IP Network Topology and the Impact of Underlying Transport Networks.

Jason Spencer*, Lionel Sacks

{ jason, lsacks } @ee.ucl.ac.uk

Department of Electronic and Electrical Engineering, University College London, London, UK

Abstract: Recent developments in Internet mapping and metrification have suggested that the topology of the Internet follows a series of power laws. In this paper we propose a simple algorithm for generating such topologies which uses growth and existing connectivity as necessary elements to generate power law networks. We also discuss the impact that the underlying transport networks and their topologies may play on the construction of IP networks.

1 Introduction

When studying aspects of networks and the effect of various ideas such as routing strategies, network dimensioning and the like, one very large variable is the topology of the network. For simplicity regular topologies are often used such as uniform grids of nodes or rings or more random graphs like a uniformly randomly connected network. Real IP network topologies aren't usually like this and this can have a large impact on results [15][16]. Recent studies [1] have shown that the Internet topology isn't one of these regular patterns or even a totally randomly connected set of nodes but rather a network which follows an emergent topology which exhibits a number of power-laws. The source of these power laws has been speculated [1] and has been attributed often to the fact that these networks grow in size. Current topology generators however consider the IP network as an independent layer and don't consider the underlying transport networks (some have considered distance cost functions [6]). The growth of networks together with the cost functions applied by the transport networks may result in more realistic topologies and offer an explanation as to the source of the power laws.

The need for more realistic topologies was demonstrated by Albert et al. who showed that scale-free topologies are resistant to random failure but very sensitive to deliberate attack [5][8] but uniformly randomly connected networks are sensitive to both attack and failure. To evaluate the realism of topology generators some mettrics will also be required.

2 Internet Topology Metrics

In any topology consisting of nodes and uni-directional links there are two immediate metrics: the number of nodes and the number of links, and hence the average connectivity of the nodes. Recent laws discovered in Internet topologies can also give us a number of metrics.

2.1.1 Power Laws in Internet Topologies

Faloutsos et al. discovered four power laws [1] in three instances of inter-domain topologies and one instance of a node-level topology. The following four laws were found to hold at both the node-level and the BGP domain-level:

Power-Law 1 (rank exponent): The outdegree (connections to a node) was found to be proportional to the rank of a node, to the power of a constant. The rank (position in table) being the position of the node in a table sorted (numerically decreasing) by the outdegree of the node.

Power-Law 2 (outdegree exponent): The frequency of an outdegree is proportional to the outdegree to the power of a constant.

Power-Law 3 (hop-plot exponent): The total number of pairs of nodes within h hops of each other, is proportional to the number of hops to the power of a constant. This is more of an approximation since it only holds for value of h which are much less than the network diameter.

Power-Law 4 (eigenvalue exponent): The sorted eigenvalues (decreasing order) of the adjacency matrix (an N node by N node matrix which is 1 when the two nodes are connected and 0 otherwise) are proportional to the index into the list, to the power of a constant. The power law was shown to hold for only the top 20 or so eigenvalues [1].

The exponents of these power laws are some reflection of the topologies and therefore we will be considering them as metrics of the topologies. The eigenvalues for example are a metric for a number of characteristics including network diameter, connectivity and cliques (minimal sub-graphs) among others. Faloutsos et al.

* Corresponding author, funded by the Engineering and Physical Sciences Research Council (EPSRC), UK and British Telecommunications Plc

showed that there were big differences in the values for these between node-level and domain-level topologies (Table 1-Table 4).

2.2 Topology Generators

Many topology generators have already been proposed in literature [6][10][11][12][13][14], some consider the known structure of the network being modelled such as the two tier architecture of the Internet, clustering and sub-nets while some consider random graphs and create more generic topologies.

2.2.1 Random Graph Models

Topologies of uniformly randomly connected nodes, first examined by Erdös and Rényi [13] are commonly used to generate test networks. They have a few shortcomings, such as the lack of internal structure, and this leads to characteristics like an average network diameter that is independent of the number of nodes and that all nodes have the same average outdegree. Other random models like the small-world model described by Watts and Strogatz [14] consists of a lattice of nodes, each one connected to n_{ws} nearest neighbours. Then with some probability p_{ws}, one of the endpoints of a random link is reconnected to a random node. It has some interesting properties such as a distinct phase transition in the average diameter of the network once p_{ws} passes a certain threshold.

None of these graph models however cater for power laws in the topologies. There are a number of ways of generating such graphs but they generally rely on preferential attachment of nodes according to existing connectivity. Nodes connect to other nodes, preferring to connect to the already more connected nodes [5]: such that the probability of connecting to a node i, of j nodes in the network, with k_i links already is $P(k_i) \sim k_i / \Sigma k_j$. Such a network model creates networks which follow all four powers laws described by Faloutsos et al. albeit with different exponents.

2.3 Current Topology Generators

A number of topology generators exist, each concentrating on various aspects of the Internet. Waxman [6] first proposed a topology generator in his examination of multicast routing trees. The generator used the euclidean distance between nodes to govern their connectivity: $P(u,v) = \beta e^{-d(u,v)/L\alpha}$, where P(u,v) is the probability of linking nodes u and v, $d(u,v)$ is the euclidean distance between u and v, L is the euclidean diameter of the network and a and ß are parameters. The use of euclidean distance now makes the geographic distribution of nodes a factor in the topology. The geographic distribution is in reality closely governed by the distribution of equipment in the transport networks.

Other topology generators include Transit-Stub [11] and Tiers [10] which attempt to emulate different aspects of Internet structure such as transit network or hierarchical topologies. More recently the BRITE [12] topology generator was created to investigate the source of power laws in Internet topologies.

2.3.1 Proposed topology generation algorithm

The algorithm to create topologies should be simple, fast, scaleable and configurable. It should also allow incremental growth and allow dynamic changes in connectivity.

The algorithm exhibits growth in network size, preferential connectivity and growth in connectivity. It was designed to emulate real network growth. A node is added, initially with a single link to a random node, favouring more connected nodes, and then as the load on the network increases links are added between randomly chosen node pairs, again favouring the more connected nodes.

The algorithm is as follows:

1. $N_l = G_l$;
2. Add a node. Connect it to only one other node, where the probability of connecting to a node is $P(k) \sim k_i / \Sigma k_j$.
3. while(N_l>1) { Choose two different unconnected nodes in the network with probability $P(k) \sim k_i / \Sigma k_j$ and connect them. $N_l = N_l - 1$; }
4. $N_l = N_l + G_l$.
5. Jump to step 2 until the network is the required size (number of nodes).

Here G_l is the growth in number of links per time epoch (per node added) (G_l is any real number > 0) and N_l is the number of links which are to be added to the network in that time epoch. So therefore, at every time epoch a node is added and preferentially connected to an existing node with a single link. Then with a ratio of G_l (link additions per node addition) add a link between any two nodes that aren't already connected, favouring the more connected nodes. More algorithms which rely on preferential connectivity can be seen in [7].

The result is a network that has a link to node ratio of G_l +1. The parameters are therefore average outdegree (link to node ratio) and the network size (number of nodes).

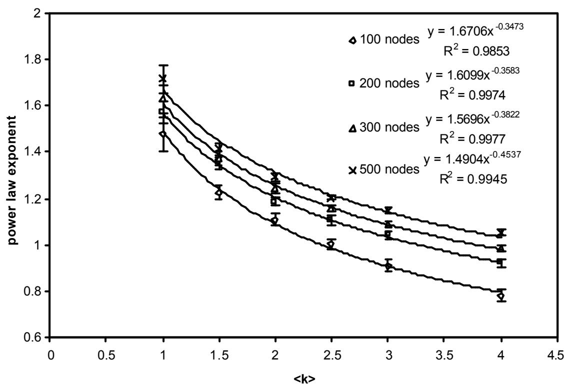

Figure 1 The outdegree frequency power law (power-law 2 [1]) exponent versus the average outdegree

Topology Instance	Num. of nodes	<k>	Measured rank exponent	Predicted (500 nodes)	Predicted (given # of nodes)
Int-11-97	3015	3.42	0.81139	1.18	1.12
Int-04-98	3530	3.65	0.82127	1.21	1.15
Int-12-98	4389	3.76	0.74496	1.22	1.15
Rout-95	3888	2.57	0.48759	1.05	0.99

Table 1 Power Law Exponent predictions and actual values for the outdegree rank exponent

Topology Instance	Num. of nodes	<k>	Measured outdegree exponent	Predicted (500 nodes)	Predicted (given # of nodes)
Int-11-97	3015	3.42	2.15632	0.85	0.16
Int-04-98	3530	3.65	2.16356	0.83	0.08
Int-12-98	4389	3.76	2.20288	0.82	-0.01
Rout-95	3888	2.57	2.48626	0.97	0.05

Table 2 Power Law Exponent predictions and actual values for the outdegree frequency exponent

Topology Instance	Num. of nodes	<k>	Measured hop-plot exponent	Predicted (500 nodes)	Predicted (given # of nodes)
Int-11-97	3015	3.42	-4.62706	-2.78	-3.95
Int-04-98	3530	3.65	-4.71768	-2.74	-4.02
Int-12-98	4389	3.76	-4.86588	-2.72	-4.15
Rout-95	3888	2.57	-2.83987	-2.95	-4.28

Table 3 Power Law Exponent predictions and actual values for the hop-plot exponent

Topology Instance	Num. of node	<k>	Measured eigenvalue exponent	Predicted (500 nodes)	Predicted (given # of nodes)
Int-11-97	3015	3.42	0.471327	0.55	0.45
Int-04-98	3530	3.65	0.502062	0.56	0.46
Int-12-98	4389	3.76	0.486946	0.57	0.45
Rout-95	3888	2.57	0.17742	0.50	0.39

Table 4 Power Law Exponent predictions and actual values for the eigenvalue exponent

3 Examination of Generated Topologies

The topology generator was run for a range of network sizes (N nodes) of 100, 200, 300 and 500 nodes with average outdegree values $<k>$ of 1.0, 1.2, 1.5, 2.0, 2.5, 3.0 (G_l = 0.0, 0.2, 0.5, 1.0, 1.5, 2.0, 3.0 respectively). The power law correlation of the resulting topologies for all four laws was good (R^2 > 0.9 for all power laws). The exponents and coefficients of the power law fits are however dependent on network size and average outdegree. To investigate this further we plotted the exponents plotted over different $<k>$ and N values, the oudegree exponent can be seen in Figure 1.

3.1 A Quantitative Examination of the Power Laws

Let us now consider the suitability of the algorithm to Internet topology generation and the choice of parameters. The network sizes were unfortunately limited to 500 nodes due to the need for shortest path routing in the hop-plot exponent calculation. Over a range of $<k>$ values the exponents fit trends of their own (Figure 1). The range of average outdegree values $<k>$ was chosen so that it covers the values measured by Faloutsos et al. We can see that as the network size gets larger the lines of the graph converge. In Table 1 to Table 4 we have the network parameters as per Faloutsos' data and two predictions as to the exponents of the resulting networks. There are three data sets describing AS topologies (Int-11-97, Int04-98, Int-12-98) and one describing node level topology (Rout-95). The first prediction we make ignores the real network size and interpolates along the 500 node trend. The second case attempts to extrapolate the coefficients of the trends to the given network size and makes an estimate to what the exponent would be if the Internet were in fact like our topology generator.

We can see that even though the topologies do follow power laws they don't match the exponent value from the Internet study. In a few cases the predictions were at least within 10-20% of the measured values, for example the AS data eigenvalue exponent or the AS data hop-plot exponents. The other power laws, especially the node-level exponents were however very different.

4 Conclusions and Discussion

The lack of similar exponents leads us to believe that while there certainly is a level of preferential connectivity when nodes are added at the node-level of the Internet this certainly isn't the only factor. We know for a fact that nodes are actually clustered within AS domains and that the connectivity of these also follows the power laws. This suggests a requirement for more structured elements in topology as well as the need to incorporate some form of geographical information and cost function, similar to Waxman's work [6] and similar to elements of the BRITE topology generator [12], especially the use of heavy-tailed node placement. Geographic restrictions would prevent the creation of links which traverse the entire network, which is unrealistic in real networks but possible in the generator.

The source of these restrictions and rules to the topological design is from two sources: The ability of the transport network, and the demands of the IP flows (and the effect of network operator policies on them). The IP flows are in turn influenced by web page connectivity, and the transport network is affected by the physical network. Certain levels of service may be required for certain links, the deliverable level of service is dependent on the transport network. The result is a feedback system because the users (or network designers) who create (or plan) demands will respond to perceived quality of service, and just the same way the physical layer will deliver levels of service according to what load it is already experiencing. This is similar to the Faloutsos et al. [1] postulate that the power laws are created by co-operative and antagonistic forces, and that the network must reconfigure itself to cope with demand. We could further hypothesise that since our feedback elements are in fact heterogeneous, that the source of our power laws is Self-Organising Criticality [17], where heterogeneity is actually a requirement [18].

REFERENCES

[1] M. Faloutsos, P. Faloutsos, C. Faloutsos, "On Power-Law Relationships of the Internet Topology", In proceedings of SIGCOMM: IEEE Conference on Communication, 1999.

[2] "Growth Dynamics of the World-Wide Web", Bernardo A. Huberman, Lada A. Adamic. Nature Vol 401, 9 Sept 1999

[3] "Scale-free characteristics of random networks: the topology of the world wide web" Albert-Laszlo Barabasi, Reka Albert and Hawoong Jeong. Physica A. 281 (2000) 69-77

[4] Abello, J., Buchsbaum, A. & Westbrook, J. "A functional approach to external graph algorithms". Lect. Notes Computer Science. 1461, 332-343 (1998).

[5] "Error and attack tolerance of complex networks", Reka Albert; Hawoong Jeong; Albert-Laszlo Barabasi; Nature Volume 406 Number 6794 Page 378 - 382 (2000)

[6] B.M.Waxman. "Routing of Mulitpoint Connections". IEEE Journal on Slected Areas in communications, 6(9):1617-1622, 1988

[7] Christopher R. Palmer and J. Gregory Steffan, "Generating Network Topologies That Obey Power Laws". Proceedings of the Global Internet Symposium, Globecom2000, November 2000, San Francisco.

[8] "Resilience of the Internet to Random Breakdowns", Reuven Cohen, Keren Erez, Daniel ben-Avraham, and Shlomo Havlin. Physical Review Letters -- November 20, 2000 -- Volume 85, Issue 21, pp. 4626-4628

[10] M. B. Doar. "A Better Model for Generating Test Networks." In Proceedings of Globecom '96, Nov. 1996

[11] E. W. Zegura, K. L. Calvert, and S. Bhattacharjee. "How to model an internetwork." In Proceedings of INFOCOM 96, 1996

[12] A. Medina, I. Matta, and J. Byers, "On the origin of power laws in Internet topologies," ACM Computer Communication Review, vol. 30, no. 2, pp. 18--28, Apr. 2000.

[13] B. Bollobas. "Random Graphs" Harcourt Brace Jovanovich, 1985.

[14] Duncan J. Watts and Steven H. Strogatz, "Collective dynamics of 'small-world' networks", Nature Vol 393 pp 440, 4 June 1998

[15] Ellen W. Zegura, Kenneth L. Calvert and Michael J. Donahoo. "A Quantitative Comparison of Graph-Based Models for Internet Topology". IEEE/ACM Transaction on Networking Vol 5. No 6. December 1997.

[16] P. Radoslavov, H. Tangmunarunkit, H. Yu, R. Govindan, S. Shenker, D. Estrin, "On characterizing network topologies and analyzing their impact on protocol design", USC-CS-TR-00-731, March 2000.

[17] Per Bak. "Self-Organized Criticality" Physica A (1990) 403-409

[18] M.G. Shnirman and E.M. Blanter "Self-Organized Criticality in a Mixed Hierarchical System". Physical Review Letters 14 December 1998, Vol. 81, Number 24. pp.5445-5448