Modeling Internet Physical Topology (Draft)

Nicolas SCHABANEL

CR CNRS – Laboratoire de l'Informatique du Parallélisme – ÉNS Lyon

(part of this document are joint works with Claire Kenyon and/or Ignacio Alvarez-Hamelin)

1 About the Internet

What is Internet? Internet is a hierarchical network composed of communication devices, the routers, interconnected with point-to-point links. The interconnection between the Autonomous systems composes the top level of the hierarchy. An autonomous system (AS) is basically a local network with a gateway to the outside (for instance, the local network of an university). Autonomous systems are of various types or, more precisely, belong to different domains (Stub domains, Transit domains, or Multi-homed domains). Basically, Transit domains interconnect Stub domains. The graph we will consider is the graph whose nodes are the ASs and whose edges are the link between them.

Links between autonomous systems are unknown: network companies keep their network structure secret. This motivates the need of a model: because the real network is unknown, one needs a model to design algorithms or calibrate structures for the Internet.

What do we know about the Internet topology? Measurements of Internet topology are mostly done using BGP tables or traceroutes. BGP tables are routing tables stored in routers to route the messages: it contains where to route (to which neighboring router) what according to its destination. Unfortunately the tables are mostly private (for the same reason as before). Traceroutes allow in some circumstances to trace the route followed by a message to its destination. Tracing a lot of messages to various destinations allows to get some picture of the Internet. Routeviews.org offers to the community the results of their monthly"crawls" on their website [14]. This is a precious tool for researchers. This picture is however partial: Shenker et al [7] estimate by analyzing known local neighborhoods that at least 25% to 50% of the links are missed in the current Internet measurements.

Broido and Claffy [3] took a picture of the internet based on traceroutes and BGP tables between 20 nodes which did accept to cooperate: they detected about 665,000 nodes in

2000 with an average degree[1] of about 4. The study of Faloutsos et al [9] revealed that the distribution of the degrees of the ASs in the Internet follows some kind of power law with exponent -2.2: the number of nodes with degree d is about $1/d^{2.2}$. This was unexpected by the network community. This initiated the quest for new network models.

On the dynamics side, Routeviews provides very useful datas. From 1998 to 2000, about 200 to 400 new ASs appeared in Internet every month, while 20 to 150 ASs dispeared every month. For the same period of time, about 400 to 1000 new links were created every month, while 200 to 800 links dispeared. When a new AS was created its initial degree is most of time 1, some times 2, rarely more. The same fact holds for the dead ASs.

2 Some Models for the Internet

Various models have been designed for the Internet topology. The very first models based on traditional random graphs produced disconnected graphs or exponential law on degrees for the desired proportion of links. Here is a list of models that produce connected graph that verify power laws on vertex degrees.

- Power law random graph by Aiello, Chung and Lu [1],
- Brite by Medina, Matta and Byers [13],
- Inet by C. Jin and Q. Chen and S. Jamin [11],
- Preferential attachment (Rich get richer) by Albert and Barabasi [2],
- GPL by Bu and Towsley [4],
- Nem by Magoni and Pansiot [12],
- HOT by Fabrikant, Kousoupias ad Papadimitriou [8].

All these models build different graphs. We will present some of them. For instance the now classical Albert and Barabasi [2] model builds a graph by adding nodes one by one: when inserted, the node is linked to a fixed number of already present nodes chosen randomly, proportionally to their current degrees. This model is based on the "rich get richer" or "preferential attachment" principle. The model in [8] constructs a rooted tree (this is thus as is, a very poor model for the internet): nodes are added one by one; each inserted node is linked to a single node, the node that minimizes a linear tradeoff between the hop distance to the root and the euclidean length of the link. Both models yield power laws on degrees. As is, because it constructs a tree, the later model is useless to model the internet; interestingly enough, it captures an important idea that will be discussed in Section 4, and it can be generalized to realistic models. Note that none of these models

[1] The degree of a node in a graph is the number of its neighbors.

takes into account the particular hierarchy of Internet autonomous systems. The other models try to include this hierarchy.

3 Observations on Internet measures and models

Remarks on target parameters. Most of the Internet topology models aim to obtain as closely as possible the value of the parameters (approximately) measured on the Internet: mainly, the exponent of the power law on degrees, the average clustering coefficient, the average path length. A lot of critics can be addressed to this approach:

Exponent of the power law on degrees. The computation of the exponent is subject to large error: small variation on the distribution has a huge impact on the exponent computation. Consider for example, a power law $1/i^2$ for i from 1 to 100. When computing the exponent of the closest power law using linear regression in the log-log space, a shift of the data of 0.5 to the left or the right induces a variation from 1.85 to 2.19 for the exponent (we use MatLab for these computations). Therefore, focusing on the exact measured exponent can not be a reasonable aim for model designers.

Clustering coefficient. The clustering coefficient of a node is the probability that two of its neighbors are connected. One striking fact on social networks is that their average clustering coefficient are "high" (at least higher than expected according to traditional random graphs). Some think that it may be an explanation of small-world phenomenon. An important critic regarding focusing on clustering coefficient is that the clustering coefficient essentially measures the proportion of triangle-like neighborhood in the graph: this coefficient drops quadratically with the degree of the nodes and has thus significant value only for small degree nodes. One can also construct graphs with high clustering coefficient that are not small-worlds. A reasonable guess to explain the relatively large value of the clustering coefficient is that social networks have a lot of very small degree nodes with high clustering coefficient (Internet graph for instance is full of triangle-like neighborhood).

mean eccentricity. Social graphs are often small-world graphs and thus have very small diameters (from 5 to 20). In that case, most of the nodes have an eccentricity of about half the diameter, independently of the real graph.

As a conclusion, these parameters may not be useful to discriminate between good and bad models for social networks. An important axis of research is certainly to find relevant and discriminating parameters to measure on networks to obtain consistent results.

A careful study of the validity of a model. The first important step in network modeling was made by Albert and Barabasi. They proposed in [2] a simple process based on the principle "rich get richer" that constructs a graph verifying a power law on vertex degrees. Other models have been proposed since but because of its simplicity, this model and its variant are widely studied. As mentioned above, analyzing the usual parameters is not really relevant. A recent study by Shenker et al [7] proposed a way around to determine whether Albert and Barabasi model is a good model for Internet topology: they confronted the dynamics of the model to the measured dynamics of the Internet. Based on nodes and links births and deaths over time, their careful analysis concludes that the dynamics are different: the model Albert and Barabasi based on "linear preference" is not a good model for the internet.

In this study, Shenker et al [7] also point out that measurement of the Internet are very partial: their analysis of local neighborhoods in Internet topology allows them to conclude that in the current Internet measurements, at least 25% to 50% of the links are missed. Furthermore when these links are added, the distribution on degrees is not a strict power law anymore as in [2].

4 HOT models

It is well-known that power laws (or more precisely heavy tailed distributions) are pretty common when one does statistics over links, objects or structures built by social or human individuals. Examples are the distributions of : the number of coauthors (degree in cocitation graph), the degree of routers in Internet, the number of inlinks and outlinks in the URLs graph (known as web graph), or even the size of forest fires in US Fish and Wildlife Service Lands, etc... This fact is usually stated as "Power laws are the signature of human activities". The "rich get richer" principle in [2] does not really explain this fact, since it is not really a social value in any human society (even if it can be claimed as a fact anyway).

Carlson and Doyle [5, 6] propose another explanation (Highly Optimized Tolerance): Power laws appear when one tries to optimize an objective function subject to tolerance to failure. For example, firemen try to draw the minimum number of roads thru the forests to limit fire spreading. This approach has been recently restated by Fabrikant, Koutsoupias and Papadimitriou [8] as follows (Heuristically Optimized Tradeoff): power laws may appear when one greedly optimizes a "balanced" tradeoff between two objective functions. They show on a very basic model for network growth (it constructs a rooted tree) that if one links every new node to the node minimizing a linear combination of the hop distance to the root and of the euclidean length of the link, one gets a graph with an heavy tailed law on the degrees of the nodes when the linear tradeoff between the two objective functions are balanced, while it is an exponential law when the linear tradeoff is very unbalanced. This tree generating process may seem basic at first, but as mentioned by the authors it has a lot of natural generalizations that yields "realistic looking" networks.

With Kenyon, we are currently working on extending this approach to more realistic models for the Internet. We have shown that the phase transitions still occur if two links are drawn instead of one at the creation of the nodes. With Alvarez-Hamelin, we have also simulated a more sophisticated variant where between nodes insertions, a number of links that minimizes the tradeoff, are drawn between existing nodes. All these models produce also to power laws on degrees. And they also “look” reasonnable.

5 Defining the good questions about Internet modeling

The main question remains open: how to validate a model of the Internet? One way may be to ask: why are we looking for models of the Internet? The answer to this question will certainly help to determine proper parameters to measure.

Since Internet models will a priori be used to simulate networks behavior, a pertinent parameter would be the “network behavior” of the model. An example of such a parameter would be: are the BGP tables close to the one in the Internet? What does close mean for BGP tables? Can we handle these notions theoretically?

Another approach is given by [7]: includes the dynamics in the study. A good model for the Internet should produce from the graph of the Internet in 1998, a graph “similar” to the graph of Internet in 2000.

References

[1] William Aiello, Fan Chung, and Linyuan Lu. A random graph model for massive graphs. In *ACM Symposium on Theory of Computing (STOC)*, pages 171–180, 2000.

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

[3] A. Broido and Mc Claffy. Internet topology: Connectivity of ip graphs. In *SPIE ITCom WWW conf.*, 2001.

[4] T. Bu and D. Towsley. On distinguishing between internet power law topology generators. In *INFOCOM*, 2002.

[5] J. M. Carlson and John Doyle. Highly optimized tolerance: A mechanism for power law in designed systems. *Physical Review E*, 60(2):1412–1426, aug 1999.

[6] J. M. Carlson and John Doyle. Power laws, highly optimized tolerance, and generalized source coding. *Physical Review Letters*, 84(24):5656–5659, jun 2000.

[7] 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*, 2002.

[8] Alex Fabrikant, Elias Koutsoupias, and Christos H. Papadimitriou. Heuristically optimized trade-offs: A new paradigm for power laws in the internet. *LNCS*, (2380):110–, Jun 2002.

[9] Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos. On power-law relationships of the internet topology. In *SIGCOMM*, pages 251–262, 1999.

[10] Ramesh Govindan and Hongsuda Tangmunarunkit. Heuristics for internet map discovery. In *IEEE INFOCOM 2000*, pages 1371–1380, Tel Aviv, Israel, March 2000. IEEE.

[11] C. Jin, Q. Chen, and S. Jamin. Inet: Internet topology generator. Technical Report CSE-TR443-00, Department of EECS, University of Michigan, 2000.

[12] Damien Magoni and Jean-Jacques Pansiot. Internet topology modeler based on map sampling. In *ISCC'02 - 7th IEEE Symposium on Computers and Communications*, pages 1021–1027, July 2002.

[13] A. Medina, I. Matta, and J. Byers. On the origin of power laws in internet topologies. *Computer Communications Review*, 30(2):18–28, apr 2000.

[14] Routeviews.org. Route views archives. http://archive.routeviews.org/.