Toward an Optimization-Driven Framework for Designing and Generating Realistic Internet Topologies

David Alderson
Management Science
and Engineering
Stanford University
alderd@stanford.edu

John Doyle
Control and Dynamical Systems
California Institute of Technology
doyle@cds.caltech.edu

Ramesh Govindan
Computer Science
USC
ramesh@usc.edu

Walter Willinger
AT&T Labs-Research
walter@research.att.com

ABSTRACT

We propose a novel approach to the study of Internet topology in which we use an optimization framework to model the mechanisms driving incremental growth. While previous methods of topology generation have focused on explicit replication of statistical properties, such as node hierarchies and node degree distributions, our approach addresses the economic tradeoffs, such as cost and performance, and the technical constraints faced by a single ISP in its network design. By investigating plausible objectives and constraints in the design of actual networks, observed network properties such as certain hierarchical structures and node degree distributions can be expected to be the natural by-product of an approximately optimal solution chosen by network designers and operators. In short, we advocate here essentially an approach to network topology design, modeling, and generation that is based on the concept of *Highly Optimized Tolerance (HOT)*. In contrast with purely descriptive topology modeling, this opens up new areas of research that focus on the *causal* forces at work in network design and aim at identifying the economic and technical drivers responsible for the observed large-scale network behavior. As a result, the proposed approach should have significantly more predictive power than currently pursued efforts and should provide a scientific foundation for the investigation of other important problems, such as pricing, peering, or the dynamics of routing protocols.

Keywords

Internet topology, network optimization, robustness, complex systems, highly optimized tolerance

1. INTRODUCTION

The Internet is a complex conglomeration of rapidly growing, fragmented, and competing networks. As a research project in its early days, the design, development and deployment of Internet infrastructure was coordinated among relatively few organizations, and the location, capacity, and interconnectivity of this infrastructure was known with relative certainty. However, after the decommissioning of the NSFNET backbone in 1995, when management of the Internet was given over to commercial entities, the number of Internet Service Providers (ISPs) and their corresponding infrastructure have grown dramatically. In this new market environment, a desire to preserve competitive advantage has provided incentives for ISPs to maintain secrecy about their network infrastructure, including detailed connectivity maps.

The resulting decentralized nature of the Internet and the complexity and diversity in the number of infrastructure owners and operators, coupled with the incentive for these providers to obscure their infrastructure topologies,[1] have made comprehensive knowledge of Internet connectivity increasingly difficult. As an example, consider the *AS* graph of the Internet which depicts the connectivity between individual autonomous systems (ASs). In such a graph, each node represents an AS, while a link between two nodes indicates that the two ASs have a "peering relationship" (i.e., there exists at least one direct router-level connection between the two ASs). In this sense, AS graphs reflect business relationships among an ever increasing number of ASs. While AS-level connectivity can, in principle, be inferred from BGP-derived measurements, the fully distributed and decentralized nature of BGP makes it very difficult to obtain "complete" AS-level connectivity [14]. A second type of graph describes Internet connectivity at the router-level. Here, nodes represent individual routers, links represent one-hop connectivity at the IP level between routers, and the resulting graphs reflect physical connectivity at the router-level as seen by IP.[2] While many providers view their router-level maps as containing proprietary and business-sensitive information, reconstructing router-level topologies of individual ASs or ISPs is, in principle, feasible and relies on information obtained from selective traceroute measurements. However, as in the case of AS-level connectivity, the available data are known to provide incomplete router-level maps, and aiming for more complete topologies remains an active area of research [19, 15, 28].

---

[1] Note that while in the present context the notion of "topology" is almost exclusively used to mean "connectivity" (i.e., links and nodes without further annotation), we use it here—whenever appropriate—to mean "connectivity" plus "resource capacity" (i.e., links and nodes are annotated with for example link speed, delays, router capacity).

[2] See also Section 2.4 regarding Level-2 technologies and our proposed framework.

Gaining a basic understanding of the existing and future Internet topologies and of their evolution over time is of primary importance for networking research. As pointed out in [30], although topology should not affect the *correctness* of networking protocols, it can have a dramatic impact on their *performance*. Topology is therefore important for the design and evaluation of networking protocols, and it may also play an important role in gaining a basic understanding of certain aspects of current large-scale network behavior. As the infrastructure of the Internet continues to evolve, trends in connectivity formation are likely to yield insight into future behavior and may suggest novel strategies for provisioning requirements, traffic engineering, and operations of tomorrow's networks.

The difficult, yet important task of designing, modeling, and generating realistic Internet topologies has attracted a great deal of research attention, especially within the last few years. The prevailing approach of most of these efforts has been to focus on matching a sequence of easily-understood metrics or observed features of interest; e.g., explicitly imposed connectivity properties or hierarchical structures (see [33] and references therein), empirical node degree distributions [21, 23, 1, 7], clustering coefficients [8], etc. (see for example [30] for additional candidate metrics). However, this type of *descriptive* or *evocative* modeling can be misleading, since the question of which metrics or features are the most important ones for judging and comparing different Internet topologies remains largely unresolved, and any particular choice tends to yield a generated topology that matches observations on the chosen metrics but looks very dissimilar on others.

These observations strongly argue for a radically different approach to designing and modeling Internet topologies, one that moves beyond evocative models and instead advocates the careful development and the diligent validation of *explanatory* models that concentrate on the *causal* forces at work in the design and evolution of real topologies. We propose formulating appropriate optimization problems to model the process by which Internet connectivity is established and evolves. The basic idea is that when deploying their infrastructures, network owners and operators are, in fact, approximately solving optimization problems (either explicitly or implicitly) that express the ways in which they build up and evolve their networks. To the extent that we can identify and formulate (even at a high level of abstraction) the objectives and constraints of these optimization problems, solving the latter can be expected to lead to the generation of realistic topologies (i.e., connectivity information as well as resource provisioning information), where the observed characteristics of the resulting graph structures are well explained and understood in terms of the underlying mechanisms that are directly reflected in the optimization formulations.

2. OUR APPROACH IN A NUTSHELL

Our approach seeks to capture and represent realistic drivers of Internet deployment and operation to create a topology generation framework that is inherently *explanatory* and will perforce be *descriptive* as well. Instead of fitting certain characteristics of measured Internet topologies, any such agreements with empirical observations would instead be evidence of a successful explanatory modeling effort. We take as the basic unit of study the solitary ISP since there are a large number of important networking issues—such as configuration, management, pricing, and provisioning—that are naturally and precisely relevant at the level of the ISP. An understanding of the key issues facing ISPs combined with the ability to generate "realistic, but fictitious" ISP topologies would greatly enhance the ability of networking researchers to address these important problems. A second reason for studying the topology of the solitary ISP is that, at an appropriate level of abstraction, the Internet as a whole is simply a conglomeration of interconnected ISPs. Our belief is that by understanding the forces driving topology evolution for the single ISP, one can make great progress to understanding Internet topology at a broader level.

2.1 Driving Forces: Econ and Tech

Our starting premise is that any explanatory framework for Internet topology modeling and generation needs to incorporate both *economic* factors and *technical* factors faced by ISPs. Because of the costly nature of procuring, installing, and maintaining the required facilities and equipment, the ISP is economically constrained in the amount of its physical plant that it can support. Indeed, the economic survival of the ISP depends on carefully balancing limited revenue streams with its capital expenditures. As a result, the buildout of the ISP's topology tends to be incremental and ongoing. At the same time, there are economic aspects of networking that are largely determined by "external" forces, but impose potentially major constraints on the design and evolution of ISP topologies. For example, ignoring economic realities (e.g. most customers reside in the big cities, most high-bandwidth pipes are found between big cities, or most national or global ISPs peer for interconnection in the big cities) can be expected to result in topologies that are too generic to be of practical value.

Similarly, the layout of ISP topologies reflects technical constraints imposed by physical or hardware realities. For example, while routers can only be directly connected to a limited number of neighboring routers due to the limited number of interfaces or line cards they allow, no such limitations exist per se when it comes to the number of peering relationships an ISP can enter in with competing ISPs. Other tech-based factors that may seriously constrain the interconnectivity of ISP topologies are certain Level 2 technologies (e.g., Sonet, ATM, WDM) or the availability and location of dark fiber. Collectively, these economic and technical factors place bounds on the network topologies that are feasible and actually achievable by ISPs.

Given that we require our approach to be driven by economic and technical factors, we next assume that many of these factors can be effectively captured and represented in a mathematical model using a combinatorial optimization framework. A challenge in using such a framework is to demonstrate that some of the crucial economic or technical factors can indeed be expressed in terms of some sort of combinatorial optimization problem. We also need to identify the appropriate optimization formulations (objectives, constraints, parameters) for representing a range of ISP behavior, from very generic to highly specific. Finally, we have to illustrate with convincing examples how knowing the causal forces expressed via these optimization formulations explains the properties of the resulting topology and advances our knowledge about the design and evolution of Internet topologies.

2.2 Modeling ISP Topology

In modeling the topology of an ISP, we are necessarily trying to construct a router-level graph that is consistent with the decisions being made by an ISP in the same market environment. Some key questions to be answered for this model include *What are the economic and technical factors facing individual ISPs?* and *What are plausible objectives, constraints, parameters for an optimization formulation?* This formulation can take one of several forms. In a *cost-based* formulation, the basic optimization problem is to build a network that minimizes cost subject to satisfying traffic demand. Alternatively, a *profit-based* formulation seeks to build a network that satisfies demand only up the point of profitability—that is, economically speaking where marginal revenue meets marginal cost.

No matter the formulation, one of the key inputs for this approach is a model for traffic demand. A natural approach to traffic demand is based on *population* centers dispersed over a *geographic* region. In this manner, one could derive the topology of a single "national ISP" from the demand for traffic between people across the country or across town. Furthermore, the tremendous size of a national ISP makes it often convenient to decompose the network into separate problems whenever possible. Most often, this decomposition comes in the form of network hierarchy. It commonly takes the form of backbone networks (wide area networks or WANs), distribution networks (metro area networks or MANs), and customers (local area networks or LANs). Using this approach, the size, location and connectivity of the ISP will depend largely on the number and location of its customers, and it is possible to generate a variety of local, regional, national, or international ISPs in this manner.

2.3 Modeling Internet Topology

Given the ability to effectively model the router-level topology of an ISP (including the placement of peering nodes or points of presence), issues about *peering* become limited to interconnecting the router-level graphs. The relevant questions in this context are *What are the economic and technical factors facing peering relationships between ISPs?* and as before *What are plausible objectives, constraints, parameters for an optimization formulation?* Here, it will be important to leverage previous work on the general economics underlying Internet peering relationships [4], optimal location of peering points between peers [3], and the gaming issues of interdomain traffic management [22]. Furthermore, we believe there may be opportunities to consider peering relationships from the perspective of competitive games.

2.4 Caveats

Using an optimization approach we will generate a solution that is a function of the problem formulation (objective and constraints), the problem data (parameter values), and in cases where the problem cannot be solved exactly *the approximation technique itself.* There are many possible reasons why this approach could fail. For example, we may not be successful in capturing the dominant economic and technical forces driving topological growth. It is also possible that real decisions are neither consistent nor rational and thus do not correspond to any abstract mathematical formulation, although the effort will even then yield beneficial insights into what *ought* to be done. In particular, we expect this approach to shed light on the question of how important the careful incorporation of Level-2 technologies and economics is. Note that current router-level measurements are all IP-based and say little about the underlying link-layer technologies.

3. WILL IT WORK?

Despite the early and somewhat speculative nature of this work, it is supported by both theoretical and empirical evidence and would significantly enrich the current attempts of developing of a unified and integrated theory of the Internet.

3.1 Theoretical Support

A major theme in the physics literature for more than a decade has been the ubiquity of power law distributions in natural and artificial complex systems [5]. Engineered systems such as the Internet have recently been added to that list (e.g., see references in [32]). However, even though the orthodox physics view tends to associate power laws unambiguously with critical phase transitions [5], it is easy to refute this apparent connection—at least in the specific case of the Internet [32]. A radically different alternative view has recently been proposed by Carlson and Doyle [11, 12], relies on the concept of HOT (for *Highly Optimized Tolerance*), and has already been proven to be far more powerful and predictive than the orthodox theory [11, 12, 32].

Highly Optimized Tolerance (HOT): By emphasizing the importance of design, structure, and optimization, the HOT concept provides a framework in which the commonly-observed highly variable event sizes (often referred to as power law behavior) in systems optimized by engineering design are the results of tradeoffs between yield, cost of resources, and tolerance to risk.[3] *Tolerance* emphasizes that robustness (i.e., the maintenance of some desired system characteristics despite uncertainties in the behavior of its component parts or its environment) in complex systems is a constrained and limited quantity that must be diligently managed; *Highly Optimized* alludes to the fact that this goal is achieved by highly structured, rare, non-generic configurations which—for highly engineered systems—are the result of deliberate design. In turn, the characteristics of HOT systems are high performance, highly structured internal complexity, apparently simple and robust external behavior, with the risk of hopefully rare but potentially catastrophic cascading failures initiated by possibly quite small perturbations [12]. The challenge alluded to in Section 4 below then consists of applying HOT in the specific context of network topology by relying on the vast body of existing literature on network optimization.

Heuristically Optimized Tradeoffs: The first explicit attempt to cast topology design, modeling, and generation

---

[3] HOT thus suggests that these tradeoffs lead to highly optimized designs that perforce allow for a wide range of event sizes, in particular for occasional extreme sizes as a result of cascading failures.

as a HOT problem was by Fabrikant et al. [16] who suggest *heuristically optimized tradeoff* as an alternative (and we think attractive) acronym for HOT. They proposed a toy model of incremental access network design that optimizes a tradeoff between connectivity distance and node centrality. They showed that by changing the relative importance of these two factors to the overall objective function, the resulting topology can exhibit a range of hierarchical structures, from simple star-networks to trees. Furthermore, by tuning the relative importance of the two factors, the authors proved that the resulting node degree distributions can be either exponential or of the power-law type.

3.2 Measurement-Based Support

The work by Faloutsos et al. [17] was the first to report observing power laws for the node degree distributions of measured AS graphs of the Internet. This unexpected finding has stirred significant interest and has led to a drastic increase in the number of studies related to Internet topology modeling and generation (e.g., see [14] and references therein). As for router-level maps, attempts to infer Internet-wide maps have been reported in [25, 9, 15]. For our purpose, the more recent studies described in [19], and especially in [28], are of particular relevance as they focus on recovering the router-level maps of individual ISPs.

While some of these and related studies claim that both the inferred router- and AS-level graphs have similar structural properties, a more careful inspection of the underlying data and their analysis (see, in particular, [14] for the case of AS graphs and [28] for the case of ISP topologies) suggests that there may be indeed different mechanisms at work for generating them. While very different optimization problems can lead to very similar topologies, we believe that the optimization formulations —their objectives, constraints, and parameters—for generating the router-level graph and AS graph are very different. At a minimum, the identification of these formulations will be useful. A second possibility is that the router- and AS-level graphs are more different than similar, but that the standard metrics used to assess network similarity are not the right ones (e.g., see [30, 31]).[4]

3.3 A Piece of a Bigger Puzzle

The Internet serves as ideal starting point for a scientific exploration of the broader issues of robustness in complex systems, partularly those throughout engineering and biology. In most of these systems, complexity is driven by the need for robustness to uncertainty in system components or the operating environment far more than by the need for increased functionality. At the same time, most of this complexity tends to be hidden, deliberately creating the illusion of superficially simple systems and inviting the subsequent development of specious theories. However, motivated largely by the study of HOT systems, recent research efforts have provided for the first time a nascent but promising foundation for a rigorous, coherent, verifiable, and reasonably complete mathematical theory underpinning Internet technology (e.g., see [13] and references therein). This new theory emphasizes the importance of protocols that organize highly structured and complex modular hierarchies to achieve system robustenss, but also tend to create fragilities to rare, neglected, or unknown perturbations. It addresses directly the performance and robustness of both the "horizontal" decentralized and asynchronous nature of control in TCP/IP as well as the "vertical" separation into the layers of the TCP/IP protocol stack from the application layer down to the link layer. At present, the theory is mainly concerned with the transport layer, where the new findings generalize notions of source and channel coding from information theory as well as decentralized versions of robust control. The resulting new theoretical insights about the Internet also combine with our understanding of its origins and evolution to provide a rich source of ideas about complex systems in general.

The work proposed in this paper aims at extending this nascent mathematical theory of the Internet to layers below the transport layer by exploring the decentralized mechanisms and forces responsible for realizing a physical Internet infrastructure, which in turn provides a major ingredient for investigating both the "horizontal" (spatially distributed) and "vertical" aspects of IP routing. The ability to make detailed measurements at the relevant layers and an in-depth knowledge of how the individual parts work and are interconnected facilitate the validation of the proposed approach, make it possible to unambiguously diagnose and "reverse engineer" any claims or findings, and allow a clean separation between sound and specious theories.

4. NETWORK ACCESS DESIGN

There are a multitude of network design problems facing the modern ISP. While the length of this position paper prohibits a comprehensive review of these interesting and important problems, we have chosen to focus our initial attention on the problem of designing a distribution network that provides local access for its customers.[5] Typically, this design problem occurs at the level of the metropolitan area, and it is subject to the service demands of the individual customers and the technical constraints of the equipment in use. The purpose of selecting this problem as a starting point is to illustrate how simple, yet reasonable formulations can provide insight into resulting topology.

The network access design problem was originally studied in the context of planning local telecommunication access (see [6], [18], and references therein).[6] In general, these formulations incorporate the fixed costs of cable installation and the marginal costs of routing, as well as the cost of installing additional equipment, such as concentrators. An emphasis on cost in these formulations leads to solutions that are *tree* (or *forest*) topologies. However, other formulations are possible and can necessarily lead to vastly different topologies.[7]

---

[4] In fact, we expect that the right optimization formulation will help to identify the appropriate metrics and help to highlight the discrepancies if they exist.

[5] In this context, a customer is anyone who wants direct network access via a dedicated connection and is not using some other infrastructure (e.g. DSL over public telephone lines, cable networks) for this purpose.

[6] To the extent that the methods and intuition resulting from these studies were used in building real telecom networks, the access design problem is of even greater importance to our study, since many of the early Internet topologies piggybacked on these design principles and existing network structures.

[7] For example, adding a path redundancy requirement

4.1 Buy-At-Bulk Access Design

The network access design problem has received renewed attention within the last few years because of the *buy-at-bulk* nature of provisioning fiber-optic cables within ISPs [26, 2]. Buy-at-bulk means that when building its network, the ISP has for its installed links a choice of several different {capacity, cost} combinations, which we refer to as *cable types*. Specifically, each cable type $k \in \{1, 2, \ldots K\}$ has an associated capacity u_k, a fixed overhead (installation) cost σ_k, and a marginal usage cost δ_k. Collectively, the cable types exhibit *economies of scale* such that for $u_1 \leq u_2 \leq \cdots \leq u_K$, one has $\sigma_1 \leq \sigma_2 \leq \cdots \leq \sigma_K$ and $\delta_1 > \delta_2 > \cdots > \delta_K$. In words, larger capacity cables have higher overhead costs, but lower per-bandwidth usage costs.

The problem of interest is to construct a graph that connects some number of spatially distributed customers to a set of central (core) nodes, using a combination of cables that satisfies the traffic needs of the customers and incurs the lowest overall cost to the ISP. Details of the formulation can be found in [26, 2]. Since the decision for the optimal traffic routes and cables choices are interdependent, both problems must be solved simultaneously.

Despite the simple nature of this formulation, problems of this type are hard to solve to optimality. Indeed, constrained network access design problems belong within the family of *minimum cost spanning tree (MCST)* and *Steiner tree* problems, and the buy-at-bulk network design problem is known to be NP-Hard [26]. The best approximation algorithm known is the randomized algorithm by Meyerson et al. [24] who provide a constant factor bound on the quality of the solution independent of problem size.

4.2 Preliminary Results

In a preliminary investigation of the buy-at-bulk network access design problem, we have found that the approximation method in [24] yields tree topologies with exponential node degree distributions. These initial results were obtained using fictitious, yet realistic parameters[8] for cable capacities and costs. We believe that these results are consistent with those in [16], however a thorough search of the parameter space remains to be completed. In either case, we are already finding that the approach embodied by these methods is yielding valuable insights and exposing directions for new research.

5. RESEARCH AGENDA

While this work is still at its early stages, we have identified a number of areas that will require novel contributions within this broad research initiative.

What are the causal relationships between the objectives and constraints of a network design problem and the resulting topology? While there is deep understanding for how to solve combinatorial network optimization problems, the emphasis has traditionally been limited to the accuracy of the solutions and the computational complexity of achieving them. The scientific challenge here is to exploit the HOT perspective of network design to associate concrete optimization formulations with specific characteristic of the resulting network topology. While [16] is a promising first step in this direction, it has little to do with designing real networks.

What is the relative importance of such concrete formulations to real ISP topology design? This includes both backbone networks and distribution networks. We believe there are significant opportunities to learn from the best practices of seasoned network operators [20] as well as the potential for valuable contributions to the large-scale planning of Internet infrastructure. Gaining a basic understanding of this issue will require close interactions with network designers and operators and can benefit from exploiting a variety of different economic data relevant to the ISP and related market sectors.

What metrics and measurements will be required to validate or invalidate the resulting class of explanatory models? Diligent model validation (as for example outlined in [32]) will be an essential aspect of the proposed explanatory topology modeling work, and empirical studies such as those in [15, 28] are a necessary first step. However, by insisting on empirical verification of the causes underlying the advocated HOT-based approach to network topology against available (or yet to be measured) data, we enter an area where significant research efforts will be required for the development of novel and scientifically sound techniques and tools. Again, the result of Fabrikant et al. in [16] and their explicit multi-objective optimization formulation provide a concrete starting point for attempting to validate the forces at work in a proposed HOT-based topology models against feasible measurements.

Is is possible to accurately, yet anonymously characterize an ISP topology? Given that the current market environment for ISPs is only likely to become more competitive, we should consider how to devise technical solutions for the current barriers to information exchange.

What can economic theory tell us about the current and future interaction between competing ISPs? As the Internet becomes integrated as a critical infrastructure for our daily lives, it will be increasingly important that the environment for these companies is stable and that their behavior is predictable.

6. ACKNOWLEDGMENTS

This work was supported by the Institute for Pure and Applied Mathematics at UCLA as part of the 2002 Program on Large-Scale Communication Networks.

7. REFERENCES

[1] W. Aiello, F. Chung, and L. Lu. A random graph model for massive graphs. In *STOC 2000*, 2000.

[2] M. Andrews and L. Zhang. The access network design problem. *39th IEEE FOCS*, 1998.

[3] D. Awduche, J. Agogbua, J. McManus. An approach to optimal peering between autonomous systems in the Internet. *Proc. of Inter. Conf. on Comp. Commun. and Networks*, 1998.

---

breaks the tree structure of the optimal solution.

[8] Parameters were chosen to be consistent with the assumptions of the algorithm and the current marketplace.

[4] P. Baake and T. Wichmann. On the economics of Internet peering. *Netnomics*, 1(1), 1999.

[5] P. Bak. *How Nature Works: The Science of Self-Organized Criticality* Copernicus, New York, 1996.

[6] A. Balakrishnan, T.L. Magnanti, A. Shulman, and R.T. Wong. Models for planning capacity expansion in local access telecommunication networks. *Ann. of Oper. Res.*, 33, pp. 239–284, 1991.

[7] A.L. Barabasi and R. Albert. Emergence of Scaling in Random Networks. *Science*, 286:509–512, 1999.

[8] T. Bu and D. Towsley. On distinguishing between Internet power law topology generators. *Proc. IEEE INFOCOM'02*, 2002.

[9] H. Burch and B. Cheswick. Mapping the Internet. *IEEE Computer*, 32(4), 1999.

[10] K. Calvert, M. Doar, and E.W. Zegura. Modeling Internet topology. *IEEE Communications Magazine*, 35, pp. 160–163, 1997.

[11] J.M. Carlson and J.C. Doyle. Highly Optimized Tolerance: Robustness and design in complex systems. *Phys. Rev. Lett.*, 84, pp. 2529–2532, 2000.

[12] J.M. Carlson and J.C. Doyle. Complexity and Robustness. *Proc. Nat. Acad. Sci.*, 99, suppl. 1, pp. 2538–2545, 2002.

[13] J. Doyle, J. Carlson, S. Low, F. Paganini, G. Vinnicombe, W. Willinger, J. Hickey, P. Parillo, and L. Vandenberghe. Robustness and the Internet: Theoretical Foundations. http://netlab.caltech.edu/pub/papers/RIPartII.pdf

[14] H. Chang, Q. Chen, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. The origins of power laws in Internet topologies revisited. *Proc. IEEE INFOCOM'02*, 2002.

[15] k. claffy, T. Monk, and D. McRobb. Internet tomography. *Nature*, 1999.

[16] A. Fabrikant, E. Koutsoupias, and C. Papadimitriou. Heuristically Optimized Trade-offs: A new paradigm for power laws in the Internet. *ICALP 2002.* pp. 110–122. 2002.

[17] C. Faloutsos, P. Faloutsos, and M. Faloutsos. On power-law relationships of the Internet topology. *Proc. ACM SIGCOMM'99*, pp. 251–262, 1999.

[18] B. Gavish. Topological design of telecommunication networks—local access design methods. *Ann. of Oper. Res.*, 33, pp. 17–71, 1991.

[19] R. Govindan and H. Tangmunarunkit. Heuristics for Internet map discovery. *Proc. IEEE INFOCOM'00*, 2000.

[20] G. Huston. *ISP Survival Guide: Strategies for Running a Competitive ISP.* John Wiley & Sons, New York, 2000.

[21] C. Jin, Q. Chen, and S. Jamin. Inet: Internet topology generator. Tech. Rep. CSE-TR-433-00, Univ. of Michigan, Ann Arbor, 2000.

[22] R. Johari and J. Tsitsiklis. Routing and peering in a competitive Internet. Presentation at *IPAM Workshop on Large-Scale Comm. Networks: Topology, Routing, Traffic, and Control*, UCLA, 2002.

[23] A. Medina, A. Lakhina, I. Matta, and J. Byers. BRITE: An approach to universal topology generation. *Proc. MASCOTS'01*, 2001.

[24] A. Meyerson, K. Mungala, S. Plotkin. Designing networks incrementally. *41st IEEE FOCS*, 2001.

[25] J. Pansiot and D. Grad. On routes and multicast trees in the Internet. *ACM Comp. Comm. Rev.*, 1997.

[26] F.S. Salman, J. Cheriyan, R. Ravi, and S. Subramanian. Buy at bulk design: approximating the single-sink installation problem. In *Proc. 8th ACM-SIAM Symp. on Discrete Algorithms*, pp. 619–628, 1997. (Revised, 1998.)

[27] H. Simon. On a class of skew distribution functions. *Biometrika*, 42:425–440, 1955.

[28] N. Spring, R. Mahajan, and D. Weatherall. Measuring ISP topologies with rocketfuel. *Proc. ACM SIGCOMM'02* (to appear).

[29] H. Tangmunarunkit, J. Doyle, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. Does AS size determine degree in AS topology? *ACM Comp. Comm. Rev.*, 31, pp. 7–10, 2001.

[30] H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker, W. Willinger. Network topology generators: Degree-based vs. structural. *Proc. ACM SIGCOMM'02* (to appear).

[31] D. Vukadinović, P. Huang, and T. Erlebach. A spectral analysis of the Internet topology. ETH TIK-NR. 118, 2001.

[32] W. Willinger, R. Govindan, S. Jamin, V. Paxson, and S. Shenker. Scaling phenomena in the Internet: Critically examining criticality. *Proc. Nat. Acad. Sci.*, 99, suppl. 1, pp. 2573–2580, 2002.

[33] E. Zegura, K. Calvert, and M. Donahoo. A quantitative comparison of graph-based model for Internet topology. *ACM/IEEE Trans. on Networking*, 5(6):770–783, 1997.