

Topology Generation based on Network Design Heuristics*

Bruno Quoitin

CSE Department, Université catholique de Louvain, Belgium

quoitin@info.ucl.ac.be

Categories and Subject Descriptors: C.2.1 Network Architecture and Design: Network topology

General Terms: Measurement, Performance, Design

Keywords: Topology generation, network design heuristics

1. INTRODUCTION

Building increasingly precise and realistic network topologies is an important issue for the purpose of evaluating networking applications. Still, the generation of router-level topologies has not been widely covered. Though, the properties of the router-level topologies have a significant impact on simulation results. For instance, the evaluation of applications such as Voice/Video over IP, P2P, routing protocols and traffic engineering methods critically depends on the properties captured by the topology model.

We believe that today's network topologies are the fruits of a careful design taking into account practical constraints. We argue that it is possible to generate realistic network topologies by reproducing and automating the work of a human network designer. Due to the computational complexity of network design, heuristics are often used to build networks. The networking literature contains lots of network design methods which are currently seldom used by researchers. One reason is that few algorithms are implemented in publicly available tools. We propose an open-source tool [1] that can be used to generate network topologies based on network design methods. Our tool currently implements a variety of network design heuristics and allows to play with a large spectra of parameters leading to different network designs.

2. RELATED WORK

Several approaches to obtain router-level topologies have been used by researchers. The first and most natural approach was to rely on existing network topologies. This approach is limited due to the difficulty of obtaining the topology of operational networks today. Proposals such as [12] exist for inferring network topologies at the router-level from measurements. Unfortunately, such methods are known to sometimes miss multiple paths between routers or to infer links and routers that do not really exist [9, 13].

*This work was supported by the Walloon Government within the WIST TOTEM project (<http://totem.info.ucl.ac.be>) and partially by the E-NEXT European Network of Excellence. We thank Thomas Telkamp for providing the GBLX topology and Nicolas Simar for the Geant topology. We also thank Jean Lepre, Cristel Pelsser and Steve Uhlig for their comments.

Another approach consists in generating synthetic topologies. One of the most used router-level topology generators is GT-ITM [4]. It allows to build router-level topologies with a backbone/access hierarchy. Nodes are placed randomly and connected using a probabilistic model. BRITE [11] is another topology generator able to produce router-level topologies.

The problem with the approach followed by GT-ITM and BRITE is that the topologies are generated in order to mimic pure graph properties of real networks. They fail to capture the optimization process that is also at the basis of the real network topologies. In [10], Li et al show that topologies generated without taking into account economical and technical constraints perform poorly and predict that future topology generators should not be built on pure graph-theoretic properties but upon more pragmatic properties such as the maximum throughput that can be achieved by the network and its resilience to failures. In [2], Alderson et al propose to formulate the network design problem as an optimization problem which takes as input a traffic demand and produces a router/host level topology. However, to date, no topology generator implements the proposed approach.

3. MEASURING REAL NETWORKS

Real networks are the result of a complex design process. The network design problem consists of multiple, sometimes contradictory objectives. No single optimal solution exists, rather a front of possible solutions. The network design problem has been fairly discussed in the literature, in particular by [3, 7]. Its objectives may be summarized in **minimizing the latency**, **dimensioning the links** and **adding redundancy** so that rerouting is possible in case of failure. Finally, the network must be designed at **the minimum cost**. None of these objectives are currently explicitly found in degree-based generators.

In order to build more realistic network topologies, we need to better understand the characteristics and the structure of real networks. For this purpose, we have selected a set of metrics from the networking literature. These metrics capture various aspects of the network design problem. They cover performance properties (delay, throughput) and pure graph properties¹.

A first metric is the distribution of the **distances between pairs of nodes** (δ) along the shortest-path route. It is an indication of the delay required to transmit packets between these nodes under the assumption that the largest part of the transmission delay is due to the propagation delay along the links. We measure the distribution of the distances between pairs by measuring the length of the shortest-paths between all pairs of nodes.

To measure the amount of redundancy offered by a network, we use the **path diversity** metric (ρ) [13]. This metric measures the

¹We present a subset of our metrics and refer the reader to [1].

availability of diverse paths between pairs of vertices. This is important for network robustness and traffic engineering. The diversity of paths between a pair of nodes is obtained by computing the shortest-path between these nodes. Then, removing this path from the graph and computing the second shortest-path, and so on until no more path can be found.

We also measure the connectivity of the graph through the distribution of the **node degrees** (d). The distribution also gives information on the existence of hubs, which are nodes with a high degree, where many other nodes connect.

We show in Table 1 the above metrics applied on a set of networks containing large carrier networks as well as regional networks: Abilene, the US research network, ISP-A and ISP-B which are US backbone networks, GBLX-US which is the US part of GBLX, Géant the European research network as well as Eu-ISP and Tiscali, two European networks.

	$ V $	$ E $	$\bar{\delta}$	$\bar{\rho}$	\bar{d}
Abilene	11	14	25.6	1.63	2.55
Géant	23	37	21.8	2.43	3.22
Eu-ISP	53	98	NA	1.97	3.69
Tiscali	39	52	16.7	1.39	2.67
GBLX-US	41	77	26.6	1.79	3.76
ISP-A	20	44	26.9	3.03	4.4
ISP-B	20	44	28.9	3	4.4

Table 1: Comparison of network metrics on various real networks.

Based on Table 1, one can assert that real networks sample the spectrum of several metrics. For instance, backbone networks that cover the United States (Abilene, ISP-A and ISP-B) have different sizes (number of vertices and edges) and different performance results. Abilene is quite sparse. ISP-A and ISP-B have a high path diversity ($\bar{\rho} \simeq 3$). The average node degree differs largely among all the topologies, ranging from 2.55 for Abilene to 4.4 in ISP-A and ISP-B.

4. NETWORK TOPOLOGY GENERATOR

In this section, we present our topology generation tool. The approach follows the tasks of a network designer. Starting with a set of geographically located nodes to interconnect and a demand matrix, the tool produces a plausible topology.

Basically, the tool uses a bottom-up approach. The tool first groups the nodes into clusters which are the Points-of-Presence (POPs) of the network. The clustering methods used to group nodes into POPs are based on the euclidian distance between the nodes, based on the traffic demand or a combination of both. The tool implements k-medoids, Ward’s hierarchical clustering algorithm as well as a threshold-based method [3].

Then, the structure of each POP is built. We propose a realistic approach to generate a POP structure which is similar to real network designs. We generate POP structures that resemble Sprint [8] or AOL networks [6]. For each POP, we select the n most central nodes as backbone routers. The backbone nodes of a POP are densely connected together using for instance a tour that guarantees 2-edge-connectivity or a clique. Then, the remaining nodes of the POP, which model access nodes are connected to the POP’s backbone nodes using at least k edges. Using $k \geq 2$ guarantees redundancy in case of failure.

Once each POP has been generated, the tool produces a design for the backbone. The backbone is a graph which will interconnect the backbone nodes of the POPs. The tool proposes various heuris-

tics to generate the backbone. The first network design heuristic we consider is known as **MENTOR** [3] and builds a hybrid minimum spanning tree/shortest-path tree (MST-SPT). MENTOR is close to the Heuristically Optimized Trade-offs (HOT) proposed in [5]. Since trees are weak networks, another heuristic called **MENTour** [3] can be used. This heuristics directly builds a 2-edge-connected network using a TSP approximation heuristic. Another way to produce a 2-edge-connected network is the **Two Trees (2T)** [7] which builds 2 MSTs. Finally an interesting mesh generation technique consists in computing a **Delaunay triangulation** of the backbone nodes. It produces a topology with alternate paths between nodes, while minimizing the number of such paths. This is an efficient way of obtaining a cost-effective topology with redundancy. The tool implements other mesh-generation techniques that we do not describe here due to space limitations.

We show in Table 2 an example of four backbone designs based on the Géant POPs. The name of the heuristic in use is written in the first column. The other columns give performance metrics as described in Section 3. The δ metric is given in kilometers.

	$ E $	\bar{d}	$\bar{\delta}$	$\max \delta$	$\bar{\rho}$	$\min \rho$
Géant	37	3.22	2167	9215	2.43	2
MENTOR 0.3	22	1.91	2107	9529	1	1
MENTour	23	2	4865	11751	2	2
TwoTrees	44	3.83	1801	9243	2.80	2
Delaunay	63	5.48	1737	9111	4.93	4

Table 2: Performances of various designs for Géant.

Finally, the tool can be used to assign IGP weights and capacities to links. Different assignment schemes are possible. We can set link weights proportional to the link length for instance. For capacities, one possible method consists in assigning link capacities that will ensure that the demand matrix can be accommodated. This link capacity assignment method computes the shortest-paths between all pairs of nodes, then routes the demands on these paths and determines the load of each link. It can therefore select the link capacity of each link in order to limit the maximum link utilization to a predefined level τ .

5. REFERENCES

- [1] B. Quoitin. IGen. <http://www.info.ucl.ac.be/~bqu/igen>, 2005.
- [2] D. Alderson, J. Doyle, R. Govindan, and W. Willinger. Toward an Optimization-Driven Framework for Designing and Generating Realistic Internet Topologies. *ACM SIGCOMM CCR*, 33(1), January 2003.
- [3] R. S. Cahn. *Wide Area Network Design: Concepts and Tools for Optimisation*. Morgan Kaufmann, 1998.
- [4] K. Calvert, J. Eagan, S. Merugu, A. Namjoshi, J. Stasko, and E. Zegura. Extending and Enhancing GT-ITM. *ACM SIGCOMM*, 2003.
- [5] A. Fabrikant, E. Koutoupas, and C. H. Papadimitriou. Heuristically Optimized Trade-offs: A New Paradigm for Power Laws in the Internet. *International Colloquium on Automata, Languages and Programming*, 2002.
- [6] V. Gill and J. Mitchell. AOL Backbone OSPF-ISIS Migration. *NANOG29*, 2003.
- [7] W. D. Grover. *Mesh-Based Survivable Networks*. Prentice Hall PTR, 2004.
- [8] G. Iannaccone, C.-N. Chuah, S. Bhattacharyya, and C. Diot. Feasibility of IP Restoration in a Tier-1 Backbone. *IEEE Network*, 18(2):13–19, March 2004.
- [9] A. Lakhina, J. W. Byers, M. Crovella, and P. Xie. Sampling Biases in IP Topology Measurements. *IEEE INFOCOM*, 2003.
- [10] L. Li, D. Alderson, W. Willinger, and J. Doyle. A First Principles Approach to Understanding the Internet’s Router-level Topology. *ACM SIGCOMM*, 2004.
- [11] A. Medina, A. Lakhina, I. Matta, and J. Byers. BRIT: An Approach to Universal Topology Generation. *MASCOTS 2001*, 2001.
- [12] N. Spring, R. Mahajan, and D. Wetherall. Measuring ISP Topologies with Rocketfuel. *ACM SIGCOMM*, 2002.
- [13] R. Teixeira, K. Marzullo, S. Savage, and G. M. Voelker. Characterizing and Measuring Path Diversity of Internet Topologies. *ACM SIGMETRICS*, 2003.