Tomography of scale-free networks and shortest path trees

Tomer Kalisky,[1,*] Reuven Cohen,[2,†] Osnat Mokryn,[3] Danny Dolev,[4] Yuval Shavitt,[3] and Shlomo Havlin[1]
[1]*Minerva Center and Department of Physics, Bar-Ilan University, 52900 Ramat-Gan, Israel*
[2]*Department of Computer Science and Applied Mathematics, Weizmann Institute of Science, Rehovot, Israel*
[3]*Department of Electrical Engineering-Systems, Tel-Aviv University, Tel-Aviv, Israel*
[4]*School of Engineering and Computer Science, Hebrew University, Jerusalem, Israel*

(Received 10 May 2006; revised manuscript received 18 October 2006; published 15 December 2006)

In this paper we model the tomography of scale-free networks by studying the structure of layers around an arbitrary network node. We find, both analytically and empirically, that the distance distribution of all nodes from a specific network node consists of two regimes. The first is characterized by rapid growth, and the second decays exponentially. We also show analytically that the nodes degree distribution at each layer exhibits a power-law tail with an exponential cutoff. We obtain similar empirical results for the layers surrounding the root of shortest path trees cut from such networks, as well as the Internet.

DOI: 10.1103/PhysRevE.74.066108 PACS number(s): 89.75.Hc, 89.20.Ff

I. INTRODUCTION

Many systems in nature have a weblike structure consisting of nodes and links connecting them, and may be described as "networks." Examples for such networks may be found in computer science (e.g., the Internet), biology, and sociology [1–3]. It was recently realized that many networks in nature exhibit a "scale-free" degree distribution, meaning that the number of nodes with degree k decays as a power law $P(k) \sim k^{-\lambda}$ where $2 < \lambda < 3$ [4]. These types of networks were found to have anomalous properties, for example, they were found to have a small radius and also to be resilient to random breakdown of nodes or links (e.g., Refs. [5,6]).

In this paper we investigate the structure of scale-free networks, and suggest an analytical derivation for some of their characteristics. We verify our findings with simulations, and compare them to real router-level Internet data obtained from the Lucent mapping project [7]. We observe similar results for shortest path trees cut from these networks. Given a specific real-world network, this analysis can be used as a tool for preliminary evaluation of the different models describing its structure.

Specifically, we introduce here a first attempt to analytically evaluate the number and degree distribution of nodes as a function of their distance from a specific network node. We show that the distribution of the number of nodes as a function of the distance consists of two regimes. The first regime is characterized by a very rapid growth in the number of nodes, leading to a maximum after a rather short distance. Then a second regime starts, characterized by an exponential decay in the number of nodes. We also show that the tail of the node degree distribution at each layer follows a power law with an exponential cutoff. Although this cutoff was observed in previous works [5,8,9], it was not characterized. Here, we show analytically that it is of an exponential nature, and decays with the distance from the origin point.

We further investigate the inner structure of scale-free networks by studying shortest path trees cut from them. The study of such trees is important, as they may represent the structure of multicast trees in the Internet. Previous results [10] show that the degree distribution of such trees exhibits a power-law tail. Here we investigate their layer structure and the distribution of number of nodes as a function of their distance from the root of the tree. Using simulations and real router-level Internet data, we show that the trees exhibit a layer behavior similar to the network they were cut from.

Note that our results conform to the "ultrasmall world" phenomenon described in Refs. [5,8,9], which means that the average distance between every two nodes in a scale-free network scales as $\ln \ln N$, where N is the total number of nodes in the network. As a result, the average distance in scale-free networks is much smaller than in traditional random Erdős-Rényi networks [11], for which the average distance scales as the logarithm of the number of nodes. However, as opposed to the ultrasmall world phenomenon, which is extremely hard to observe in reasonable size graphs, the results we give here are applicable for finite size networks and allow a numerical calculation of several aspects of the layer structure. As we show, these analytical results are in agreement with both simulations and real data.

A possible application of our results is to use them as an indicative tool for the assessment of the accuracy of modeling mechanisms of the Internet and multicast trees, in addition to other metrics widely used today (e.g., the network degree distribution). This understanding may also aid in the development of better routing and transport layer protocols, as well as structural sensitive application level algorithms, such as placement problems.

II. BACKGROUND

A. Graph generation

As mentioned above, recent studies have shown that many real world systems may be described as scale-free networks. That is, their degree distribution follows a power law, $P(k) = ck^{-\lambda}$, where c is an appropriate normalization factor, and λ is the exponent of the power law.

Several techniques for generating such scale-free graphs were introduced [12–14]. The model we use here is usually

*Electronic address: kaliskt@mail.biu.ac.il
†Electronic address: cohenr@bu.edu

termed the *configuration model* (CM) [15] or the *generalized random graph* (GRG) model [16], and is used for generating random graphs with a given degree distribution. The construction method is as follows: For a given graph size N, a set V of N nodes is generated. The degree sequence is determined by randomly choosing a degree for each one of the nodes from the given degree distribution. At first, each node is given a predetermined number of outgoing links. We denote the set of these links (which are still unmatched) by C, and term them *open connections*. Then, we randomly match each open link with another open link to create an edge connecting two nodes. The set of edges is denoted by E.

Notice that this procedure may lead to a multigraph, i.e., a graph with self-loops and multiple edges between two sites. However, the fraction of such links is statistically insignificant and we ignore them [17]. Notice also that the order by which we match the open links is insignificant as long as the choice of the matching link is random.

There are other methods for generating scale-free networks, such as the Barabási-Albert (BA) model [12,18], which involves growth and preferential attachment. However, while in the BA model the graph degree distribution function emerges only at the end of the process, in the configuration model the distribution is known *a priori*, thus enabling us to use it in our analysis during the construction of the graph.

B. Graph radius and distribution cutoff

Recent work [5,19] has shown that the radius r (defined as the *average* distance between the highest degree node and all other nodes) of scale-free graphs with $2<\lambda<3$ is extremely small and scales as $r\sim\ln\ln N$. The meaning of this is that even for very large networks, finite size effects must be taken into account, because algorithms for traversing the graph will arrive at the network edge after a small number of steps.

Since the scale-free distribution has no typical degree, its behavior is influenced by externally imposed cutoffs, i.e., minimum and maximum values for the allowed degrees, k. The fraction of sites having degrees above and below the threshold is assumed to be zero. The lower cutoff, m, is usually chosen to be of order $O(1)$, since it is natural to assume that in real world networks many nodes of interest have only one or two links. The upper cutoff, K, can also be enforced externally (say, by the maximum number of links that can be physically connected to a router). However, in situations where no such cutoff is imposed, we assume that the system has a natural cutoff.

To estimate the natural cutoff of a network, we assume that the network consists of N nodes, each of which has a degree randomly selected from the distribution $P(k)=ck^{-\lambda}$. An estimate of the average value of the largest of the N nodes can be obtained by looking for the smallest possible tail that contains a single node on the average [6],

$$\sum_{k=K}^{\infty} P(k) \approx \int_K^{\infty} P(k)dk = 1/N. \tag{1}$$

Solving the integral yields $K\approx mN^{1/(\lambda-1)}$, which is the approximate natural upper cutoff of a scale-free network [6,20,21].

In the rest of this paper, in order to simplify the analysis presented, we will assume that this natural cutoff is imposed on the distribution by the exponential factor $P(k)=ck^{-\lambda}e^{-k/K}$.

III. TOMOGRAPHY OF SCALE-FREE NETWORKS

In this section we study the statistical behavior of layers surrounding the maximally connected node in the network. First, we describe the process of generating the network and define our terminology. Then, we analyze the degree distribution at each layer surrounding the maximally connected node.

A. Model description

We base our construction on the CM [13], also described in Sec. II A. The construction process tries to gradually "expose" the network [5,19], and is forcing a hierarchy on the CM, thus enabling us to define layers in the graph.

We start by setting the number of nodes in the network, N. We then choose the nodes degrees according to the scale-free distribution function $P(k)=ck^{-\lambda}$, where $c\approx(\lambda-1)m^{\lambda-1}$ is the normalizing constant and k is in the range $[m,K]$, for some chosen minimal degree m and the natural cutoff $K=mN^{1/(\lambda-1)}$ of the distribution [6,20].

At this stage each node in the network has a given number of outgoing links, which we term *open connections*, according to its chosen degree. Using our definitions in Sec. II A, the set of edges in E is empty at this point, while the set of open links in C contains all unmatched outgoing links in the graph.

We proceed as follows: we start from the maximal degree node, which has a degree K, and connect it randomly to K available open connections, thus removing these open connections from C [see Fig. 1(a)]. We have now exposed the first *layer* (or *shell*) of nodes, indexed as $l=1$. We now continue to fill out the second layer $l=2$ in the same way: We connect all open connections emerging from nodes in layer No. 1 to randomly chosen open connections. These open connections may be chosen from nodes of layer No. 1 (thus creating a loop) or from other links in C. We continue until all open connections emerging from layer No. 1 have been connected, thus filling layer $l=2$ [see Fig. 1(b)]. Generally, to form layer $l+1$ from an arbitrary layer l, we randomly connect all open connections emerging from layer l to other open connections, either those emerging from l or chosen from the other links in C [see Fig. 1(c)]. Note, that when we have formed layer $l+1$, layer l has no more open connections. The process continues until the set of open connections, C, is empty.

B. Analysis

We proceed now to evaluate the probability for nodes with degree k to reside outside the first l layers, denoted by $P_l(k)$.

The number of open connections outside layer No. l, is given by

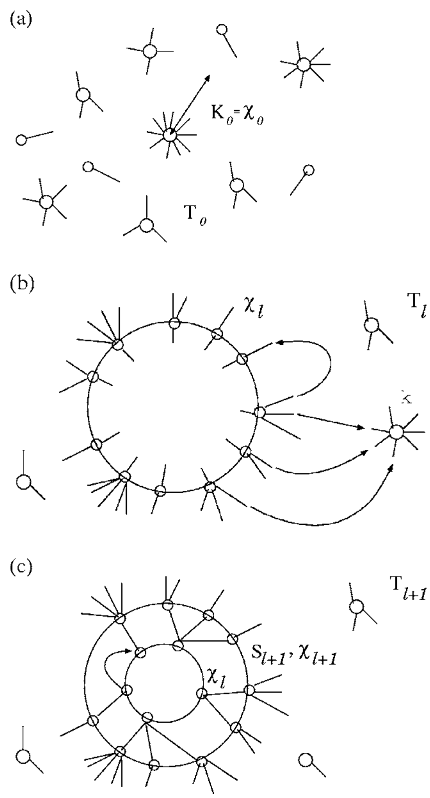

FIG. 1. Illustration of the exposure process. The large circles denote exposed layers of the giant component, while the small circles denote individual sites. The sites outside the circles have not been reached yet. (a) We begin with the highest degree node and fill out layer No. 1. (b) In the exposure of layer No. $l+1$ any open connection emerging from layer No. l may connect to any open node (T_l connections) or loop back into layer No. l (χ_l connections). (c) The number of connections emerging from layer No. $l+1$ is the difference between T_l and T_{l+1} after reducing the incoming connections S_{l+1} from layer No. l.

$$T_l = N\sum_k kP_l(k). \tag{2}$$

Thus, we can define the probability that a detached node with degree k will be connected to an open connection emerging from layer l by $\frac{k}{\chi_l+T_l}$, where χ_l is the number of open connections emerging from layer l [see Fig. 1(b)].

Therefore, the conditional probability for a node with degree k to be also outside layer $l+1$, given that it is outside layer l, is the probability that it does not connect to *any* of the χ_l open connection emerging from layer l, that is

$$P(k,l+1|l) = \left(1 - \frac{k}{\chi_l + T_l}\right)^{\chi_l} \approx \exp\left(-\frac{k}{1+\frac{T_l}{\chi_l}}\right), \tag{3}$$

for large enough values of χ_l.

Thus, the probability that a node of degree k will be outside layer No. $l+1$ is

$$P_{l+1}(k) = P_l(k)P(k,l+1|l) = P_l(k)\exp\left(-\frac{k}{1+\frac{T_l}{\chi_l}}\right). \tag{4}$$

Thus we derive the exponential cutoff

$$P_l(k) = P(k)\exp\left(-\frac{k}{K_l}\right), \tag{5}$$

where

$$\frac{1}{K_{l+1}} = \frac{1}{K_l} + \frac{1}{1+\frac{T_l}{\chi_l}}. \tag{6}$$

An alternate method for deriving the above relationship is given in the Appendix.

Now let us find the behavior of χ_l and S_l, where S_{l+1} is the number of links incoming to the $l+1$ layer from layer No. l (and approximately equals N_{l+1}, the number of nodes in the $l+1$th layer [22]). The number of incoming connections to layer $l+1$ equals the number of connections emerging from layer l, minus the number of connections looping back into layer No. l. The probability for a connection to loop back into layer l is

$$P(\text{loop}|l) = \frac{\chi_l}{\chi_l + T_l}, \tag{7}$$

and therefore

$$S_{l+1} = \chi_l\left(1 - \frac{\chi_l}{\chi_l + T_l}\right). \tag{8}$$

The number of connections emerging from all the nodes in layer No. $l+1$ is $T_l - T_{l+1}$. This includes the number of incoming connections from layer l into layer $l+1$, which is equal to S_{l+1}, and the number of outgoing connections χ_{l+1}. Therefore

$$\chi_{l+1} = T_l - T_{l+1} - S_{l+1}. \tag{9}$$

At this point we have the following relations: $T_{l+1}(K_{l+1})$ [Eqs. (2) and (5)], $S_{l+1}(\chi_l, T_l)$ [Eq. (8)], $K_{l+1}(K_l, \chi_l, T_l)$ [Eq. (6)], and $\chi_{l+1}(T_l, T_{l+1}, S_{l+1})$ [Eq. (9)]. These relations may be solved numerically. Note that approximate analytical results for the limit $N \to \infty$ can be found in Refs. [5,8,19,23].

C. Empirical results on networks

Figure 2 shows results from simulations (filled symbols) for the number of nodes at layer l, which can be seen to be in agreement with the analytical curves of S_l (lines). We can see that starting from a given layer $l=L$ the number of nodes decays exponentially. We believe that the layer index L is related to the radius of the graph [5,19]. It can be seen that S_l is a good approximation for the number of nodes at layer l. This is true in cases when only a small fraction of sites in each layer l have more than one incoming connection. An example for this case is when $m=1$ so that most of the sites in the network have only one connection. Figure 3 shows results for $P_l(k)$ with similar agreement [24]. Note the exponential cutoff which becomes stronger with l.

It is important to note that the simulation results give the probability distribution for the nodes of the giant component (the largest connected component of the network), while the

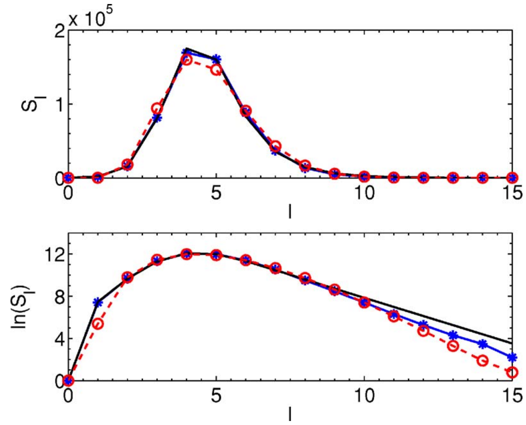

FIG. 2. (Color online) Approximate number of nodes (S_l) vs layer index l for a network with $N=10^6$ nodes, $\lambda=2.85$, and $m=1$. Filled symbols represent simulation results while black lines are a numerical solution for the derived recursive relations. Empty symbols represent the gamma function with $\alpha=10.576$ and $\beta=2.225$. Bottom: from the semilog plot we see that there is an exponential decay of S_l for layers $l>L$ starting from a given layer L which we believe is related to the radius of the graph.

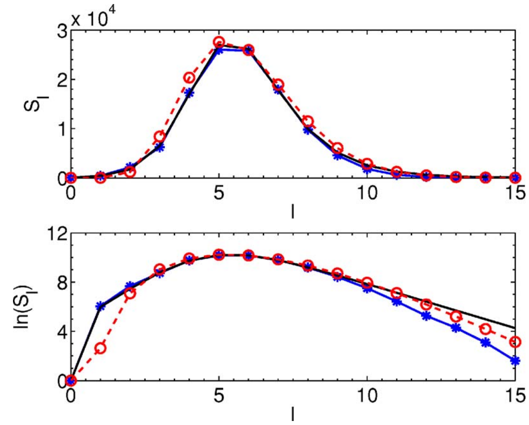

FIG. 4. (Color online) Number of nodes at each layer for a router level cut of the Internet with $N=112\,969$ nodes (LC topology). Analytical reconstruction for S_l is done with $\lambda=3$ and $m=1$. The fit to the gamma function was done with parameters $\alpha=9.99$ and $\beta=1.701$.

analytical reconstruction gives the probability distribution for the whole graph. This may explain the difference in the probability distributions for lower degrees: many low-degree nodes are not connected to the giant component and therefore the probability distribution derived from the simulations is smaller for low degrees. However, as can be seen from Fig. 3, the finite size effects due to the difference between the giant component and the whole network do not affect the (power-law) tail of the distribution function.

Figures 4 and 5 show the same analysis for a cut of the Internet at router-level (Lucent mapping project [7], LC topology, see Table I). The actual probability distribution is not a pure power law, rather it can be approximated by $\lambda=2.3$ for small degrees and $\lambda=3$ at the tail. Our analytical reconstruction of the layer statistics assumes $\lambda=3$, because the tail of a power-law distribution is the important factor in determining properties of the system. This method results in a good reconstruction for the number of nodes in each layer, and a qualitative reconstruction of the probability distribution in each layer.

Previous measurements [25] have shown that the number of nodes at each layer from a chosen node on the Internet follows a gamma distribution:

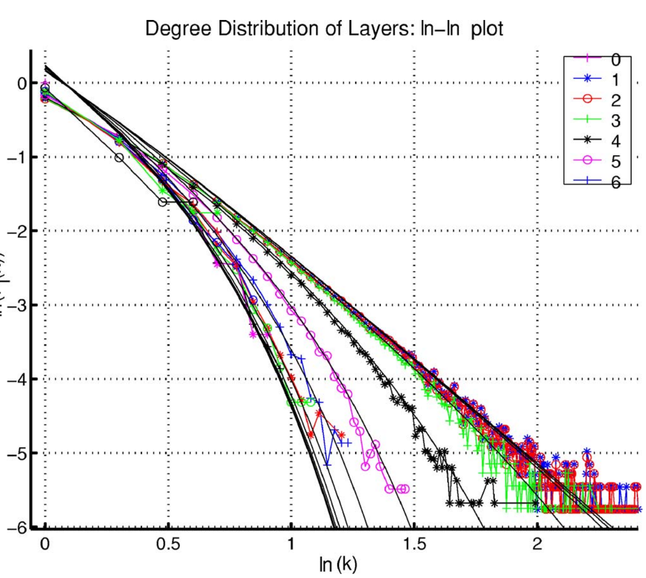

FIG. 3. (Color online) Log-log plot of $P_l(k)$ for different layers $l=0,1,2,\ldots$, for a network with $N=10^6$ nodes, $\lambda=2.85$, and $m=1$. Symbols represent simulation results while black lines are a numerical solution for the derived recursive relations.

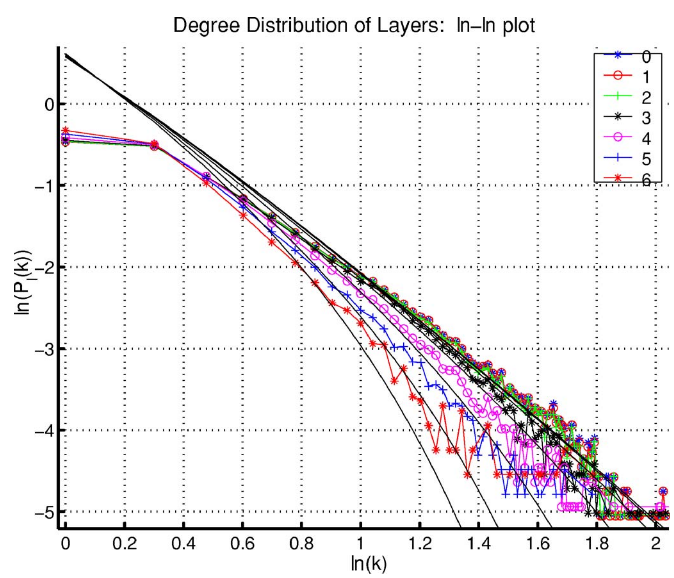

FIG. 5. (Color online) Log-log plot of $P_l(k)$ for different layers $l=0,1,2,\ldots$, for a router level cut of the Internet with $N=112\,969$ nodes (LC topology). Qualitative analytical reconstruction is done with $\lambda=3$, and $m=1$.

TABLE I. The types of underlying network topologies used for the analysis of the tomography of networks and multicast trees (BA, Barabási-Albert; CM, configuration model). For all networks generated by the Barabási-Albert model we used $m_0=6$ and $q=0$. Note that for the real Internet data the degree distribution is not a pure power law, rather, it can be approximated by $\lambda=2.3$ for small degrees and $\lambda=3$ at the tail.

Name	Type	Parameters	No. of Nodes	Avg. Node degree	Power-law exponent (λ)
VS	Generated (BA)	$m=1$; $p\in 0:0.05:0.5$	10000	1.99–3.98	4–3.5
LS	Generated (BA)	$m=3$; $p\in 0:0.05:0.5$	10000	5.98–12.04	3.3333–3.16667
IS	Generated (BA)	$m=2$; $p\in 0:0.05:0.5$	10000	3.99–7.9	3.5–3.25
IS	Generated (BA)	$m=1.5,2$; $p=0.1$	50000;100000	3.3;4.4	3.6;3.45
IS	Generated (BA)	$m=1.5$, $p=0.1$	1000000	3.335	3.6
CM	Generated (CM)	$\lambda=2.85$; $m=1$	1000000	2.17	2.85
CM	Generated (CM)	$\lambda=2.5$; $m=1$	1000000	2.97	2.5
LC	Real data	Internet	112969	3.2	3 (approximate)

$$f_{\alpha,\beta}(l)=\frac{1}{\Gamma(\alpha)}\beta^{\alpha}l^{(\alpha-1)}\exp(-\beta l), \qquad (10)$$

where the free parameters α and β may be found by the mean $E(f)$ and variance $V(f)$ of the distribution, $\beta=\frac{E(f)}{V(f)}$, $\alpha=\frac{[E(f)]^2}{V(f)}$. As can be seen in Figs. 2 and 4, the gamma distribution function gives a reasonable approximation to our results.

Note that in our model, large degree nodes of the network tend to reside in the lower layers, while the layers further away from the source node are populated mostly by low degree nodes [10]. This implies that the tail of the distribution affects the lower layers, while the distribution function for lower degrees affects the outer layers. Thus the deviations in the analytical reconstruction of the number of nodes per layer for the higher layers may be attributed to the deviation in the assumed distribution function for low degrees (that is $\lambda=3$ instead of $\lambda=2.3$). Note also that our model does not take into account the correlations in node degrees, which were observed in the Internet [26], and hierarchical structures [27]. All these considerations may explain the deviation of our measurements from the model predictions (see also Refs. [28–31]).

As a final cautionary point, it should be mentioned that some recent papers [32–35] discussing the influence of the measurement bias on the measured topology suggest that the measured node degrees are not the actual ones. Usually, traceroute data is used to study the Internet topology. It is known that this kind of measurement tends to miss a high number of links far from the measurement center and therefore to underestimate the degree of distant nodes. This may lead to a picture giving the central nodes a higher degree than the peripheral ones. Therefore, it is possible that the empirical data does not represent a complete view of the Internet, but rather a combination of the Internet structure with measurement bias.

IV. TOMOGRAPHY OF SHORTEST PATH TREES (EMPIRICAL RESULTS)

In this section, we study the tomography of shortest path trees embedded in scale-free networks.

A motivation for studying shortest path trees is that they may be seen as analogs for Internet multicast trees [36]. Multicast trees are trees built over the Internet topology and used to transmit data (in particular multimedia) from one source to many destinations (clients). Each new client connecting to this tree is connected through its shortest path to the source. The benefit of multicast trees is that data is transmitted only once to each node in the tree, whereas if communication was established with each client individually, it is highly probable that each node along the path would receive each packet multiple times, once for each shortest path going through it.

Hence, we study the structure and characteristics of the depth rings around the root node of shortest path trees embedded in the network. We will show that in this aspect, shortest path trees have similar characteristics to the networks they were cut from.

A. Topology and tree generation

Our method for producing trees is the following. First, we generate power-law topologies based on the CM (described in Sec. II A) and the BA model [18].

The BA model is based on two principles: "growth" and "preferential attachment." The model specifies four parameters, m_0, m, p, and q, where m_0 is the initial number of nodes, and m is the initial connectivity (degree) of each node. When a link is added (or rewired), one of its endpoints is chosen randomly, and the other is chosen with a probability that is proportional to the node degree. This "preferential attachment" reflects the fact that new links often attach to popular (high-degree) nodes. The growth model is the following: we start with m_0 detached nodes. Then, with probability p, m new links are added to the topology. With probability q, m links are rewired, and with probability $1-p-q$ a new node with m links is added. The process repeats itself until we reach the required number of nodes. Note that m, p, and q determine the average degree of the nodes and the power-law scaling exponent—see Ref. [18] for details [37].

We created a vast range of topologies, but concentrated on several parameter combinations that can be roughly described as very sparse (VS), Internet like sparse (IS) and less sparse (LS). We also created a topology using the CM, described in Sec. II A. Table I summarizes the main characteristics of the topologies used in this paper.

From these underlying topologies, we create the trees in the following manner. For each predetermined size of client population we choose a root node and a set of clients. Using Dijkstra's algorithm we build the shortest path tree from the root to the clients. To create a set of trees that realistically resemble Internet trees, we define four basic tree types. These types are based on the degree of the root node and the client nodes. In order to represent the case of a tree rooted at a big ISP site, we choose a root node with a high degree with respect to the underlying topology. Then, we either choose the clients as low-degree nodes, or at random as a control group. Note, that due to the characteristic of the power-law distribution, a random selection of a node has a high probability of choosing a low-degree node. The next two tree types have a low-degree root, which corresponds to a multicast session from an edge router. Again, the two types differ by the clients degrees, which are either low, or picked at random.

The tree client population is chosen at the range [50,4000] for the 10 000 node generated topology, [50,10 000] for the 100 000 node generated topology, and 100 000 clients for all 10^6 nodes topologies [38].

There are two underlying assumptions made in the tree construction. The first is that the multicast routing protocol delivers a packet from the source to each of the destinations along a shortest path tree. This scenario conforms with current Internet routing (it should be noted that the actual BGP routing in the Internet is not guaranteed to be shortest path, and contains several other policy considerations). For example, IP packets are forwarded based on the reverse shortest path, and multicast routing protocols such as source specific multicast [39] deliver packets along the shortest path route. In addition, we assume that client distribution in the tree is uniform [7,40–42].

B. Tree characteristics—simulations

In previous work we have shown that shortest path trees cut from a power-law network topology obey a similar power law for the degree distribution, as well as the subtrees sizes [10]. The results were shown to hold for all trees cut from all generated topologies, even for trees as small as 200 nodes. Here we investigate the tomography of the trees, and look at the degree distribution of nodes at different depth rings around the root, i.e., tree layers.

Figure 6 shows the number of nodes and the degree distribution at each layer l from the root of a shortest path tree cut from a CM topology (10^6 nodes). The root was chosen with a high degree, and the clients with a low degree. It can be seen that the degree distribution obeys a power law with an exponential cutoff, similar to what was observed in networks. This phenomenon was found to be stable regardless of the tree type, and the client population size [43]. As seen from Fig. 6, the number of nodes at each layer can be approximated by a gamma distribution [25], similar to what was found in networks (Sec. III C).

Figure 7 shows the same analysis for a shortest path tree cut from a BA topology (10^6 nodes). Again, the degree distributions for the different layers obey a power law with progressively stronger exponential cutoffs, similar to what was observed in networks.

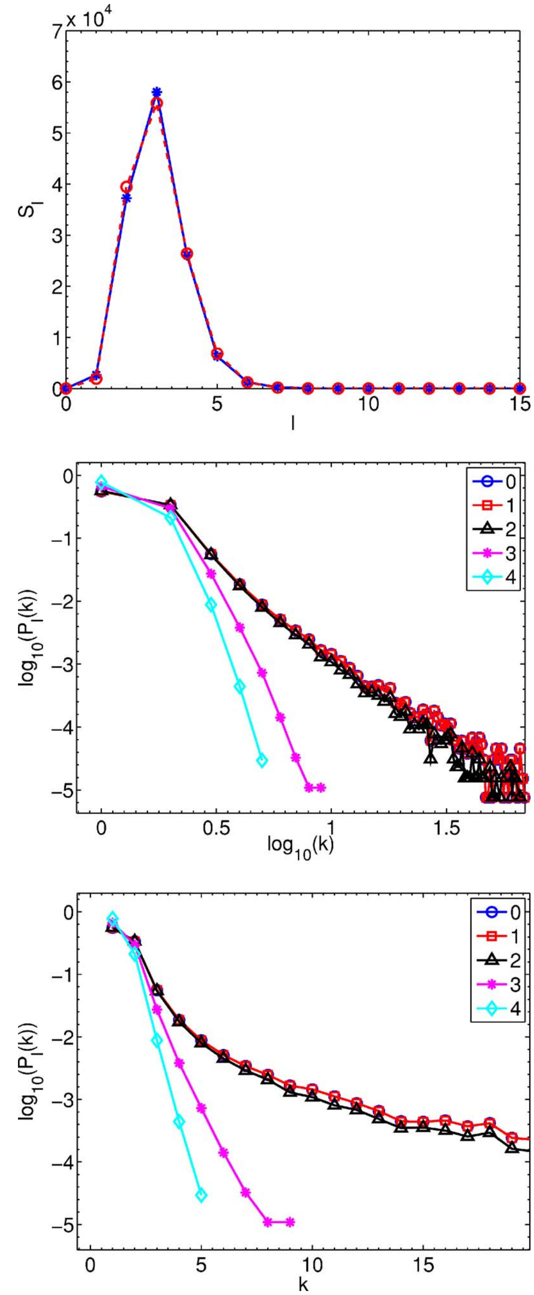

FIG. 6. (Color online) Tomography of a 100 000 client multicast tree cut from a CM topology of 10^6 nodes, with $m=1$ and $\lambda=2.5$. The root was chosen with a high degree, and the clients were chosen randomly. Top: Number of nodes at each layer (filled symbols). Empty symbols represent a gamma distribution function with $\alpha=10.302$ and $\beta=3.424$. Middle: Log-log plot of the degree distributions $P_l(k)$ for different layers $l=0,1,2,\ldots$. Bottom: A semilog plot of the same distributions. The degree distributions exhibit a power-law tail ($\lambda_{\text{tree}}\approx 2.5$) with an exponential cutoff which is becoming stronger at each layer.

In order to understand the exact relationship of the degree of a node on its layer, we plot the number of nodes of each degree at different layers. Figures 8 and 9 show the distribu-

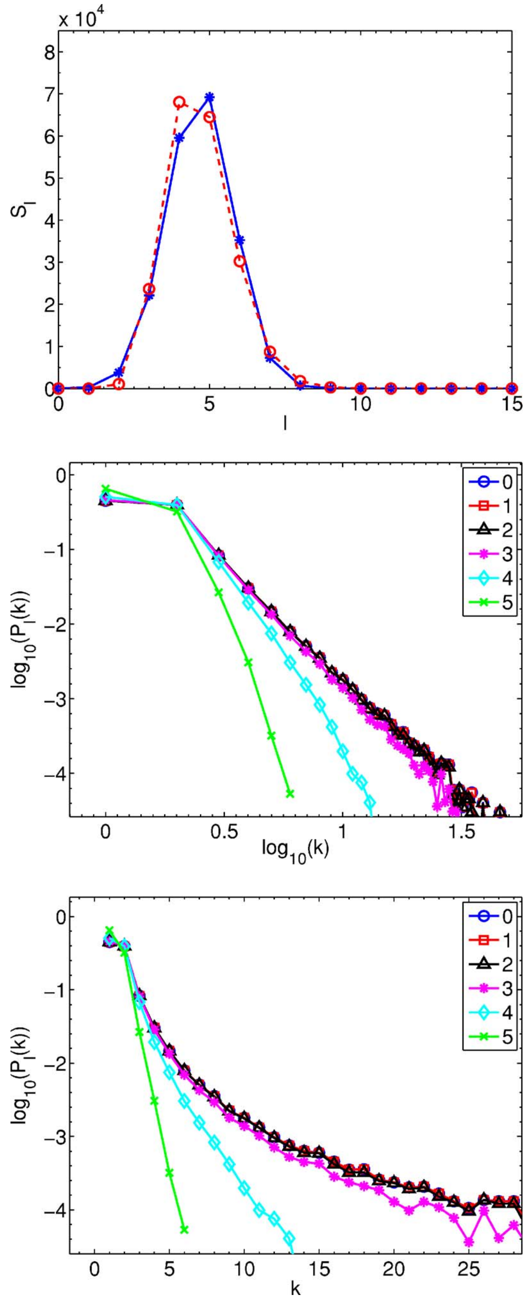

FIG. 7. (Color online) Tomography of a 100 000 client multicast tree cut from a BA topology of 10^6 nodes with $m_0=6$, $m=1.5$, $p=0.1$, and $q=0$. The root was chosen with a high degree, and the clients were chosen randomly. Top: Number of nodes at each layer (filled symbols). Empty symbols represent a gamma distribution function with $\alpha=18.215$ and $\beta=3.896$. Middle: Log-log plot of the degree distributions $P_l(k)$ for different layers $l=0,1,2,\ldots$. Bottom: A semilog plot of the same distributions. The degree distributions exhibit a power-law tail ($\lambda_{\text{tree}}\approx 3.2$) with an exponential cutoff which is becoming stronger at each layer.

tion of the distance of degree two ($k=2$) nodes, and high degree nodes ($k\geqslant 7$ or $k\geqslant 8$) for CM and BA topologies. It can be seen that for both models the high-degree nodes tend to reside much closer to the root than the low-degree nodes, and in adjacent layers [44,45,48].

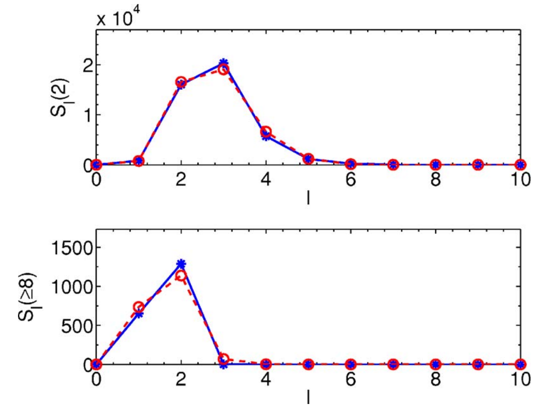

FIG. 8. (Color online) Top: Number of degree two nodes ($k=2$) at each layer of a multicast tree cut from CM topology with $\lambda=2.5$ and 10^6 nodes. Empty symbols represent a fit to the gamma distribution with $\alpha=11.173$ and $\beta=3.984$. Bottom: Number of high-degree nodes ($k\geqslant 8$) at each layer of a multicast tree cut from the same underlying topology. Empty symbols represent the corresponding gamma distribution with $\alpha=12.25$ and $\beta=7.36$. It can be seen that the high-degree nodes tend to reside much closer to the root than the low-degree nodes.

We also checked the distribution of the lengths of the paths to the clients. Our results show that the less connected the underlying topology, the higher is the average tree cut from the topology. For a 10 000 node underlying topology with an average degree of 3 and higher, the height of the trees was not more than 10. On an underlying topology of 100 000 nodes, the height of the trees was not more than 12. In accordance with our findings of a core of high-degree nodes, the trees were higher on the average when the root was a low-degree node, compared to trees with a high-degree root.

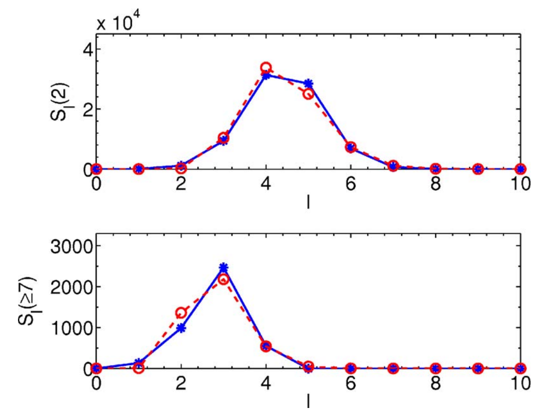

FIG. 9. (Color online) Top: Number of degree two nodes ($k=2$) at each layer of a multicast tree cut from BA topology with $m_0=6, m=1.5, p=0.1, q=0$, and 10^6 nodes. Empty symbols represent a fit to the gamma distribution with $\alpha=23.846$ and $\beta=5.397$. Bottom: Number of high-degree nodes ($k\geqslant 7$) at each layer of a multicast tree cut from the same underlying topology. Empty symbols represent the corresponding gamma distribution with $\alpha=16.862$ and $\beta=5.959$. It can be seen that the high-degree nodes tend to reside much closer to the root than the low-degree nodes.

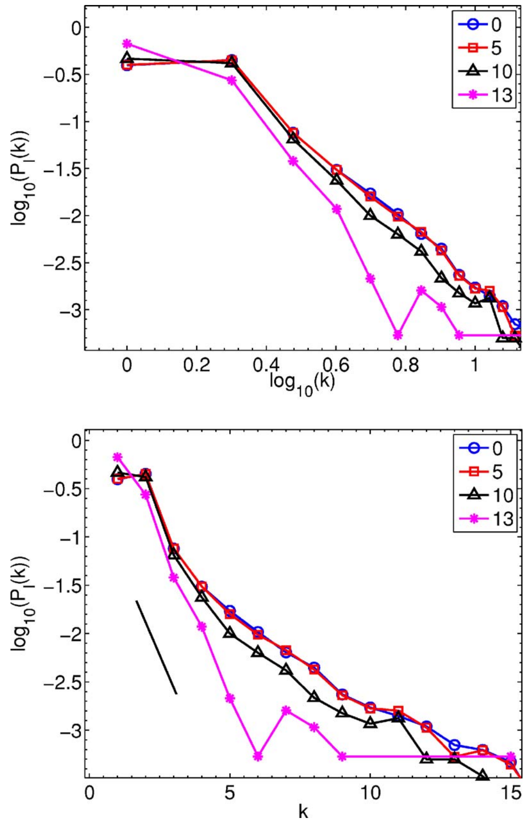

FIG. 10. (Color online) Tomography of Internet tree with 12 810 nodes and 5072 clients. Top: Log-log plot of the degree distributions $P_l(k)$ for different layers $l=0,5,10,13$. Bottom: A semilog plot of the same distributions. There is strong indication that the degree distributions exhibit a power-law tail ($\lambda_{\text{tree}} \approx 3.18$ [10]) with an exponential cutoff which is becoming stronger at each layer.

C. Tree characteristics—Internet data

We support the above findings with results obtained from a real Internet data set. Since we have no access to multicast tree data we use the client population of a medium sized web site with scientific and/or engineering content. This may represent the potential audience of a multicast of a program with scientific content. Two lists of clients were obtained, and traceroute was used to determine the paths from the root to the clients (see Ref. [10] for more details). It is important to note that the first three levels of the tree consist of routers that belong to the site itself, and therefore might be treated as the root point of the tree, although in our graphs they appear separately.

Figure 10 shows the degree distribution of consecutive layers of the Internet tree. Although the tree is rather small and the data sets are very noisy, there is a strong indication that the degree distribution behaves similar to what was observed in networks, i.e., a power law with an exponential cutoff that is becoming stronger at each layer [46].

V. SUMMARY AND CONCLUSIONS

In this paper we define a "layer" in a network as the set of nodes at a given distance from a chosen node. For scale-free networks, we find that the degree distribution of the nodes at each layer exhibits a power-law tail with an exponential cutoff. We derive equations for this exponential cutoff and compare them with empirical results. We also model the behavior of the number of nodes at each layer, and explain the observed exponential decay in the outer layers of the network. We obtain similar results for layers surrounding the root of shortest path trees cut from such networks, as well as for router-level data from the Internet.

It is possible that our findings may help in devising better network algorithms for the Internet that take advantage of the network structure. For example, we presented in the past [10] an algorithm for fast estimation of the multicast group size that is based on our previous findings regarding the distribution of high-degree nodes in Internet multicast trees.

Furthermore, our analytical results show that certain attributes observed in the Internet may be explained by simple models. Although the Internet is most probably a structured network, designed for optimal function, and not an arbitrary realization of a random graph [28,29], our analytical and empirical results enable us to reconstruct some of the central characteristics of the layer structure observed in the Internet.

ACKNOWLEDGMENTS

The authors wish to thank Rakesh Sinha for useful assistance. The authors thank the Israel Science Foundation for support. One of the authors (T.K.) thanks the Israeli Center for Complexity Science for financial support. This work was also supported by the European research NEST Contract No. DYSONET012911. After the introduction of this paper [47], several related papers discussing the structure of the Internet [28–31] and the effect of traceroute sampling [32–34] have appeared.

APPENDIX: DERIVING THE EXPONENTIAL CUTOFF USING AN ALTERNATIVE ANALYTIC APPROXIMATION

We treat each node independently, and the interaction between the nodes is inserted through the expected number of incoming connections (this is also called the "mean field" approach). At each node, the process is treated as equivalent to randomly distributing χ_l independent points on a line of length $\chi_l + T_l$ and counting the resultant number of points inside a *small* interval of length k. Thus, the number of incoming connections k_{in} from layer l to a node with k open connections is distributed according to a Poisson distribution with

$$\langle k_{\text{in}} \rangle = \frac{k}{\chi_l + T_l} \chi_l \tag{A1}$$

and

$$P_{l+1}(k_{\text{in}}|k) = e^{-\langle k_{\text{in}} \rangle} \frac{\langle k_{\text{in}} \rangle^{k_{\text{in}}}}{k_{\text{in}}!}. \tag{A2}$$

The probability for a node with k open connections *not* to be connected to layer l, i.e., to be outside layer $l+1$ also, is

$$P(k,l+1|l) = P_{l+1}(k_{\text{in}}=0|k) = e^{-\langle k_{\text{in}}\rangle} = \exp\left(-\frac{k}{1+\frac{T_l}{\chi_l}}\right). \quad \text{(A3)}$$

Thus the total probability to find a node of degree k outside layer $l+1$ is

$$P_{l+1}(k) = P_l(k)P(k,l+1|l) = P_l(k)\exp\left(-\frac{k}{1+\frac{T_l}{\chi_l}}\right), \quad \text{(A4)}$$

and one obtains an exponential cutoff.

[1] A.-L. Barabási, *Linked: The New Science of Networks* (Perseus, Cambridge, MA, 2002).

[2] S. N. Dorogovtsev and J. F. F. Mendes, *Evolution of Networks* (Oxford University Press, Oxford, 2003).

[3] R. Pastor-Satorras and A. Vespignani, *Evolution and Structure of the Internet: A Statistical Physics Approach* (Cambridge University Press, Cambridge, 2004).

[4] Scale-free networks have been a popular framework for studying the structure of complex networks in general, including the Internet. However, it is important to note that there remains a controversy if the general models for generating scale-free networks (e.g., the Barabási-Albert model and the configuration model used throughout this paper) are appropriate representations of the actual real world networks—see Refs. [28,29] for details.

[5] R. Cohen and S. Havlin, Phys. Rev. Lett. **90**, 058701 (2003).

[6] R. Cohen, K. Erez, D. ben-Avraham, and S. Havlin, Phys. Rev. Lett. **85**, 4626 (2000).

[7] H. Burch and B. Cheswick, IEEE Comput. **32**(4), 97 (1999).

[8] S. N. Dorogovtsev, J. F. F. Mendes, and A. N. Samukhin, Nucl. Phys. B **653**, 307 (2003).

[9] F. Chung and L. Lu, Internet Math. **1**, 91 (2003).

[10] D. Dolev, O. Mokryn, and Y. Shavitt, *IEEE INFOCOM'03* (IEEE, San Francisco, CA, 2003).

[11] P. Erdös and A. Rényi, Publ. Math., Inst. Hautes Etud. Sci. **5**, 17 (1960); B. Bollobás, *Random Graphs* (Cambridge University Press, Cambridge, 2001).

[12] A.-L. Barabási and R. Albert, Science **286**, 509 (1999).

[13] B. Bollobas, Eur. J. Comb. **1**, 311 (1980).

[14] M. Molloy and B. Reed, Combinatorics, Probab. Comput. **7**, 295 (1998).

[15] M. E. J. Newman, in *Handbook of Graphs and Networks*, edited by S. Bornholdt and H. G. Schuster (Wiley-VCH, Berlin, 2003), Chap. 2.

[16] D. S. Callaway, M. E. J. Newman, S. H. Strogatz, and D. J. Watts, Phys. Rev. Lett. **85**, 5468 (2000).

[17] Indeed, simulations show that explicitly removing self-loops and multiple edges from the network does not affect our results.

[18] R. Albert and A.-L. Barabási, Phys. Rev. Lett. **85**, 5234 (2000).

[19] R. Cohen, D. ben-Avraham, and S. Havlin, in *Handbook of Graphs and Networks*, edited by S. Bornholdt and H. G. Schuster (Wiley-VCH, Berlin, 2003), Chap. 4.

[20] S. N. Dorogovtsev, J. F. F. Mendes, and A. N. Samukhin, Phys. Rev. E **63**, 062101 (2001).

[21] A. A. Moreira, J. S. Andrade, and L. A. Nunes Amaral, Phys. Rev. Lett. **89**, 268703 (2002).

[22] This holds true assuming that almost no site in layer $l+1$ is reached by two connections from layer l. This is justified in the case where $m=1$, and also for the first layers in case of $m>1$.

[23] An approximate expression for the upper cutoff was found to be [5] $K_l \sim A\left\{\frac{[(\lambda-2)^{l-1}-1]}{(3-\lambda)}\right\}N\left[\frac{(\lambda-2)^l}{(\lambda-1)}\right]$, where $A=\langle k\rangle m^{\lambda-2}/(3-\lambda)=\frac{(\lambda-1)m}{(\lambda-2)(3-\lambda)}$.

[24] Although the configuration model (CM) gives a power-law degree distribution only in expectation, we find that for sufficiently large N the power-law form of the distribution does not vary significanly from one realization to the other, and there is no need to average over many realizations. Note also that in our simulations we allowed for two (or more) parallel links between two nodes. However, we found that whether we allow this or not, there is no significant effect on our results.

[25] W. Cheswick, J. Nonnenmacher, C. Sahinalp, R. Sinha, and K. Varadhan, Tech. Rep. Technical Memorandum 113410–991116–18TM, Lucent Technologies, 1999.

[26] M. E. J. Newman, Phys. Rev. Lett. **89**, 208701 (2002).

[27] R. Pastor-Satorras, A. Vazquez, and A. Vespignani, Lect. Notes Phys. **650**, 425 (2004).

[28] L. Li, D. Alderson, W. Willinger, and J. Doyle, *ACM SIGCOMM 2004* (ACM, Portland, OR, 2004).

[29] D. Alderson, J. Doyle, L. Li, and W. Willinger, Internet Math. **2**, 431 (2006).

[30] H. Chang, S. Jamin, and W. Willinger, *IEEE INFOCOM'06* (IEEE, Barcelona, Spain, 2006).

[31] W. Muhlbauer, A. Feldmann, O. Maennel, M. Roughan, and S. Uhlig, *ACM SIGCOMM'06* (ACM, Pisa, Italy, 2006).

[32] T. Petermnann and P. D. L. Rios, Eur. Phys. J. B **38**, 201 (2004).

[33] A. Clauset and C. Moore, Phys. Rev. Lett. **94**, 018701 (2005).

[34] L. Dall'Asta, I. Alvarez-Hamelin, A. Barrat, A. Vazquez, and A. Vespignani, Lect. Notes Comput. Sci. **3405**, 140 (2005).

[35] R. Cohen and D. Raz, *IEEE INFOCOM'06* (IEEE, Barcelona, Spain, 2006).

[36] Note that shortest path trees do not reflect the actual multicast protocols or router topology. However, some basic characteristics of multicast trees may be captured by these simplified models.

[37] In some of our simulations we use $m=1.5$. This is done by randomly setting m, at each iteration to be either $m=1$ or $m=2$ with equal probability.

[38] Precise data on the typical number of clients in real multicast trees is currently lacking. However, we found that most of the

tree characteristics described in this paper can be observed clearly when the number of clients is approximately 10% of the total number of network nodes and above.

[39] H. Holbrook and B. Cain, RFC 4607, IETF (unpublished).

[40] G. Philips, S. Shenker, and H. Tangmunarunkit, *ACM SIGCOMM'99* (ACM, Cambridge, MA, 1999).

[41] R. C. Chalmers and K. Almeroth, *IEEE INFOCOM'01* (IEEE, Anchorage, Alaska, 2001).

[42] P. Krishnan, D. Raz, and Y. Shavitt, IEEE/ACM Trans. Netw. **8**, 568 (2000).

[43] Note that the number of nodes at each layer in shortest path trees may be much smaller than the corresponding number in networks. Therefore very large network topologies must be taken in order to observe the exponential cutoff in shortest path trees. For small networks the exponential cutoff cannot be seen directly due to poor statistics.

[44] This phenomenon was observed in all cases, but was more obvious when the tree root was a high-degree node.

[45] Note that this may not be the case for any general network with a power-law degree distribution. Is was shown by Li *et al.* [29] that different graphs having the same power-law degree distribution have different properties if the high-degree nodes tend to reside in the graph center or in the graph periphery. These correlations may affect the layer structure—see also Ref. [48] for details.

[46] Fitting the degree distribution of each layer to a power law gives a roughly increasing power-law exponent for consecutive layers (results not shown). We believe this is actually an indication for the presumed exponential cutoff. However, current data sets of Internet trees do not allow for better statistical significance.

[47] R. Cohen, D. Dolev, S. Havlin, T. Kalisky, O. Mokryn, and Y. Shavitt, e-print cond-mat/0305582.

[48] R. Xulvi-Brunet, W. Pietsch, and I. M. Sokolov, Phys. Rev. E **68**, 036119 (2003).