Network Topologies, Power Laws, and Hierarchy

Hongsuda Tangmunarunkit, USC-ISI
Ramesh Govindan, USC-ISI
Sugih Jamin, Univ. of Michigan
Scott Shenker, ACIRI
Walter Willinger, AT&T

It has long been thought that the Internet, and its constituent networks, are *hierarchical* in nature. Consequently, the network topology generators most widely used by the Internet research community, GT-ITM and Tiers, create networks with a deliberately hierarchical structure. However, recent work by Faloutsos *et al.* revealed that the Internet's degree distribution — the distribution of the number of connections routers or Autonomous Systems (ASs) have — is a power-law. The degree distributions produced by the GT-ITM and Tiers generators are not power-laws.

To rectify this problem, several new network generators have recently been proposed that produce more realistic degree distributions; these new generators do not attempt to create a hierarchical structure but instead focus solely on the degree distribution. There are thus two families of network generators, *structural* generators that treat hierarchy as fundamental and *degree-based* generators that treat the degree distribution as fundamental. The guiding principles of the two families of generators are very different; the degree distribution is a very local property, having to do with the numbers of connections at individual nodes, while hierarchy is very nonlocal and describes the overall structure of the network.

In this research [1], we attempt to gain more understanding about the topological structure of the Internet. Underlying our investigation is the issue of whether the hierarchical structure or the degree distribution is more fundamental in modeling the Internet. More specifically, we address the following questions:

Question #1 *Which generated networks most closely resemble the Internet?* To answer this question, we use several topology metrics to compare networks produced by these two families of generators to current measurements of the Internet graph—the router-level (RL) and Autonomous System (AS) connectivity maps. These metrics are intended to capture some basic *large-scale* or *overall* properties of network structures, as opposed to purely local properties; we make a judgment that the large-scale properties are more fundamental than the local properties. The degree distribution is a rather local property. Even though the structural generators do not match the degree distribution of the Internet topologies, they could nonetheless create networks with large-scale properties that were similar to the Internet graphs.

Our results suggest that the AS and RL graphs have similar properties and the degree-based generators are significantly better at representing the Internet than the structural generators. This leaves us with the seeming paradox that while the Internet certainly has hierarchy, it appears that the Internet graphs are better modeled by network generators that completely ignore hierarchy! Resolving this paradox leads us to our second question.

Question #2 *Is there any relationship between hierarchical structure and power-law degree distributions?* We introduce a novel measure of hierarchy, and use this to investigate the nature of hierarchy in the generated and measured networks. We find that while the degree-based generators do not explicitly inject hierarchy into the network, the power-law nature of the degree distribution results in a substantial level of hierarchy—not as strict as the hierarchy present in the structural generators, but significantly more hierarchical than random graphs. This relatively loose form of hierarchy, produced merely by the presence of the power-law degree distribution, more accurately reflects the nature of hierarchy in the Internet than the strict hierarchy produced by the structural generators.

Question #3 *Where does the power-law degree distribution in the AS topology come from?* [2] From an Internet router-level topology, an AS-overlay can be computed. This computation allows us to estimate the AS size distribution, where size is defined by the number of routers in the AS. We find that the distribution of AS sizes exhibits a power-law. Moreover, there is a strong correlation between AS size and degree. Based on the ubiquity of highly-variable size distributions in real-world entities such as cities by population size, file size, web document, etc., we conjecture that the power-law degree distribution in the AS topology may simply follow from its power-law size distribution.

The work presented here is preliminary, and is only a first step towards an understanding of the issues we've addressed. Thus far, our metrics and analysis rely on shortest paths on the graph. We plan to incorporate policy constraints into our work especially in the context of measured networks.

*This work was supported in part by the Defense Advanced Research Projects Agency under grant F30602-00-2-055. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the Defense Advanced Research Projects Agency.

[1] http://www.isi.edu/~hongsuda/publication contains full papers.

[2] In collaboration with John Doyle, Caltech.