Generating Realistic ISP-Level Network Topologies

Oliver Heckmann, Michael Piringer, Jens Schmitt, and Ralf Steinmetz

***Abstract*—Simulations are an important tool in network research. As the selected topology often influences the outcome of the simulation, realistic topologies are needed to produce realistic simulation results. Using several similarity metrics to compare artificially generated topologies with real world topologies this letter gives hints how to use the wide-spread topology generators BRITE, TIERS, and GT-ITM to create realistic topologies.**

***Index Terms*—Communication networks, simulation, topology.**

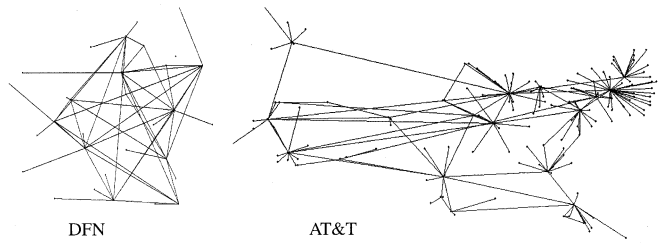

Fig. 1. The DFN and AT&T topologies.

I. Introduction

THE wide-spread topology generators BRITE [1], TIERS [2], and GT-ITM [3] offer a big range of configuration parameter. GT-ITM, for example, has 16 different configuration parameters (for the transit-stub model). How realistic a generated topology is depends on the combination of these parameters. Usually the generated topologies are judged realistic or not by pure visual inspection. In this letter, we define objective criteria (similarity metrics). Based on those we search for parameter combinations of the generators mentioned above to generate topologies that are similar with respect to the metrics to two real world ISP topologies. Those real world topologies are: 1) the rather large U.S. AT&T continental IP backbone and 2) the smaller DFN G-Win (German research network), see Fig. 1. With these results we can compare how realistic artificially created topologies are and derive parameter combinations for the generators. They can act as a starting point for anyone who wants to do ISP level simulations using topology generators.

Similar work was done in [4] on AS level graphs with at least 1000 nodes in order to evaluate topology generators for AS level graphs.

The paper is structured as follows: In Section II we present our similarity metrics. After that we present as results the best combinations for the two example topologies and three generators. We conclude with a short summary and an outlook.

II. Similarity Metrics

To measure the similarity of two network topologies we define the following metrics that capture the basic connectivity properties of the topology graph. We are interested in graphs with the same connectivity properties but not in equivalent graphs. In the graphs we distinguish between edge nodes (which are connected to end-users and other networks) and core/backbone nodes (which are only connected to nodes of the same network). We define the following metrics.

1) The first metric uses the hop-plot of all nodes. For each graph g we look at all n nodes and calculate how many other nodes can be reached within $h = 1, 2, 3\ldots$ hops. From this we derive the relative frequency distribution F_h^g. We then compare the frequency distributions of both graphs.
2) The second metric F_h^{ge} is similar to the first but only looks at edge-nodes.
3) Next, from the outdegree d_i of each node i we derive the relative frequency distribution of all nodes for both graphs, and use the significance level of a Wald–Wolfowitz test for the similarity of the two distributions.
4) We also compare the rank exponent $\Re$.
5) The outdegree exponent O of the first and second power-laws are as defined in [5].
6) We also used the relative difference $|n_{\text{generated}} - n_{\text{reference}}/n_{\text{reference}}|$ in the number of nodes.
7) The relative difference in the number of links are considered as additional metrics.

To express the difference in two distributions we sum up the accumulated absolute difference over all classes. Every metric is normalized to return a value between 0 and 1 with 1 resembling the highest similarity. All metrics are added to a combined metric and the result is normalized again.

We used a heuristic similar to Hook and Jeeves [6] to search for the parameter combination that yields the maximum combined similarity metric. If multiple topologies are created with the same parameter combination the resulting combined similarity metric varies less than 2% for all tested topology generators.

III. Results

The parameters of Table I were found for Brite and the DFN and AT&T topologies with a high and satisfying combined and normalized similarity of 0.972 rsp. 0.951. Please note that the values α and β do not seem to significantly influence the outcome of the measurements when the parameter links/node is set to 2.

Manuscript received December 20, 2002. The associate editor coordinating the review of this paper and approving it for publication was Prof. C. Douligeris. This work is supported in part by the German research network provider DFN (http://www.dfn.de) as part of the LETSQoS project (http://www.letsqos.de/).

The authors are with the Multimedia Communications Lab (KOM), Department of Electronic Engineering and Information Technology, Darmstadt University of Technology, 64283 Darmstadt, Germany (e-mail: Heckmann@kom.tu-darmstadt.de; Piringer@kom.tu-darmstadt.de; Schmitt@kom.tu-darmstadt.de; Steinmetz@kom.tu-darmstadt.de).

Digital Object Identifier 10.1109/LCOMM.2003.814708