measurement data.

We start with the degree distribution. Note that (WIT) directly correlates degree with the underlying wealth, which means that simply changing the economic model of wealth evolution allows one to achieve a different degree distribution in the graph. As discussed in Section III-C, a variety of wealth distributions ranging from exponential to Gamma to power-law can be generated from so-called *wealth exchange* models [20], where ISPs trade wealth among themselves in a more explicit fashion than modeled earlier in this paper. Thus, by replacing the multiplicative wealth process with an economic exchange model, one can produce a much wider variety of degree distributions using the general framework of (WIT), which is an excellent direction for future work.

Extension of (WIT) in regard to clustering is even more straightforward given the results in the previous section. As shown in Section IV-D, the clustering of (WIT) graphs is decided by walk length l during neighbor selection. If a different clustering coefficient is desired, one can vary l to achieve any value of $E[\gamma(G)]$ between 0.038 and 0.68 according to Fig. 10. These results provide evidence that the proposed model is capable of producing a wide range of clustering coefficients and thus potentially cover the currently unknown graph structure of the true Internet.

2) Hierarchical Structure: It should be noted that inter-AS Internet links are generally believed to fall into one of the two categories depending on the contractual relationship between the neighbors: *peer-peer* and *customer-provider* [17]. However, the model introduced in Section IV-B treats all links equally and does not explicitly create a hierarchical structure commonly attributed to the AS graph. One solution to this issue, which we call *implicit hierarchy*, is to perform processing of WIT topologies for presence of hierarchical structure after the graph has been built and assign to each link a label "peer-peer" or "customer-provider" using heuristics from related work [21], [40]. The second approach, which we call *explicit hierarchy*, is to determine link types during the construction process as part of WIT. We next briefly expand on this idea.

Assume that during join, each node in the graph is labeled as ISP or non-ISP using some random determination (e.g., a stochastic Bernoulli process). As in the regular WIT model, random walks during attachment find nodes with enough wealth to reciprocate link formation, but now the targets of these attachments are limited to ISPs only (i.e., non-ISPs do not attach to each other). If a non-ISP node connects to an ISP, the link is automatically labeled "customer-provider." If two ISPs decide to attach to each other, the type of link may be based on a purely random event (e.g., a coin flip) or the wealth ratios of the two ISPs (e.g., if ISPs' wealth is within 50% of each other, the link is of type "peer-peer" and otherwise, of type "customer-provider"). To leverage this hierarchy more explicitly, a wealth exchange model can be naturally incorporated into this framework. Specifically, during each time step, wealth travels from customers to providers, which emulates payment for services rendered. No wealth is exchanged along peering links, which assumes that peering ISPs provide equal amounts of traffic and service to each other. For non-ISPs to not go bankrupt and for ISP wealth to randomly fluctuate in response to the market, the usual multiplicative wealth model (6) applies to each node before the wealth is distributed along all customer-provider links. More detailed exposition of these algorithms and related simulations are left to future work.

VII. Conclusion

In this paper, we presented an alternative theory of the Internet evolution and developed a new topology generator based on wealth evolution and random walks. We showed that the generated graphs exhibited power-law degree distributions with flexible α and high, nondecreasing clustering coefficients. The characteristic path length and degree-degree correlation of the proposed model were also close to that of the Internet and demonstrated invariance with respect to n. This combination of (WIT)'s properties indicates that the proposed topology algorithm is viable in explaining the structural evolution of the Internet, at least to the extent possible in a very simple model.

Future work includes verifying the proposed algorithm against the complete graph of the Internet and extending WIT to construct hierarchical topologies.

References

[1] W. Aiello and F. R. K. L. ChungL, "A random graph model for massive graphs," in *Proc. ACM STOC*, May 2000, pp. 171–180.
[2] R. Albert and A. Barabási, "Topology of evolving networks: Local events and universality," *Phys. Rev. Lett.*, vol. 85, no. 24, pp. 5234–5237, Dec. 2000.
[3] A.-L. Barabasi and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, no. 5439, pp. 509–512, Oct. 1999.
[4] D. Ben-Avraham, A. Rozenfeld, and R. C. Havlin, "Geographical embedding of scale-free networks," *Physica A*, vol. 330, pp. 107–116, 2003.
[5] S. Bornholdt and H. Ebel, "World wide web scaling exponent from Simon's 1955 model," *Phys. Rev. E*, vol. 64, no. 3, p. 035104, Aug. 2001.
[6] R. Brunet and I. Sokolov, "Evolving networks with disadvantaged long-range connections," *Phys. Rev. E*, vol. 66, p. 026118, 2002.
[7] T. Bu and D. Towsley, "On distinguishing between Internet power law topology generators," in *Proc. IEEE INFOCOM*, Jun. 2002, pp. 638–647.
[8] J. M. Carlson and J. Doyle, "Highly optimized tolerance: A Mechanism for power laws in designed systems," *Phys. Rev. E*, vol. 60, pp. 1412–1427, 1999.
[9] H. Chang and S. J. Willinger, "Internet connectivity at the AS-level: An optimization-driven modeling approach," in *Proc. ACM SIGCOMM Workshop*, 2003, pp. 33–46.
[10] H. Chang and S. J. Willinger, "To peer or not to peer: Modeling the evolution of the Internet's AS-level topology," in *Proc. IEEE INFOCOM*, Apr. 2006, pp. 1–12.

[11] F. K. Chung, *Spectral Graph Theory*, ser. CBMS Regional Conference Series in Mathematics, No. 92. Providence, RI: Amer. Math. Soc., 1997.
[12] F. R. K. Chung and L. Lu, "Connected components in random graphs with given expected degree sequences," *Ann. Combinator.*, vol. 6, no. 2, pp. 125–145, Nov. 2002.
[13] S. N. Dorogovtsev, "Networks with given correlations," 2003, cond-mat/0308336v1.
[14] P. Erdös and A. Rényi, "On random graphs I," *Publication Math. Debrecen*, vol. 6, pp. 290–297, 1959.
[15] A. Fabrikant, E. Koutsoupias, and C. Papadimitriou, "Heuristically optimized tradeoffs," in *Proc. ICALP*, Jul. 2002, pp. 110–122.
[16] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On Power-law Relationships of the Internet Topology," in *Proc. ACM SIGCOMM*, Aug. 1999, pp. 251–262.
[17] L. Gao, "On inferring autonomous system relationships in the Internet," *IEEE/ACM Trans. Netw.*, vol. 9, no. 6, pp. 733–745, Dec. 2001.
[18] K. I. Goh and B. K. Kim, "Fluctuation-driven dynamics of the Internet topology," *Phys. Rev. Lett.*, vol. 88, p. 108701, 2002.
[19] S. Gorman and R. Kulkarni, "Spatial Small Worlds: New Geographic Patterns for an Information Economy," *Environ. Plann. B, Plann. Design*, vol. 31, no. 2, pp. 273–296, 2003.
[20] A. Gupta, "Models of wealth distributions—A perspective," in *Econophys. Sociophys.*, B. Chakrabarti, A. Chakraborti, and A. Chatterjee, Eds. Hoboken, NJ: Wiley-VCH, 2007, pp. 161–190.
[21] Y. He, G. Siganos, M. Faloutsos, and S. V. Krishnamurthy, "A systematic framework for unearthing the missing links: Measurements and impact," in *Proc. USENIX/SIGCOMM NSDI*, Apr. 2007, pp. 3632–3637.
[22] S. Jaiswal, A. Rosenberg, and D. Towsley, "Comparing the structure of power-law graphs and the Internet as graph," in *Proc. IEEE ICNP*, Oct. 2004, pp. 294–303.
[23] C. Jin, Q. Chen, and S. Jamin, "INET: Internet topology," Univ. Michigan, 2002, Tech. Rep. CSE-TR-456-02.
[24] S. Jin and A. Bestavros, "Small-world Internet topologies: Possible causes and implications on scalability of end-system multicast," presented at the IEEE/ACM Int. Symp. Model., Anal. Simul. (MASCOTS), Oct. 2003.
[25] K. Klemm and V. M. Eguiluz, "Highly clustered scale-free networks," *Phys. Rev. E*, vol. 65, p. 036123, 2002.
[26] P. L. Krapivsky and S. Redner, "Finiteness and fluctuations in growing networks," *J. Phys. A*, vol. 35, pp. 9517–9534, 2002.
[27] P. L. Krapivsky, G. J. Rodgers, and S. Redner, "Degree distributions of growing random networks," *Phys. Rev. Lett.*, vol. 86, no. 23, pp. 5401–5404, 2001.
[28] A. Lakhina, J. W. Byers, and M. C. Matta, "On the geographic location of Internet resources," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 8, pp. 934–948, Aug. 2003.
[29] M. Levy and S. Solomon, "Dynamical explanation for the emergence of power law in a stock market model," *Int. J. Mod. Phys. C*, vol. 7, pp. 65–72, 1996.
[30] M. Levy and S. Solomon, "Power laws are logarithmic Boltzmann laws," *Int. J. Mod. Phys. C*, vol. 7, pp. 595–601, 1996.
[31] X. Li and G. Chen, "A local-world evolving network model," *Physica A*, vol. 328, pp. 274–286, 2003.
[32] H. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani, "iPlane: An information plane for distributed services," in *Proc. OSDI*, 2006, pp. 367–380.
[33] D. Magoni and J. Pansiot, "Analysis of the autonomous system network topology," *Comput. Commun. Rev.*, vol. 31, pp. 26–37, 2001.
[34] A. Medina, I. Matta, and J. Byers, "On the origin of power-laws in Internet topologies," *Comput. Commun. Rev.*, vol. 30, pp. 18–28, 2000.
[35] M. Mihail and C. H. Papadimitriou, "On the eigenvalue power law," in *Proc. Random*, Sep. 2002, pp. 254–262.
[36] M. Mihail and N. Visnoi, "On generating graphs with prescribed degree sequences for complex network modeling applications," presented at the ARACNE, Sep. 2002.
[37] B. Mohar, "The Laplacian spectrum of graphs," *Graph Theory, Combinator., Appl.*, vol. 2, pp. 871–898, 1991.
[38] M. Molloy and B. Reed, "A critical point for random graphs with a given degree sequence," *Random Struct. Algor.*, vol. 6, pp. 161–179, 1995.
[39] M. E. J. Newman, "Assortative mixing in networks," *Phys. Rev. Lett.*, vol. 89, no. 20, p. 208701, Oct. 2002.
[40] R. V. Oliveira, D. Pei, W. Willinger, and B. Z. Zhang, "In Search of the elusive ground truth: the Internet's AS-level connectivity structure," in *Proc. ACM SIGMETRICS*, Jun. 2008, pp. 217–228.
[41] V. Pareto, *Cours D'Ѐconomique Politique*. London, U.K.: Macmillan, 1897.
[42] A. R. Puniyani, R. M. Lukose, and B. A. Huberman, "Intentional walks on scale free small worlds," 2001, cond-mat/0107212.
[43] S. Redner, "How popular is your paper? An empirical study of the citation distribution," *Eur. Phys. J. B*, vol. 4, pp. 131–134, 1998.
[44] "RIPE," [Online]. Available: http://www.ripe.net/
[45] "Route Views project," [Online]. Available: http://www.routeviews.org/
[46] P. O. Seglen, "The skewness of science," *J. Amer. Soc. Inf. Sci.*, vol. 43, pp. 628–638, 1992.
[47] Y. Shavitt and E. Shir, "DIMES: Let the Internet measure itself," *Comput. Commun. Rev.*, vol. 35, no. 5, pp. 71–74, Oct. 2005.
[48] H. A. Simon, "On a class of skew distribution functions," *Biometrika*, vol. 42, no. 3–4, pp. 425–440, Dec. 1955.
[49] H. A. Simon, "Some further notes on a class of skew distribution functions," *Inf. Contr.*, vol. 3, pp. 80–88, 1960.
[50] "Skitter," [Online]. Available: http://www.caida.org/
[51] D. Sornette, "Multiplicative processes and power laws," *Phys. Rev. E*, vol. 57, pp. 4811–4813, 1998.
[52] H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, "Network topology generators: Degree-based vs. structural," in *Proc. ACM SIGCOMM*, Aug. 2002, pp. 147–159.
[53] D. Vukadinovic, P. Huang, and T. Erlebach, "A spectral analysis of the Internet topology," ETH, 2001, Tech. Rep. TIK-NR.118.
[54] C. Warren, L. Sander, and I. Sokolov, "Geography in a scale-free network model," *Phys. Rev. E*, vol. 66, p. 056105, 2002.
[55] D. J. Watts, *Small Worlds: The Dynamics of Networks between Order and Randomness*. Princeton, NJ: Princeton Univ. Press, 1999.
[56] D. J. Watts and S. Strogatz, "Collective dynamics of small world networks," *Nature*, vol. 393, pp. 440–442, Jun. 1998.
[57] B. Waxman, "Routing of multipoint connections," *IEEE J. Sel. Areas Commun.*, vol. 6, no. 9, pp. 1617–1622, Dec. 1988.
[58] R. W. Wolff, *Stochastic Modeling and the Theory of Queues*. Englewood Cliffs, NJ: Prentice-Hall, 1989.
[59] S. Yook, H. Jeong, and A. Barabási, "Modeling the Internet's large-scale topology," *Natl. Acad. Sci.*, vol. 99, no. 21, pp. 13382–13386, Oct. 2002.
[60] S. Zhou and R. J. Mondragon, "Accurately modeling the Internet topology," *Phys. Rev. E*, vol. 70, no. 6, p. 066108, Dec. 2004.

Xiaoming Wang (S'04) received the B.S. degree in computer science and the M.S. degree in electronic engineering from Beijing University of Posts and Telecommunications, Beijing, China, in 1999 and 2002, respectively.

He is currently working toward the Ph.D. degree at Texas A&M University, College Station. During 2002–2003, he was with the Samsung Advanced Institute of Technology, South Korea. His research interests include peer-to-peer systems, probabilistic analysis of computer networks, and topology modeling.

Mr. Wang has been a Student Member of the Association for Computing Machinery (ACM) since 2007.

Dmitri Loguinov (S'99–M'03–SM'08) received the B.S. degree (with honors) in computer science from Moscow State University, Moscow, Russia, in 1995, and the Ph.D. degree in computer science from the City University of New York, New York, in 2002.

Between 2002 and 2007, he was an Assistant Professor with the Department of Computer Science at Texas A&M University, College Station. He is currently a tenured Associate Professor and Director of the Internet Research Lab (IRL) in the same department. His research interests include peer-to-peer networks, video streaming, congestion control, Internet measurement, and modeling.

Dr. Loguinov is a Senior Member of the Association for Computing Machinery (ACM).