On the Tomography of Networks and Multicast Trees

Reuven Cohen, Danny Dolev, Shlomo Havlin, Tomer Kalisky, Osnat Mokryn, Yuval Shavitt

***Abstract*— 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 is a power law with an exponential cut-off. We obtain similar empirical results for the layers surrounding the root of multicast trees cut from such networks, as well as the Internet.**

I. Introduction

In recent years there is an extensive effort to model the topology of the Internet. While the exact nature of the Internet topology is in debate [5], it was found that many realistic networks posses a power law, or scale free degree distribution [13], [14], [18], [6], [10]. Albert and Barabási [2], [1] suggested a dynamic graph generation model for such networks. One of their main findings was the self similarity characteristic of such networks. Interestingly, empirical findings on partial views obtained similar results, which may lead to the assumption that due to the self similarity nature of the Internet structure, this characteristic would be exposed through different cuts and filters.

In this paper we study the tomography of scale free networks and multicast trees cut from them. We use the Molloy Reed graph generation method [19] in conjunction with similar techniques to study the layer structure (tomography) of networks. Specifically, we study the number and degree distribution of nodes at a given (shortest path) distance from a chosen network node. We show analytically that the distance distribution of all nodes from a specific network node consists of two regimes. The first can be described as a very rapid growth, while the second is found to decay exponentially. We also show that the node degree distribution at each layer obeys a power law with an exponential cut-off. We back our analytical derivations with simulations, and show that they match.

As noted by Lakhina *et al.* [17], it is a significant challenge to test and validate hypotheses about the Internet topology, because of lack of highly accurate maps. Our analytical findings suggest a simple local test for the validity of the power law model as an exact model of the Internet. Indeed our findings suggest that there is a good agreement of the empirical and analytical results. The slight difference we had can be attributed to bias in data collection and to second order phenomena such as, degree correlation, hierarchies, and geographical considerations.

We also study shortest path trees cut from scale free networks, as they may represent the structure of multicast trees. We investigate their layer structure and distribution. We show that the structure of a multicast tree cut from a scale free network exhibits a layer behavior similar to the network it was cut from. We validate our analysis with simulations and real Internet data. We believe that enriching our understanding of the structure of multicast trees, can aid us in developing better multicast algorithms, e.g., in the past we used the statistics of high degree nodes to devise better algorithms for estimating the multicast group size [10].

The paper is organized as follows. Section II details previous findings. In section III we introduce the process used for generating scale free graphs and their layers. Then, we analyze the resulting tomography of such networks, and back the results with simulations and real data in section IV. In Section V we investigate the tomography of multicast trees cut from such networks, and back our findings with real Internet data.

II. Background

A. Graph Generation

Recent studies have shown that many real world networks, and, in particular, the Internet, are 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 [2], [19]. Molloy and Reed suggested an interesting construction method for scale free networks in [19]. The construction was part of a model describing an "exposure" process used to evaluate the size of the largest component in a random scale free network. We term this model the *MR model*. The construction method is as follows. A graph with a given degree distribution is generated out of the probability space (ensemble) of possible graph instances. For a given graph size N, the degree sequence is determined by randomly choosing a degree for each of the N nodes from the degree distribution. Let us define V as the set of N chosen nodes, C as the set of unconnected outgoing links from the nodes in V, and E as the set of edges in the graph. Initially, E is empty. Then, the links in C are randomly matched, such that at the end of the process, C is empty, and E contains all the matched links $< u, v >$, $u, v \in V$. Throughout this paper, we refer to the set of links in C as *open connections*.

Note, that while in the BA model the graph degree distribution function emerges only at the end of the process, in the MR model the distribution is known a-priori, thus enabling us to use it in our analysis during the construction of the graph.

B. Distribution Cut-Off

Recent work [9], [7], has shown that the radius[1], r, of scale free graphs with $2 < \lambda < 3$ is extremely small and scales as $r \sim \log\log N$. 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 get to 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 0. 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.

[1]We define the radius of a graph, r, as the average distance of all nodes in the graph from the node with the highest degree (if there is more than one we will arbitrarily choose one of them). The average hop distance or diameter of the graph, d, is restricted to:

$$r \leq d \leq 2r, \tag{1}$$

Thus the average hop sequence is bound from above and from below by the radius.

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 [8]:

$$\sum_{k=K}^{\infty} P(k) \approx \int_{K}^{\infty} P(k)dk = 1/N. \tag{2}$$

Solving the integral yields $K \approx mN^{1/(\lambda-1)}$, which is the approximate natural upper cutoff of a scale free network [8], [11], [20].

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 maximal 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 Molloy-Reed model [19], also described in section II. The construction process tries to gradually expose the network, following the method introduced in [9], [7], and is forcing a hierarchy on the Molloy-Reed model, 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 [8], [11].

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 II, the set of links in E is empty at this point, while the set of outgoing open links in C contains all unconnected 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 figure 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 figure 1(b)). Generally, to form layer $l+1$ from an arbitrary layer l, we randomly connect all open connections emerging from l to either other open connections emerging from l or chosen from the other links in C (see figure 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:

$$T_l = N \sum_k k P_l(k) \tag{3}$$

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 figure 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:

$$P(k, l+1|l) = \left[1 - \frac{k}{\chi_l + T_l}\right]^{\chi_l} \approx \approx exp\left(-\frac{k}{1+\frac{T_l}{\chi_l}}\right), \tag{4}$$

for large enough values of χ_l.

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

$$P_{l+1}(k) = P_l(k) P(k, l+1|l) = = P_l(k) exp\left(-\frac{k}{1+\frac{T_l}{\chi_l}}\right) \tag{5}$$

Thus we derive the exponential cutoff:

$$P_l(k) = P(k) exp\left(-\frac{k}{K_l}\right) \tag{6}$$

Where :

$$\frac{1}{K_{l+1}} = \frac{1}{K_l} + \frac{1}{1+\frac{T_l}{\chi_l}} \tag{7}$$

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

Now let us find the behavior of χ_l and S_l, where S_l is the number of links incoming to the $l+1$ layer (and approximately[2] equals N_{l+1}, the number of nodes in the $l+1$ layer). 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:

$$P(\text{loop}|l) = \frac{\chi_l}{\chi_l + T_l} \tag{8}$$

and Therefore:

$$S_{l+1} = \chi_l \left(1 - \frac{\chi_l}{\chi_l + T_l}\right) \tag{9}$$

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:

$$\chi_{l+1} = T_l - T_{l+1} - S_{l+1} \tag{10}$$

At this point we have the following relations: $T_{l+1}(K_{l+1})$ Eq. (3) and Eq. (6), $S_{l+1}(\chi_l, T_l)$ Eq. (9), $K_{l+1}(K_l, \chi_l, T_l)$ Eq. (7), and $\chi_{l+1}(T_l, T_{l+1}, S_{l+1})$ Eq. (9) and (10). These relations may be solved numerically. Note that approximate analytical results for the limit $N \to \infty$ can be found in [9], [7], [12]. [3]

IV. Empirical Results on Networks

Figure 2 shows results from simulations (colored 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 [9], [7]. 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

[2]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 $m > 1$.

[3]An approximate expression for the upper cutoff was found to be [9]:

$$K_l \sim A^{\frac{(\lambda-2)^{l-1}-1}{3-\lambda}} N^{\frac{(\lambda-2)^l}{\lambda-1}}. \tag{11}$$

where $A = \langle k \rangle m^{\lambda-2}/(3-\lambda) = \frac{(\lambda-1)m}{(\lambda-2)(3-\lambda)}$.

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. Note the exponential cutoff which becomes stronger with l.

It is important to note that the simulation results give the probability distribution for the giant percolation cluster, while the 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 percolation cluster and therefore the probability distribution derived from the simulation is smaller for low degrees.

Figure 4 and Figure 5 show the same analysis for a cut of the Internet at router level (Lucent mapping project [3], 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.

In general, large degree nodes of the network mostly 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 $\lambda = 2.3$).

Our model does not take into account the correlations in node degrees, which were observed in the Internet [21], and hierarchical structures [24]. This may also explain the deviation of our measurements from the model predictions.

V. Empirical Findings on the Tomography of Multicast Trees

In this section, we detail some of our findings on the structure and characteristics of the depth rings around the root node of shortest path trees. All of our findings were also validated on real Internet data.

A. Topology and Tree Generation

Our method for producing trees is the following. First, we generate power law topologies based on the Barabási-Albert model [1]. The model specifies 4 parameters: a_0, a, p and q [4]. Where a_0 is the initial number of detached nodes, and a is the initial connectivity of a node. When a link is added, one of its end points is chosen randomly, and the other with probability that is proportional to the nodes degree. This reflects the fact that new links often attach to popular (high degree) nodes. The growth model is the following: with probability p, a new links are added to the topology. With probability q, a links are rewired, and with probability $1 - p - q$ a new node with a links is added. Note that a, p and q determine the average degree of the nodes. 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). 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 defined four basic tree types. These types are based on the rank of the root node and the clients nodes. The rank of a node is its location in a list of descending degree order, in which the lowest rank, one, corresponds to the node with the highest degree in the graph. For the case of a tree rooted at a big ISP site, we choose a root node with a low rank, thus ensuring the root is a high degree node with respect to the underlying topology. Then, we either choose the clients as high ranked nodes, or at random, as a control group. Note, that due to the characteristic of the power law distribution, a random selection of a rank has a high probability of choosing a low degree node. The next two tree types have a high ranked root, which corresponds to a multicast session from an edge router. Again, the two types differ by the clients degree distribution, which is either low, or picked at random.

The tree client population is chosen at the range $[50, 4000]$ for the 10000 node generated topology, $[50, 10000]$ for the 100000 node generated topology, and $[500, 50000]$ for the trees cut from real Internet data. For each client population size, 14 realizations were generated for each of the four tree types. All of our results are averaged over these realizations. The variance of the results was always negligible.

[4]The notations in [1] are m_0, m, p and q.

Name	Type	Parameters	No. of Nodes	Avg. Node degree
VS	generated	$a = 1; p \in 0 : 0.05 : 0.5$	10000	$1.99 - 3.98$
IS	generated	$a = 2; p \in 0 : 0.05 : 0.5$	10000	$3.99 - 7.9$
LS	generated	$a = 3; p \in 0 : 0.05 : 0.5$	10000	$5.98 - 12.04$
Big IS	generated	$a = 1.5, 2; p = 0.1$	50000;100000	3.3,4.4
BL[1,2]	real data	–	Internet	3.2 [5]
LC	real data	–	Internet	3.2 [6]

TABLE I

TYPE OF UNDERLYING TOPOLOGIES USED

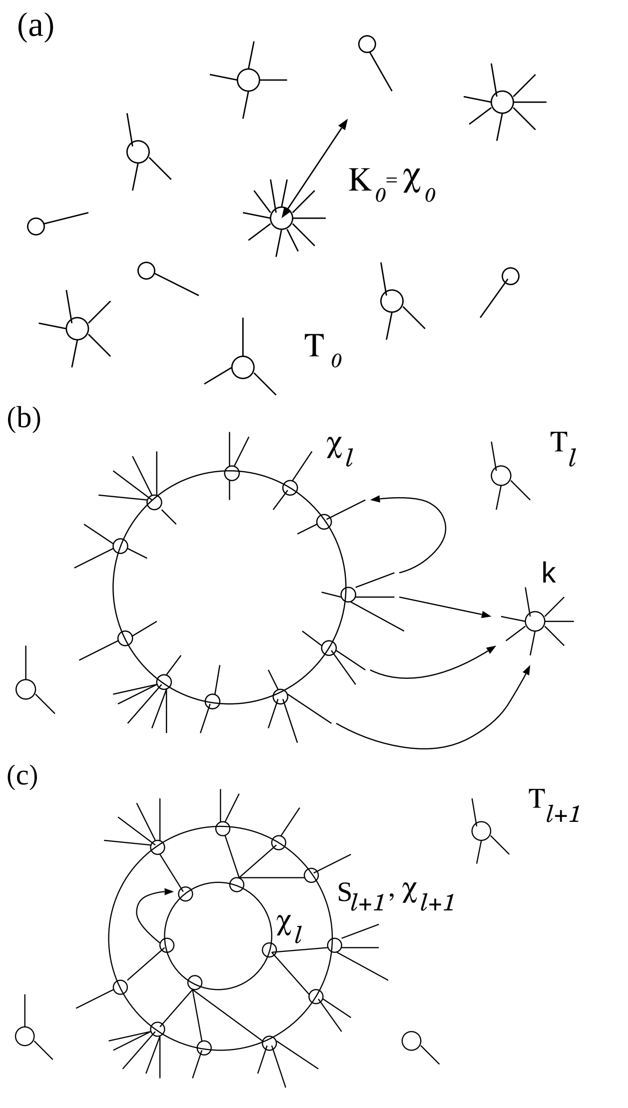

Fig. 1. Illustration of the exposure process. The large circles denote exposed layers of the giant component, while the small circles denote individual sites. The sites outside the circles have not been reached yet. (a) We begin with the highest degree node and fill out layer No.1. (b) In the exposure 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.

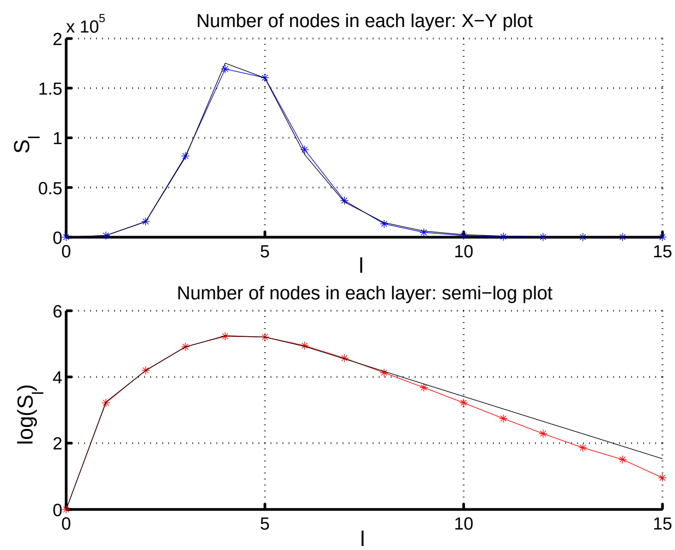

Fig. 2. Approximate number of nodes (S_l) vs. layer index l for a network with $N = 10^6$ nodes, $\lambda = 2.85$, and $m = 1$. Symbols represent simulation results while black lines are a numerical solution for the derived recursive relations. Bottom: from the semi-log plot we see that there is an exponential 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.

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. For example, IP packets are forwarded based on the reverse shortest path, and multicast routing protocols such as Source Specific Multicast [15] deliver packets along the shortest path route. In addition, we assume that client distribution in the tree is uniform, as has been shown by [23], [4].

B. Tree Characteristics

Our results show that trees cut from a power law topology obey a similar power law for the degree distribution, as well as the sub-trees sizes [10]. The results

[5] based on [16]

[6] based on [3]

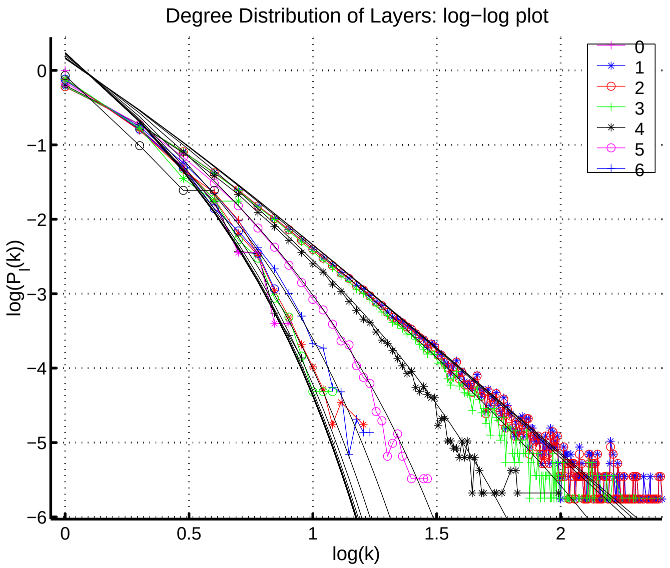

Fig. 3. Log-log plot of $P_l(k)$ for different layers $l = 0, 1, 2, \ldots$, for a network with $N = 10^6$ nodes, $\lambda = 2.85$, and $m = 1$. Symbols represent simulation results while black lines are a numerical solution for the derived recursive relations.

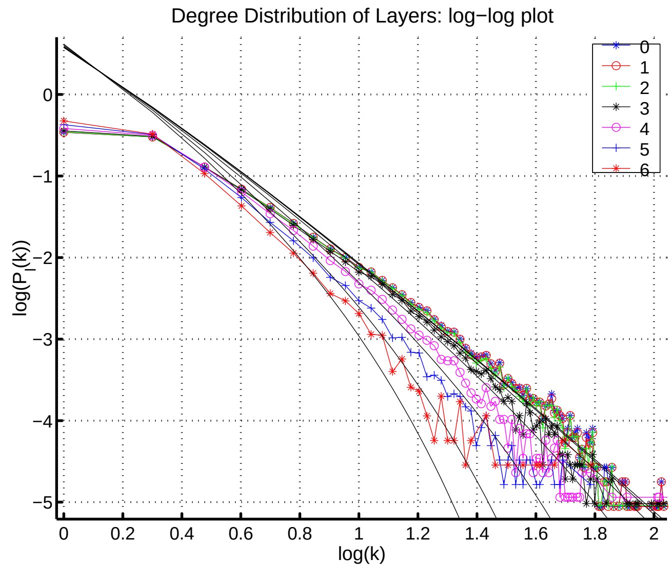

Fig. 5. Log-log plot of $P_l(k)$ for different layers $l = 0, 1, 2, \ldots$, for a router level cut of the Internet with $N = 112,969$ nodes (LC topology). Qualitative analytical reconstruction is done with $\lambda = 3$, and $m = 1$.

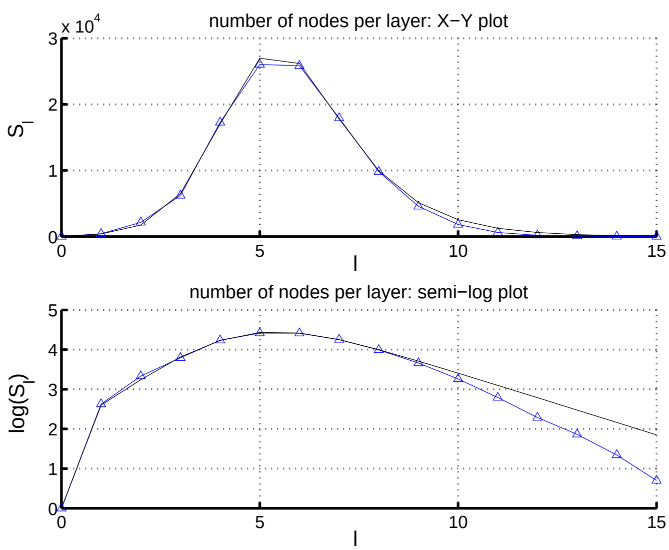

Fig. 4. Number of nodes at each layer for a router level cut of the Internet with $N = 112,969$ nodes (LC topology). Analytical reconstruction for S_l is done with $\lambda = 3$, and $m = 1$.

were shown to hold for all trees cut from all generated topologies, even for trees as small as 200 nodes.

In this work we further investigated the tomography of the trees, and looked at the degree distribution of nodes at different depth rings around the root, i.e., tree layers. It was rather interesting to observe that any layer with sufficient number of nodes to create a valid statistical sample obeyed a degree-frequency relationship which was similar to a power law, although with different slopes. We suspect that this is due to the exponential cut-off phenomenon discussed in the previous sections. Figure 6 shows this for the third layer around the root (i.e., nodes at distance three from the root) of a 300 client tree cut from a big IS topology (100000 nodes). The root was chosen with a high degree, and the clients with a low degree. Although the number of nodes is quite small, we see a very good fit with the power law. Figure 7 shows an excellent fit to the power law for the fifth layer around the root of a 10000 client tree, cut from the same topology. This phenomenon is stable regardless of the tree type, and the client population size. Note that the range of the power laws seen in figures 6 and 7 is less than one order of magnitude. This could indicate a crossover to exponential behavior.

To understand the exact relationship of the degree-frequency at different layers, we plotted the distribution of each degree at different layers. Cheswick at al. [6] found a gamma law for the number of nodes at a certain distance from a point in the Internet. Our results show that the distribution of nodes of a certain degree at a certain distance (layer) from the root seems close to a gamma distribution, although we did not determine its exact nature. Figure 8 shows the distribution of the distance of two degree nodes, and Figure 9 the distribution of the distance of high degree nodes, i.e., nodes with a degree six and higher. In both figures the root is a low degree node, and the tree has 1000 low degree clients. As can be seen, the high degree nodes tend to reside much closer to the root than the low degree nodes, and in adjacent layers. In this example, most of

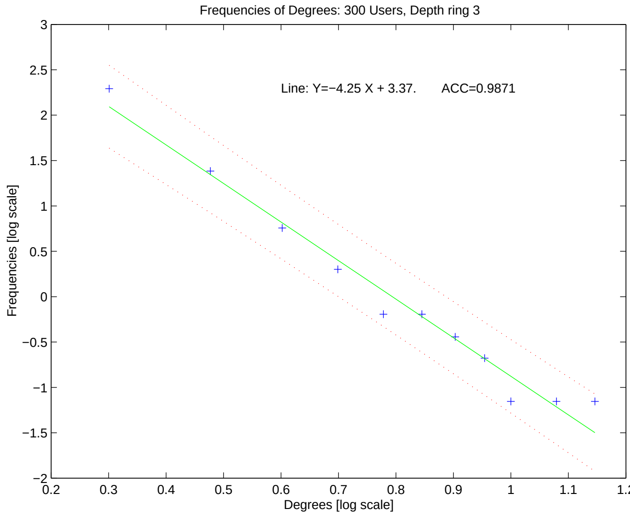

Fig. 6. Third layer of a 300 client tree cut from topology $a_0 = 6, a = 1.5, p = 0.1, q = 0$

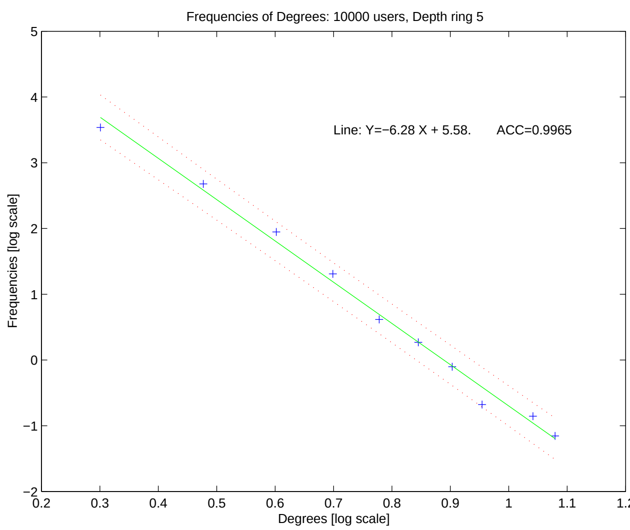

Fig. 7. Fifth layer of a 10000 client tree cut from topology $a_0 = 6, a = 1.5, p = 0.1, q = 0$

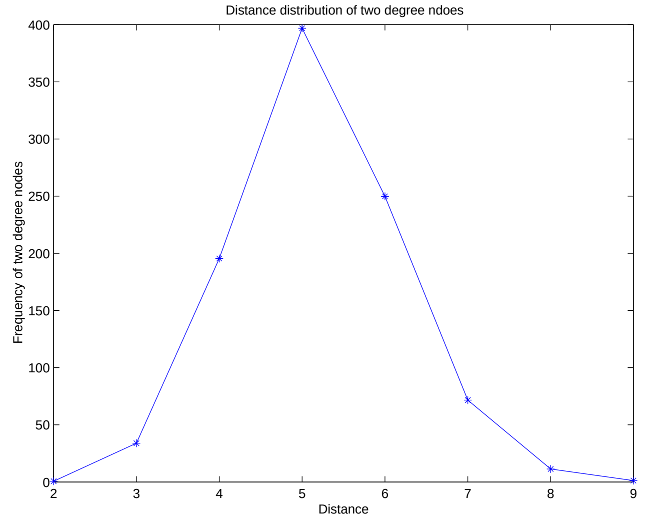

Fig. 8. Distribution of degree two nodes in a tree cut from topology $a_0 = 6, a = 1, p = 0.3, q = 0$.

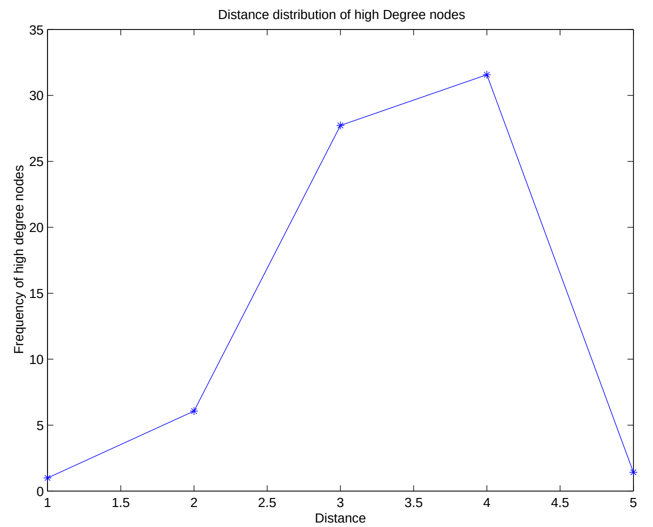

Fig. 9. Distribution of the high degree nodes ($k \geq 6$) in a tree cut from topology $a_0 = 6, a = 1, p = 0.3, q = 0$.

them are in the second to forth layers around the root, with only two more at layer five. This phenomenon was even more obvious when the root was a high degree node.

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 10000 node underlying topology with an average degree of three and higher, the height of the trees was not more than ten. On an underlying topology of 100000 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.

We verify 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/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. 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 these graphs they appear separately. Figure 10 shows the frequency of degrees in the tree.

The linear fit of the log-log ratio is excellent, with a correlation coefficient of 0.9829. The exponent is very close to the exponent we derived for trees cut from topologies that resemble the Internet.

Figures 11 and 12 show the frequency of degrees at layers 5 and 10 of the tree, respectively. It can be seen that the slope λ of the distribution increases with the layer number, e.g., layer 5 has a slope $\lambda \approx 2.34$, and layer 10 has a slope $\lambda \approx 2.99$. As we claimed, the shortest path tree cut from a scale free topology inherit many of the characteristics of the network topology. Moreover, we found for networks that the frequency-degree for each separate distance around the root can be approximated by a power law with an exponential cut-off, which is becoming stronger with the layer number. In Fig. 13 we plotted the slope of the distribution in the layer against the layer number and found a very good linear fit (note the outlier at $l = 7$ which was not included in the fit). The linear fit indicates that for the first layer the slope will be -1.83 ± 0.125. For scale free networks, it has been shown [22] that the first layer surrounding a chosen network node has a distribution $kP(k) \sim k^{-\lambda+1}$. Therefore, we can expect that in the first tree layer surrounding the tree root will have a frequency-degree slope of approximately $-3.18 + 1 = -2.18$ ($\lambda = -3.18$, the slope of the tree, is taken from Fig. 10) which is close to the linear prediction. While the results for the degree distribution in the first layer did not have statistical significance the slope for the second layer was -2.09 which conforms to the above numbers.

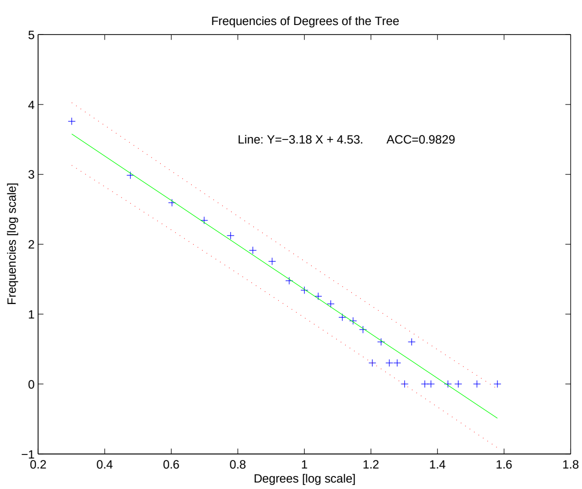

Fig. 10. Frequency of degrees of the Internet tree.

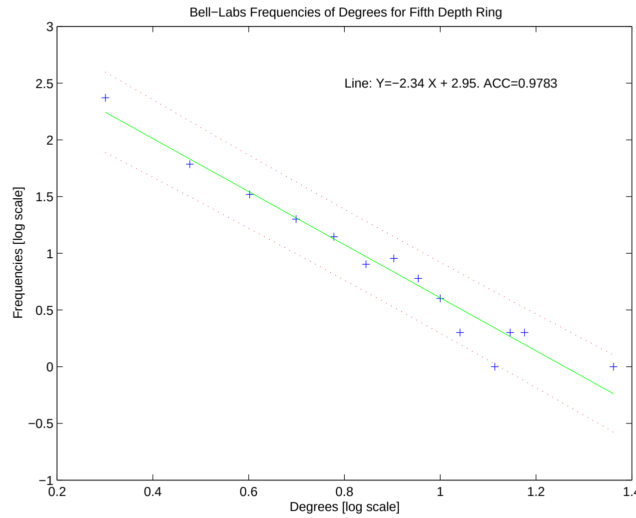

Fig. 11. Frequency of degrees at layer 5 of the Internet tree.

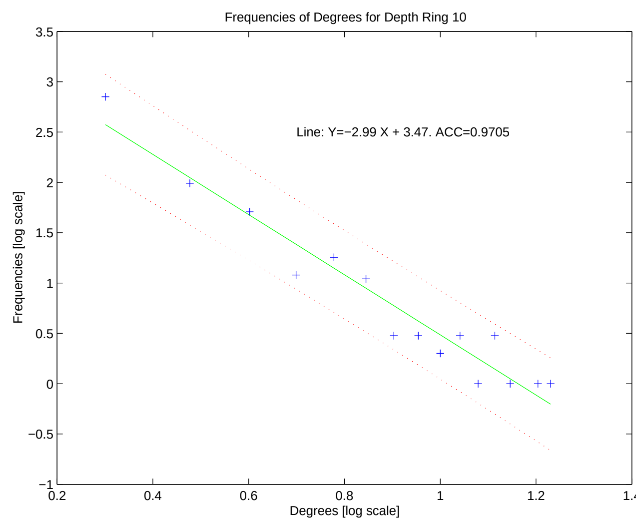

Fig. 12. Frequency of degrees at layer 10 of the Internet tree.

VI. Conclusions

We define a "layer" in a network as the set of nodes at a given distance from a chosen node. We find that the degree distribution of the nodes of a scale free network at each layer obeys a power law 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 multicast trees cut from such networks, as well as the Internet.

We believe our findings can have dual importance. First, they can help in devising better network algorithms 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

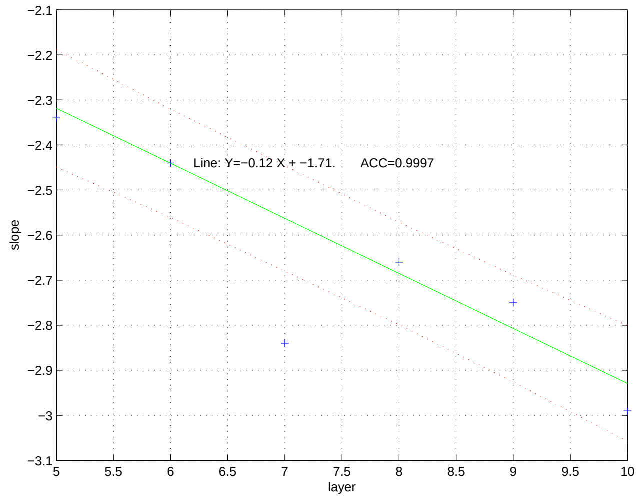

Fig. 13. Slope of the degree distribution at specific layer as a function of the layer number.

our previous finding regarding the distribution of high degree nodes in Internet multicast trees. Second, our analytical findings suggest a simple local test for the validity of the power law model as an exact model of the Internet. Indeed our findings suggest that there is a good agreement of the empirical and analytical results.

REFERENCES

[1] R. Albert and A.-L. Barab´asi. Topology of evolving networks: local events and universality. *Physical Review Letters*, 85(24):5234–5237, 11 Dec. 2000.

[2] A.-L. Barab´asi and R. Albert. Emergence of scaling in random networks. *SCIENCE*, 286:509 – 512, 15 Oct. 1999.

[3] H. Burch and B. Cheswick. Mapping the internet. *IEEE Computer*, 32(4):97–98, 1999.

[4] R. C. Chalmers and K. Almeroth. Modeling the branching characteristics and efficiency gains in global multicast trees. In *IEEE INFOCOM'01*, Anchorage, Alaska, 2001.

[5] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. The origin of power-laws in internet topologies revisited. In *IEEE Infocom 2002*, New-York, NY, USA, Apr. 2002.

[6] W. Cheswick, J. Nonnenmacher, C. Sahinalp, R. Sinha, and K. Varadhan. Modeling internet topology. Technical Report Technical Memorandum 113410-991116-18TM, Lucent Technologies, 1999.

[7] R. Cohen, D. ben Avraham, and S. Havlin. Structural properties of scale free networks. In S. Bornholdt and H. G. Schuster, editors, *Handbook of graphs and networks*, chapter 4. Wiley-VCH, 2002.

[8] R. Cohen, K. Erez, D. ben Avraham, and S. Havlin. Resilience of the internet to random breakdowns. *Physical Review Letters*, 85:4626–4628, 2000.

[9] R. Cohen and S. Havlin. Ultra small world in scale free graphs. *Physical Review Letters*, 90:058701, 2003.

[10] D. Dolev, O. Mokryn, and Y. Shavitt. On multicast trees: Structure and size estimation. In *IEEE INFOCOM'03*, San-Francisco, CA, USA, 2003.

[11] S. Dorogovtsev, J. Mendes, and A. Samukhin. Size-dependent degree distribution of a scale-free growing network. *Phys. Rev. E*, 63:062101, 2001.

[12] S. Dorogovtsev, J. Mendes, and A. Samukhin. Metric structure of random networks. *NUCLEAR PHYSICS B*, 653 (3):307–338, 2003.

[13] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *ACM SIGCOMM 1999*, Boston, MA, USA, Aug./Sept. 1999.

[14] R. Govindan and H. Tangmunarunki. Heuristics for internet map discovery. In *IEEE Infocom 2000*, pages 1371–1380, Tel-Aviv, Israel, Mar. 2000.

[15] H. Holbrook and B. Cain. Source-specific multicast for IP. In *Internet Draft, IETF*, Feb. 2002.

[16] P. Krishnan, D. Raz, and Y. Shavitt. The cache location problem. *IEEE/ACM Transactions on Networking*, 8(5):568–582, Oct. 2000.

[17] A. Lakhina, J. Byers, M. Crovella, and P. Xie. Sampling biases in ip topology measurements. In *IEEE INFOCOM'03*, San Francisco, CA, USA, Mar. 2003.

[18] A. Medina, I. Matta, and J. Byers. On the origin of power laws in internet topologies. *ACM Computer Communications Review*, 30(2):18–28, 1 Jan. 1998.

[19] M. Molloy and B. Reed. The size of the giant component of a random graph with a given degree sequence. *Combinatorics, Probability and Computing*, 7:295–305, 1998.

[20] A. A. Moreira, J. Jos S. Andrade, and L. A. N. Amaral. Extremum statistics in scale-free network models. *Physical Review Letters*, 89:268703, 2002.

[21] M. E. J. Newman. Assortative mixing in networks. *Phys. Rev. Lett*, 89:208701, 2002.

[22] M. E. J. Newman, S. H. Strogatz, and D. J. Watts. Random graphs with arbitrary degree distributions and their applications. *Phys. Rev. E*, 64:026118, 2001.

[23] G. Philips, S. Shenker, and H. Tangmunarunkit. Scaling of multicast trees: Comments on the Chuang-Sirbu scaling law. In *ACM SIGCOMM'99*, Cambridge, MA, USA, 1999.

[24] A. Vazquez, R. Pastor-Satorras, and A. Vespignani. Internet topology at the router and autonomous system level. *cond-mat/0206084*, 2002.

APPENDIX

A. Deriving the Exponential Cutoff Using Alternative Analytic Approximation

Each node is treated independently, where the *interaction* between nodes is inserted through the expected number of incoming connections. 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:

$$< k_{in} >= \frac{k}{\chi_l + T_l}\chi_l \qquad (12)$$

and :

$$P_{l+1}(k_{in}|k) = e^{-<k_{in}>}\frac{< k_{in} >^{k_{in}}}{k_{in}!} \qquad (13)$$

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:

$$P(k, l+1|l) = P_{l+1}(k_{in} = 0|k) = e^{-<k_{in}>} = \\ = exp\left(-\frac{k}{1+\frac{T_l}{\chi_l}}\right) \qquad (14)$$

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

$$P_{l+1}(k) = P_l(k)P(k, l+1|l) = P_l(k)exp\left(-\frac{k}{1+\frac{T_l}{\chi_l}}\right) \qquad (15)$$

And one obtains an exponential cutoff.