

Generating Representative ISP Topologies From First-Principles

Chong Wang
Computer Science Department
Boston University
Boston, MA 02215
wangch@cs.bu.edu

John Byers
Computer Science Department
Boston University
Boston, MA 02215
byers@cs.bu.edu

ABSTRACT

Understanding and modeling the factors that underlie the growth and evolution of network topologies are basic questions that impinge upon capacity planning, forecasting, and protocol research. Early topology generation work focused on generating network-wide connectivity maps, either at the AS-level or the router-level, typically with an eye towards reproducing abstract properties of observed topologies. But recently, advocates of an alternative “first-principles” approach question the feasibility of realizing representative topologies with simple generative models that do not explicitly incorporate real-world constraints, such as the relative costs of router configurations, into the model. Our work synthesizes these two lines by designing a topology generation mechanism that incorporates first-principles constraints. Our goal is more modest than that of constructing an Internet-wide topology: we aim to generate representative topologies for single ISPs. However, our methods also go well beyond previous work, as we annotate these topologies with representative capacity and latency information. Taking only demand for network services over a given region as input, we propose a natural cost model for building and interconnecting PoPs and formulate the resulting optimization problem faced by an ISP. We devise hill-climbing heuristics for this problem and demonstrate that the solutions we obtain are quantitatively similar to those in measured router-level ISP topologies, with respect to both topology properties and fault-tolerance.

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

General Terms: Algorithms.

Keywords: Network topology modeling, network design, optimization.

1. INTRODUCTION

Topology modeling has significant impact on other aspects of network research. A prevalent aim of topology research was to match observed graph-theoretic properties of the underlying network. However, a first-principles theory [3] proposed recently challenges this approach. The first-principles theory attempts to model the engineering issues faced when a network is designed and therefore can both reflect certain

properties found in the network and explain the process by which a network evolved.

Motivated by an optimization-driven framework [1], we attempt to build a topology generation model that starts from first-principles explanation, and use it to produce topologies for single ISPs. In our framework, the topology of an ISP can be viewed as the result of a concrete optimization problem: namely, a network design problem in which costs of deploying equipment (links, various router configurations and various PoP configurations) are known and the objective is to build a capacitated network to satisfy an initial customer demand specified as a traffic demand matrix. We show that from a simple and natural optimization framework with relatively few inputs and parameters, networks similar to measured ISP topologies from various aspects can be heuristically generated. Though proprietary information of ISPs limits the extent to which we can quantify discrepancies between our model and those used in reality, we argue that our model captures many of the scaling costs present in large-scale ISP design.

2. PROBLEM STATEMENT

In a two-dimensional plane, we are given a set of client locations C , a set of candidate facility locations F , and a traffic demand matrix $T_{|C| \times |C|}$ defined on $C \times C$. Denote the traffic demand between i and j by $t_{i,j}$.

The objective is to select a set of facilities $S \subseteq F$, a set of edges $E \subseteq (S \times C) \cup (S \times S)$ and to build an undirected graph $G = (S \cup C, E)$ with the minimal value of $cost(G) = \sum_{v \in S \cup C} cost(v) + \sum_{e \in E} cost(e)$.

Let the function $B : E \rightarrow R^+$ denote the capacities of edges. We model the costs of edges (links) and nodes (routers) in the following way: for an edge, based on the commonly used bandwidth-delay product, we model the cost of an edge as $cost(e) = (b + B(e)^{d_1})(l(e))^{d_2}$, $b > 0$, $d_1, d_2 \in (0, 1]$, $e \in E$. Here, $l(e)$ is the Euclidean distance between the nodes incident to e , while d_1 and d_2 model the decrease of the marginal cost per unit length (resp. bandwidth) as the length (resp. bandwidth) increases. For a node, we model the cost as linear in the sum of the adjoining link capacities, i.e. $cost(v) = a \sum_{(u,v) \in E} B(u,v)$, $a > 0$, $v \in S \cup C$.

The graph is subject to the following constraints:

1. Clients are single-homed: $\forall c \in C$, $degree(c) = 1$.
2. At least one path exists between every pair of clients. Denote the shortest path between i, j under OSPF by $P_{i,j}$.
3. The network has sufficient capacity to carry the traffic demands across shortest paths: $\forall k, \sum_{(i,j) \mid e_k \in P_{i,j}} t_{i,j} \leq B(e_k)$.




Figure 1: Measured and modeled network topologies (AS 7018) of 1 representative experiment.

network type	density	clustering coefficient	avg. distance	diameter	betweenness centrality	closeness centrality
RocketFuel	0.035	0.325	3.32	6	35.64%	32.82%
Two-phase Optimization	0.0376(0.0007)	0.181(0.048)	3.58(0.09)	7.63(0.52)	57.67%(7.97%)	38.73%(3.57%)
Incremental Approach	0.0366(0.0004)	0.132(0.012)	3.86(0.01)	8(0)	61.05%(0.48%)	36.78%(0.31%)

Table 1: Graph-theoretic metrics of networks (AS 7018). Experimentally computed values depict the average of 8 experiments followed by the standard deviation. The values of parameters are: $a = 0.2797$, $b = 825000$.

3. APPROXIMATION ALGORITHMS

A natural approach to solve the problem is to aggregate the customer demands at PoPs (Points of Presence) and optimize the PoP-level graph. Doing so avoids the high cost of long-haul connections between small routers. So we take a 2-phase approach to solve the optimization problem.

Phase 1: Generate Base PoPs: connect every client with a PoP. This is an instance of the universal facility location problem: the cost of PoPs corresponds to facility cost in the facility location problem; the cost of edges corresponds to shipping cost. We use the term “base PoPs” to denote the PoPs generated in this first phase.

Phase 2: Optimize the PoP-level network: Starting from the set of base PoPs, with traffic demand aggregated from the customers, we optimize the PoP-level network. From an initial feasible solution, i.e. a connected graph annotated by the capacities determined by the constraints, we run a sequence of iterations. In an iteration we evaluate each possible operation from the set below and greedily apply the one that minimizes the cost of the resulting graph:

1. Add a capacitated edge that is not in the current graph.
2. Remove an edge that does not disconnect the graph.
3. Add a transit PoP to make traffic demands between different PoPs share links to the added transit PoP instead of using separate links.
4. Remove a transit PoP that does not disconnect the graph.

Employing a Priority Queue: In practice, re-evaluating the cost of every operation at each iteration is overkill. The cost of many operations does not change, and the overall rank of most operations does not change dramatically. Therefore, we manage the set of possible operations in a priority queue, where we re-evaluate the highest priority operations constantly and sample the lower-priority ones.

Incremental Approach: An alternative to phase 2 is an *incremental* strategy: insert the base PoPs into the network in batches, ordered by the amount of customer traffic demands that they serve. After each insertion we re-optimize the intermediate network to minimize the cost. This approach has two benefits: 1) we have observed that this method leads to faster experimental convergence, and 2) this approach mimics aspects of historical evolution of networks.

4. EXPERIMENTAL RESULTS

We apply our approach to generate different topologies. In our model, d_1 and d_2 are set to 0.96 and 0.9 respectively. The remaining parameters are a and b in our problem statement. In practice, a is a deciding factor of the number of base PoPs generated by our first phase. So we adjust a to generate the topologies whose sizes are close to the measured ISP topologies [2] and compare various aspects of the measured topologies with the modeled ones. Figure 1 shows the visual comparison of the measured and modeled topologies of AS 7018 (AT&T), along with the modeled costs of the topologies. In Table 1 we report some graph-theoretic metrics. We can see that the modeled topologies show a reasonable similarity with the measured ones. See [4] for a detailed comparison and analysis of other metrics as well as experiments on other topologies and settings.

5. CONCLUSION

Our work treats the problem of building an ISP topology as a natural network design problem. We demonstrate that it is possible to apply various heuristics to prune the search space to enable generation of moderately large topologies. The resulting topologies are visually and quantitatively similar to topologies measured by Rocketfuel. With our model, we generate *annotated* topologies, which is crucial to network research but not available in current topology generators and measured topologies. Another power of our model is its ability to explain the generation process of networks and predict the topologies if some conditions change.

6. REFERENCES

- [1] D. Alderson and J. Doyle and R. Govindan and W. Willinger, “Towards an optimization-driven framework for designing and generating realistic Internet topologies”, *ACM Computer Communication Review*, 33(1), 2003.
- [2] N. Spring and R. Mahajan and D. Wetherall and T. Anderson, “Measuring ISP topologies with Rocketfuel”, *IEEE/ACM Transactions on Networking*, 12(1):2-16, 2004.
- [3] L. Li and D. Alderson and W. Willinger and J. Doyle, “A first-principles approach to understanding the Internet’s router-level topology”, In *Proceedings of ACM SIGCOMM*, August 2004.
- [4] C. Wang and J. W. Byers, “Generating Representative ISP Topologies From First-Principles”, Tech. Report BUCS-TR-2007-001, Boston University, March 2007.