A First-Principles Approach to Understanding the Internet's Router-level Topology

Lun Li
California Institute
of Technology
lun@cds.caltech.edu

David Alderson
California Institute
of Technology
alderd@cds.caltech.edu

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

John Doyle
California Institute
of Technology
doyle@cds.caltech.edu

ABSTRACT

A detailed understanding of the many facets of the Internet's topological structure is critical for evaluating the performance of networking protocols, for assessing the effectiveness of proposed techniques to protect the network from nefarious intrusions and attacks, or for developing improved designs for resource provisioning. Previous studies of topology have focused on interpreting measurements or on phenomenological descriptions and evaluation of graph-theoretic properties of topology generators. We propose a complementary approach of combining a more subtle use of statistics and graph theory with a first-principles theory of router-level topology that reflects practical constraints and tradeoffs. While there is an inevitable tradeoff between model complexity and fidelity, a challenge is to distill from the seemingly endless list of potentially relevant technological and economic issues the features that are most essential to a solid understanding of the intrinsic fundamentals of network topology. We claim that very simple models that incorporate hard technological constraints on router and link bandwidth and connectivity, together with abstract models of user demand and network performance, can successfully address this challenge and further resolve much of the confusion and controversy that has surrounded topology generation and evaluation.

Categories and Subject Descriptors

C.2.1 [**Communication Networks**]: Architecture and Design—*topology*

General Terms

Performance, Design, Economics

Keywords

Network topology, degree-based generators, topology metrics, heuristically optimal topology

1. INTRODUCTION

Recent attention on the large-scale topological structure of the Internet has been heavily focused on the *connectivity* of network

components, whether they be machines in the router-level graph [26, 10] or entire subnetworks (Autonomous Systems) in the AS-level graph [24, 14]. A particular feature of network connectivity that has generated considerable discussion is the prevalence of heavy-tailed distributions in node *degree* (e.g., number of connections) and whether or not these heavy-tailed distributions conform to power laws [23, 31, 16, 32]. This macroscopic statistic has greatly influenced the generation and evaluation of network topologies. In the current environment, degree distributions and other large-scale statistics are popular metrics for evaluating how representative a given topology is [42], and proposed topology generators are often evaluated on the basis of whether or not they can reproduce the same types of macroscopic statistics, especially power law-type degree distributions [11].

Yet, from our viewpoint, this perspective is both incomplete and in need for corrective action. For one, there exist many different graphs having the *same distribution of node degree*, some of which may be considered *opposites* from the viewpoint of network engineering. Furthermore, there are a variety of distinctly different random graph models that might give rise to a given degree distribution, and some of these models may have no network-intrinsic meaning whatsoever. Finally, we advocate here an approach that is primarily concerned with developing a basic understanding of the observed high variability in topology-related measurements and reconciling them with the reality of engineering design. From this perspective, reproducing abstract mathematical constructs such as power law distributions is largely a side issue.

In this paper, we consider *a first-principles approach* to understanding Internet topology at the *router-level*, where nodes represent routers and links indicate one-hop connectivity between routers. More specifically, when referring in the following to router-level connectivity, we always mean Layer 2, especially when the distinction between Layer 2 vs. Layer 3 issues is important for the purpose of illuminating the nature of the actual router-level connectivity (i.e., node degree) and its physical constraints. For router-level topology issues such as performance, reliability, and robustness to component loss, the physical connectivity between routers is more important than the virtual connectivity as defined by the higher layers of the protocol stack (e.g., IP, MPLS). Moreover, we use here the notion of "first-principles approach" to describe an attempt at identifying some *minimal* functional requirements and physical constraints needed to develop simple models of the Internet's router-level topology that are at the same time illustrative, representative, insightful, and consistent with engineering reality. Far from being exhaustive, this attempt is geared toward accounting for very basic network-specific aspects, but it can readily be enhanced if some new or less obvious functional requirements or physical constraints are found to play a critical role. Also, in the

process of developing models of the Internet router-level connectivity that are "as simple as possible, but not simpler", we focus on single ISPs or ASes as the Internet's fundamental building blocks that are designed largely in isolation and then connected according to both engineering and business considerations.

While there are several important factors that contribute to the design of an ISP's router-level topology (e.g., available technology, economic viability, customer demands, redundancy and geography) and while opinions will vary about which and how many of these factors matter, we focus here on a few critical technological and economic considerations that we claim provide insight into the types of network topologies that are possible. In essence, we argue the importance of explicit consideration of the basic tradeoffs that network designers must face when building real networks. In parallel, we provide evidence that network models of router-level connectivity whose construction is constrained by macroscopic statistics but is otherwise governed by randomness are inherently flawed. To this end, we introduce the notions of *network performance* and *network likelihood* as a new means for discerning important differences between generated and real network topologies. In so doing, we show that incorporating fundamental design details is crucial to the understanding and evaluation of Internet topology.

This paper is organized in the following manner. In Section 2, we review previous approaches to generating realistic topologies as well as some of the previous metrics directed at understanding and evaluating these topologies. In Section 3, we provide an alternate approach to understanding topology structure that explicitly incorporates link capacities, router technology constraints, and various economic constraints at work in the construction of real networks. Then in Section 4, we discuss several metrics (e.g., performance and likelihood) for comparing and contrasting networks, particularly networks having the same degree distribution. We present our findings in Section 5 and conclude in Section 6 by discussing implications and shortcomings of the proposed first-principles approach.

2. BACKGROUND AND RELATED WORK

Understanding the large-scale structural properties of the Internet has proved to be a challenging problem. Since the Internet is a collection of thousands of smaller networks, each under its own administrative control, there is no single place from which one can obtain a complete picture of its topology. Moreover, because the network does not lend itself naturally to direct inspection, the task of "discovering" the Internet's topology has been left to experimentalists who develop more or less sophisticated methods to infer this topology from appropriate network measurements. Because of the elaborate nature of the network protocol suite, there are a multitude of possible measurements that can be made, each having its own strengths, weaknesses, and idiosyncrasies, and each resulting in a distinct view of the network topology.

Two network topologies that have received significant attention from these experimental approaches are the *AS graph* (representing organizational interconnectivity between subnetworks) and the *router-level graph* of the Internet. Despite the challenges associated with the careful collection and interpretation of topology-related network measurements, significant efforts by the networking community are yielding an emerging picture of the large-scale statistical properties of these topologies [23, 26, 19, 10, 40, 41].

The development of abstract, yet informed, models for network topology evaluation and generation has followed the work of empiricists. The first popular topology generator to be used for networking simulation was the Waxman model [44], which is a variation of the classical Erdös-Rényi random graph [21]. The use of this type of random graph model was later abandoned in favor of models that explicitly introduce non-random structure, particularly hierarchy and locality, as part of the network design [20, 12, 48]. The argument for this type of approach was based on the fact that an inspection of real networks shows that they are clearly not random but do exhibit certain obvious hierarchical features. This approach further argued that a topology generator should reflect the design principles in common use. For example, in order to achieve desired performance objectives, the network must have certain connectivity and redundancy requirements, properties which are not guaranteed in random network topologies. These principles were integrated into the Georgia Tech Internetwork Topology Models (GT-ITM).

These *structural topology generators* were the standard models in use until power law relationships in the connectivity of both the AS-level and router-level graphs of the Internet were reported by Faloutsos et al. [23]. Since then, the identification and explanation of power laws has become an increasingly dominant theme in the recent body of network topology literature [47, 16, 31, 45]. Since the GT-ITM topology generators fail to produce power laws in node degree, they have often been abandoned in favor of new models that explicitly replicate these observed statistics.[1] Examples of these generators include the INET AS-level topology generator [28], BRITE [30], BA[47], AB [3], GLP[11], PLRG [2], and the CMU power-law generator [36].

Each of the aforementioned degree-based topology generators uses one of the following three probabilistic generation methods. The first is *preferential attachment* [7] which says (1) the growth of the network is realized by the sequential addition of new nodes, and (2) each newly added node connects to some existing nodes preferentially, such that it is more likely to connect with a node that already has many connections. As a consequence, high-degree nodes are likely to get more and more connections resulting in a power law in the distribution of node degree. For a precisely defined model that incorporates the key features of preferential attachment and is amenable to rigorous mathematical analysis, we refer to [8] and references therein. The second generation method is due to Chung and Lu [17] who considered a *general model of random graphs (GRG) with a given expected degree sequence*. The construction proceeds by first assigning each node its (expected) degree and then probabilistically inserting edges between the nodes according to a probability that is proportional to the product of the degrees of the two given endpoints. If the assigned expected node degree sequence follows a power-law, the generated graph's node degree distribution will exhibit the same power law. The third generation method, the *Power Law Random Graph (PLRG)* [2], also attempts to replicate a given (power law) degree sequence. This construction involves forming a set L of nodes containing as many distinct copies of a given vertex as the degree of that vertex, choosing a random matching of the the elements of L, and applying a mapping of a given matching into an appropriate (multi)graph[2].

One of the most important features of networks that have power law degree distributions and that are generated according to one of these probabilistic mechanisms is that they all tend to have a few centrally located and highly connected "hubs" through which essentially most traffic must flow. For the networks generated by preferential attachment, the central "hubs" tend to be nodes added early in the generation process. In the GRG model as well as in the PLRG model, the nodes with high (expected) degree have higher probability to attach to other high degree nodes and these highly

---

[1] See however a comment by E. Zegura on router-level topology modeling, http://www.caida.org/analysis/topology/router-level-topology.xml.

[2] It is believed that the PLRG and GRG models are "basically asymptotically equivalent, subject to bounding error estimates" [2].

connected nodes form a central cluster. When using these models to represent the Internet, the presence of these highly connected central nodes in these networks has been touted its "Achilles' heel" because network connectivity is highly vulnerable to attacks that target the high-degree hub nodes [4]. It has been similarly argued that these high-degree hubs are a primary reason for the epidemic spread of computer worms and viruses [37, 9]. The presence of highly connected central nodes in a network having a power law degree distribution is the essence of the so-called *scale-free* network models, which have been a popular theme in the study of complex networks, particularly among researchers inspired by statistical physics [34].

However, this emphasis on power laws and the resulting efforts to generate and explain them with the help of these degree-based methods have not gone without criticism. For example, there is a long-standing but little-known argument originally due to Mandelbrot which says in short[3] that power law type distributions should be expected to arise ubiquitously for purely mathematical and statistical reasons and hence require no special explanation. These distributions are a parsimonious null hypothesis for high variability data (i.e. when variance estimates fail to converge) just as Gaussians are for low variability data (i.e where moment estimates converge robustly and where mean and variance estimates tend to describe the measurements adequately) even though the data is not necessarily Gaussian (see [29, pp. 79–116] for details). Another more widely known deficiency is that degree-based methods for topology generation produce merely descriptive models that are in general not able to provide correct physical explanations for the overall network structure [45]. The claim is that, in the absence of an understanding of the drivers of network deployment and growth, it is difficult to identify the causal forces affecting large-scale network properties and even more difficult to predict future trends in network evolution. Nevertheless, in the absence of concrete examples of such alternate models, degree-based methods have remained popular representations for large-scale Internet structure.

This paper follows the previous arguments of [6] in favor of the need to explicitly consider the technical drivers of network deployment and growth. In spirit, it delivers for degree-based networks a similar message as [48] did for the random graph-type models [44] that were popular with networking researchers in the early 1990s. While [48] identified and commented on the inherent limitations of the various constructs involving Erdös-Rényi-type random graphs, our work points toward similar shortcomings and unrealistic features when working with probabilistic degree-based graphs.

3. A FIRST PRINCIPLES APPROACH

A key challenge in using large-scale statistical features to characterize something as complex as the topology of an ISP or the Internet as a whole is that it is difficult to understand the extent to which any particular observed feature is "fundamental" to its structure. Here, we consider a complementary approach for thinking about network topology, in which we explore some of the practical constraints and tradeoffs at work in the construction of real networks. In essence, we are asking the question, "What really matters when it comes to topology construction?" and argue that minimally one needs to consider the role of router technology and network economics in the network design process of a single ISP. The hope is that even a preliminary understanding of key factors, when combined with a more subtle use of statistics and graph theory, can provide a perspective that is more consistent both with observed measurements and the engineering principles at work in network

---

[3]We will elaborate on this argument in a forthcoming paper.
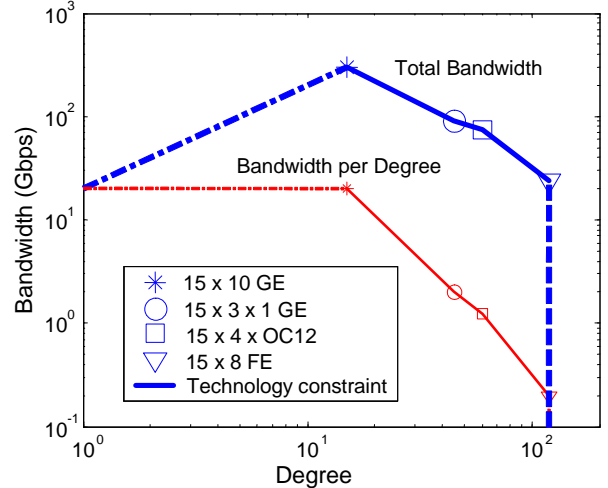

Figure 1: Technology constraint for Cisco 12416 Gigabit Switch Router(GSR): degree vs. bandwidth as of June 2002. Each point on the plot corresponds to a different combination of line cards and interfaces for the same router. This router has 15 available line card slots. When the router is configured to have less than 15 connections, throughput per degree is limited by the line-card maximum speed (10 Gbps) and the total bandwidth increases with the number of connections, while bandwidth per degree remains the same (dash-dot lines). When the number of connections is greater than 15, the total router bandwidth and bandwidth per degree decrease as the total number of connections increases (solid lines), up to a maximum of 120 possible connections for this router (dotted line). These three lines collectively define the feasible region for configuring this router.

design than with the current, at times conflicting, claims about the real Internet topology. In particular, given the current emphasis on the presence of power laws in the connectivity of the router-level Internet, it is important to understand whether such variability is plausible, and if so, where it might be found within the overall topology. Fortunately, such an explanation is possible if one considers the importance of router technology and network economics in the design process.

3.1 Technology Constraints

In considering the physical topology of the Internet, one observes that the underlying *router technology constraints* are a significant force shaping network connectivity. Based on the technology used in the cross-connection fabric of the router itself, a router has a maximum number of packets that can be processed in any unit of time. This constrains the number of link connections (i.e., node *degree*) and connection speeds (i.e., bandwidth) at each router. This limitation creates a "feasible region" and corresponding "efficient frontier" of possible bandwidth-degree combinations for each router. That is, a router can have a few high bandwidth connections or many low bandwidth connections (or some combination in between). In essence, this means that routers must obey a form of *flow conservation* in the traffic that they can handle. While it is always possible to configure the router so that it falls below the efficient frontier (thereby under-utilizing the router capacity), it is not possible to exceed this frontier (e.g., by having many high bandwidth connections). Figure 1 shows the technology constraint for the Cisco 12416 GSR, which is one of the most expensive and highest bandwidth routers available from a 2002 Cisco product catalog

Figure 2: Aggregate picture of router technology constraints. In addition to the Cisco 12000 GSR Series, the constraints on the somewhat older Cisco 7000 Series is also shown. The shared access technology for broadband cable provides service comparable to DSL when the total number of users is about 100, but can only provide service equivalent to dialup when the number of users is about 2000. Included also is the Linksys 4-port router, which is a popular LAN technology supporting up to 5 100MB Ethernet connections. Observe that the limits of this less expensive technology are well within the interior of the feasible region for core network routers.

[43]. Although engineers are constantly increasing the frontier with the development of new routing technologies, each particular router model will have a frontier representing its feasible region, and network architects are faced with tradeoffs between capacity and cost in selecting a router and then must also decide on the quantity and speed of connections in selecting a router configuration. Until new technology shifts the frontier, the only way to create throughput beyond the frontier is to build networks of routers.[4]

The current Internet is populated with many different router models, each using potentially different technologies and each having their own feasible region. However, these technologies are still constrained in their overall ability to tradeoff total bandwidth and number of connections. Consider an aggregate picture of many different technologies (shown in Figure 2), used both in the network core and at the network edge. Edge technologies are somewhat different in their underlying design, since their intention is to be able to support large numbers of end users at fixed (DSL, dialup) or variable (cable) speeds. They can support a much greater number of connections (upwards of 10,000 for DSL or dialup) but at significantly lower speeds. Collectively, these individual constraints form an overall aggregate constraint on available topology design.

We are not arguing that limits in technology fundamentally preclude the possibility of high-degree, high-bandwidth routers, but simply that the product offerings recently available to the marketplace have not supported such configurations. While we expect that companies will continue to innovate and extend the feasible region for router configuration, it remains to be seen whether or not the economics (including configuration and management) for these products will enable their wide deployment within the Internet.

3.2 Economic Considerations

Even more important than the technical considerations affecting

---

[4]Recent product announcements from router manufacturers such as Juniper Networks, Avici Systems, and Cisco Systems suggest that the latest trend in technology development is to build scaleable multi-rack routers that do exactly this.
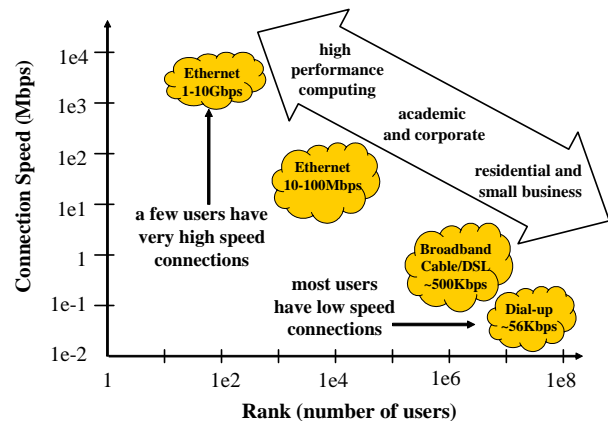

Figure 3: Aggregate picture of end user connection bandwidths for the Internet. Most users of the current Internet have relatively slow (56Kbps) connections, while only a relative few have high speed (10Gbps) connections.

router use are the economic considerations of network design and deployment, which are driven by customer demands and ultimately direct the types of technologies that are developed for use by network providers. For example, the cost of installing and operating physical links in a network can dominate the cost of the overall infrastructure, and since these costs tend to increase with link distance, there is tremendous practical incentive to design wired networks such that they can support traffic using the fewest number of links. The ability to share costs via multiplexing is a fundamental driver underlying the design of networking technologies, and the availability of these technologies enables a network topology in which traffic is aggregated at all levels of network hierarchy, from its periphery all the way to its core.

The development of these technologies has similarly followed the demands of customers, for whom there is wide variability in the willingness to pay for network bandwidths (Figure 3). For ex-
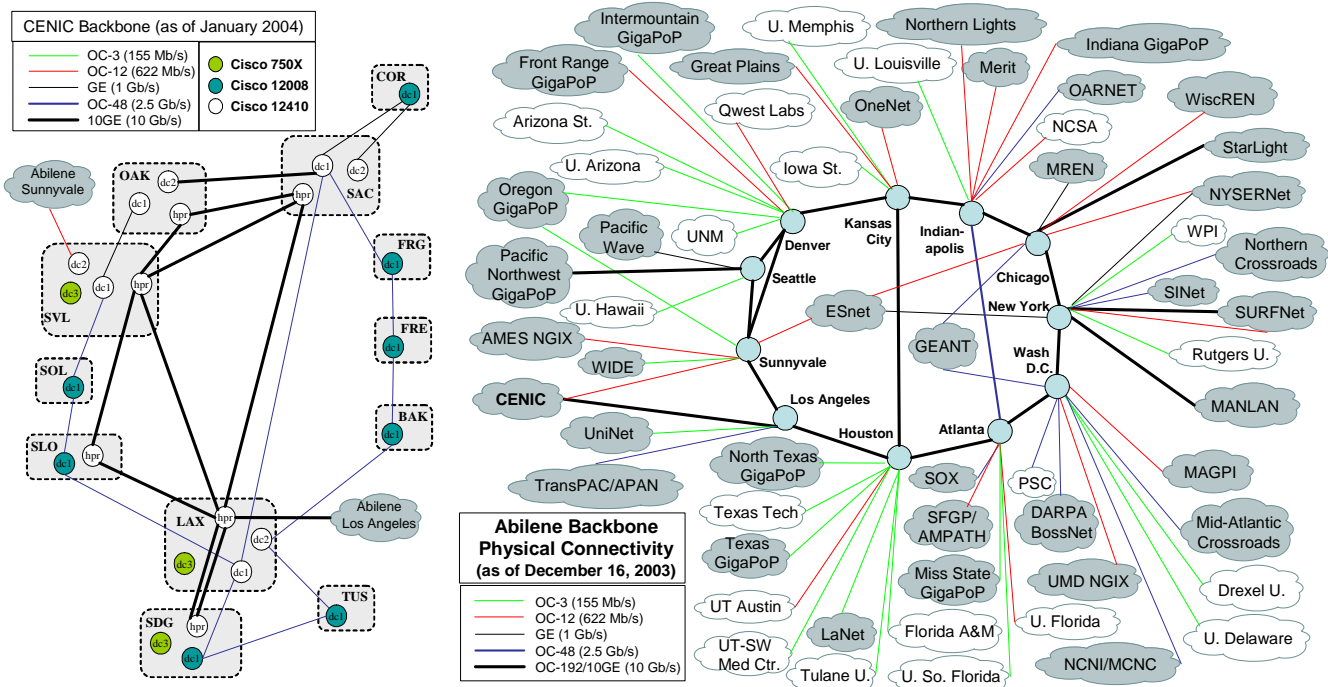
Figure 4: CENIC and Abilene networks. (Left): CENIC backbone. The CENIC backbone is comprised of two backbone networks in parallel—a high performance (HPR) network supporting the University of California system and other universities, and the digital California (DC) network supporting K-12 educational initiatives and local governments. Connectivity within each POP is provided by Layer-2 technologies, and connectivity to the network edge is not shown. (Right): Abilene network. Each node represents a router, and each link represents a physical connection between Abilene and another network. End user networks are represented in white, while peer networks (other backbones and exchange points) are represented in gray. Each router has only a few high bandwidth connections, however each physical connection can support many virtual connections that give the appearance of greater connectivity to higher levels of the Internet protocol stack. ESnet and GEANT are other backbone networks.

ample, nearly half of all users of the Internet in North America still have dial-up connections (generally 56kbps), only about 20% have broadband access (256kbps-6Mbps), and there is only a small number of users with large (10Gbps) bandwidth requirements [5]. Again, the cost effective handling of such diverse end user traffic requires that aggregation take place as close to the edge as possible and is explicitly supported by a common feature that these edge technologies have, namely a special ability to support high connectivity in order to aggregate end user traffic before sending it towards the core. Based on variability in population density, it is not only plausible but somewhat expected that there exist a wide variability in the network node connectivity.

Thus, a closer look at the technological and economic design issues in the network core and at the network edge provides a consistent story with regard to the forces (e.g., market demands, link costs, and equipment constraints) that appear to govern the build-out and provisioning of the ISPs' core networks. The tradeoffs that an ISP has to make between what is technologically feasible versus economically sensible can be expected to yield router-level connectivity maps where individual link capacities tend to increase while the degree of connectivity tends to decrease as one moves from the network edge to its core. To a first approximation, core routers tend to be fast (have high capacity), but have only a few high-speed connections; and edge routers are typically slower overall, but have many low-speed connections. Put differently, long-haul links within the core tend to be relatively few in numbers but their capacity is typically high.

3.3 Heuristically Optimal Networks

The simple technological and economic considerations listed above suggest that a reasonably "good" design for a single ISP's network is one in which the core is constructed as a loose mesh of high speed, low connectivity routers which carry heavily aggregated traffic over high bandwidth links. Accordingly, this mesh-like core is supported by a hierarchical tree-like structure at the edges whose purpose is to aggregate traffic through high connectivity. We will refer to this design as *heuristically optimal* to reflect its consistency with real design considerations.

As evidence that this heuristic design shares similar qualitative features with the real Internet, we consider the real router-level connectivity of the Internet as it exists for the educational networks of Abilene and CENIC (Figure 4). The Abilene Network is the Internet backbone network for higher education, and it is part of the Internet2 initiative [1]. It is comprised of high-speed connections between core routers located in 11 U.S. cities and carries approximately 1% of all traffic in North America[5]. The Abilene backbone is a sparsely connected mesh, with connectivity to regional and local customers provided by some minimal amount of redundancy. Abilene is built using Juniper T640 routers, which are configured to have anywhere from five connections (in Los Angeles) to twelve

---

[5]Of the approximate 80,000 - 140,000 terabytes per month of traffic in 2002 [35], Abilene carried approximately 11,000 terabytes of total traffic for the year [27]. Here, "carried" traffic refers to traffic that traversed an Abilene router. Since Abilene does not peer with commercial ISPs, packets that traverse an Abilene router are unlikely to have traversed any portion of the commercial Internet.
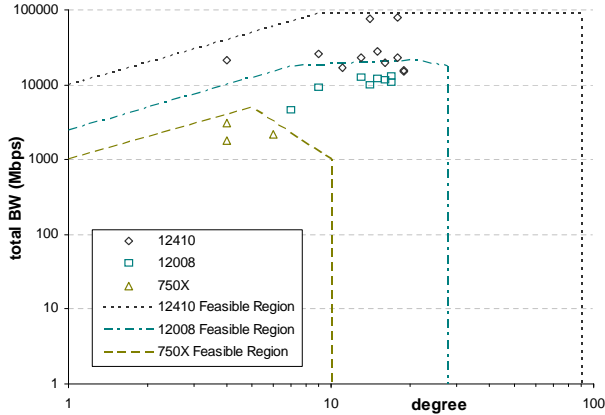
Figure 5: Configuration of CENIC routers as of January 2004. In the time since the Cisco catalog [43] was published, the introduction of a new line card (supporting 10x1GE interfaces) has shifted the feasible region for the model 12410 router. Since this router has nine available slots, this router can achieve a maximum of 90 Gbps with either nine 10GE line cards or nine 10x1GE line cards. Although the shape of the feasible region may continue to change, its presence and corresponding implications for router configuration and deployment will remain qualitatively the same.

connections (in New York). Abilene maintains peering connections with other higher educational networks (both domestic and international) but does not connect directly to the commercial Internet.

Focusing in on a regional level, we consider California, where the Corporation for Education Network Initiatives in California (CENIC) acts as ISP for the state's colleges and universities [18]. Its backbone is similarly comprised of a sparse mesh of routers connected by high speed links (Figure 4). Here, routing policies, redundant physical links, and the use of virtual private networks support robust delivery of traffic to edge campus networks. Similar observations are found when examining (where available) topology-related information of global, national, or regional commercial ISPs.

In view of recent measurement studies [26, 19, 40], it is important to recognize that the use of technologies at layers other than IP will affect what traceroute-like experiments can measure. For example, the use of shared media at Layer 2 (e.g. Ethernet, FDDI rings) either at the network edge or at exchange points between ISPs can give the appearance of high degree nodes. In an entirely different fashion, the use of Multiprotocol Label Switching (MPLS) at higher levels of the protocol stack can also give the illusion of one-hop connectivity at the lower layers when, in fact, there is none. Abilene is an ideal starting point for understanding heuristically optimal topologies, because within its backbone, there is no difference between the link layer topology and what is seen by IP. In contrast, the use of Ethernet and other link layer switching technologies within the CENIC POPs makes the interpretation and visualization of the physical intra-CENIC connectivity more difficult, but inferring the actual link layer connectivity is greatly facilitated by knowing the configurations of the individual CENIC routers as shown in Figure 5. The extent to which high degree nodes observed in traceroute-like studies is due to effects at POPs or Internet Exchange Providers (IXPs), as opposed to edge-aggregation effects, is not clear from current measurement studies.

We also recall that the emphasis in this paper is on a reasonable network design at the level of a single ISP. However, we recognize that the broader Internet is a collection of thousands of ASes

that interconnect at select locations. Thus, an important issue that is not addressed in this paper is understanding how the large-scale structure of the Internet relates to the heuristically optimal network design of single ISPs. We speculate that similar technology constraints and economic drivers will exist at peering points between ISPs, but that the complexity of routing management may emerge as an additional consideration. As a result, we fully expect border routers to again have a few relatively high bandwidth physical connections supporting large amounts of aggregated traffic. In turn, high physical connectivity at the router level is expected to be firmly confined to the network edge.

4. TOPOLOGY METRICS

4.1 Commonly-used Metrics

Previous metrics to understanding and evaluating network topologies have been dominated by graph-theoretic quantities and their statistical properties, e.g., node-degree distribution, expansion, resilience, distortion and hierarchy [11, 42]. However we claim here that these metrics are inherently inadequate to capture the essential tradeoffs of explicitly engineered networks.

Node degree distribution. In general, there are many networks having the same node degree distribution, as evidenced by the process of *degree-preserving rewiring*. This particular rewiring operation rearranges existing connections in such a way that the degrees of the nodes involved in the rearrangement do not change, leaving the resulting overall node degree distribution invariant. Accordingly, since the network can be rewired step-by-step so that the high degree nodes appear either at the network core or at its edges, it is clear that radically different topologies can have one and the same degree distribution (e.g., power law degree distribution). In this fashion, degree-preserving rewiring is a means for moving within a general "space of network graphs," all having the same overall degree distribution.

Expansion, Resilience, Distortion. Introduced in [42], these metrics are intended to differentiate important aspects of topology. *Expansion* is intended to measure the ability of a node to "reach" other nodes within a given distance (measured by hops), *resilience* is intended to reflect the existence of alternate paths, and *distortion* is a graph theoretic metric that reflects the manner in which a spanning tree can be embedded into the topology. For each of these three metrics, a topology is characterized as being either "Low" (L) or "High" (H). Yet, the quantitative values of expansion, resilience, and distortion as presented in [42] are not always easy to interpret when comparing qualitatively different topologies. For example, the measured values of expansion for the AS-level and router-level topologies show a relatively big difference (Figure 2(d) in [42]), however both of them are classified as "High", suggesting that the degree-based generators compare favorably with measured topologies. In contrast, it could be argued that Tiers generates topologies whose expansion values match that of the measured router-level graph reasonably well (Figure 2(g) in [42]), but Tiers is classified to have "Low" expansion. Such problems when interpreting these metrics make it difficult to use them for evaluating differences in topologies in a consistent and coherent manner.

Nonetheless, these metrics have been used in [42] to compare measured topologies at the autonomous system (AS) level and the router level (RL) to topologies resulting from several generators, including degree-based methods (PLRG, BA, BRITE, BT, INET) and structural methods (GT-ITM's Tiers and Transit-Stub), as well as several "canonical" topologies (e.g., random, mesh, tree, complete graph). It was observed that AS, RL, and degree-based networks were the only considered networks that share values "HHL" for ex-

pansion, resilience, and distortion respectively. Furthermore, of the canonical topologies, this "HHL" characterization was shared only by the complete graph (all nodes connected to each other). However, one canonical topology that was not considered was the "star" topology (i.e., having a single central hub), which according to their metrics would also be characterized as "HHL", and which explains why the degree-based graphs (having high degree central hubs) fit this description. Yet, the fact that both a complete graph and a star could have the same characterization illustrates how this group of metrics is incomplete in evaluating network topology.

Hierarchy. For evaluating hierarchy, [42] considers the distribution of "link values", which are intended to mimic the extent to which network traffic is aggregated on a few links (presumably, backbone links). However, the claim that degree-based generators, such as PLRG, do a better a job of matching the observed hierarchical features of measured topologies is again based on a qualitative assessment whereby previous structural generators (e.g., Tiers in GT-ITM) create hierarchy that is "strict" while degree-based generators result, like measured topologies, in hierarchies that are "moderate". This assessment is based on a model in which end-to-end traffic follows shortest path routes, however it also ignores any constraints on the ability of the network to simultaneously carry that end-to-end traffic.

From the perspective of this paper, these previous metrics appear to be inadequate for capturing what matters for real network topologies. Many of them lack a direct networking interpretation, and they all rely largely on qualitative criteria, making their application somewhat subjective. In what follows, we use the experience gained by these previous studies to develop metrics that are consistent with our first principles perspective. In particular, we consider several novel measures for comparing topologies that we show provide a minimal, yet striking comparison between degree-based probabilistic networks and networks inspired by engineering design.

4.2 Performance-Related Metrics

Recognizing that the primary purpose for building a network is to carry effectively a projected overall traffic demand, we consider several means for evaluating the performance of the network.

Throughput. We define *network performance* as the maximum throughput on the network under heavy traffic conditions based on a gravity model [38]. That is, we consider flows on all source-destination pairs of edge routers, such that the amount of flow X_{ij} between source i and destination j is proportional to the product of the traffic demand x_i, x_j at end points i, j, $X_{ij} = \alpha x_i x_j$, where α is some constant. We compute the maximum throughput on the network under the router degree bandwidth constraint,

$$\max_{\alpha} \quad \sum_{ij} \alpha x_i x_j$$
$$s.t \quad RX \leq B,$$

where X is a vector obtained by stacking all the flows $X_{ij} = \alpha x_i x_j$ and R is the routing matrix (defined such that $R_{kl} = \{0, 1\}$ depending on whether or not flow l passes through router k). We use shortest path routing to get the routing matrix, and define B as the vector consisting of all router bandwidths according to the degree bandwidth constraint (Figure 2). Due to a lack of publicly available information on traffic demand for each end point, we assume the bandwidth demand at a router is proportional to the aggregated demand of any end hosts connected to it. This assumption allows for good bandwidth utilization of higher level routers[6]. While

---

[6]We also tried choosing the traffic demand between routers as the

other performance metrics may be worth considering, we claim that maximum throughput achieved using the gravity model provides a reasonable measure of the network to provide a *fair* allocation of bandwidth.

Router Utilization. In computing the maximum throughput of the network, we also obtain the total traffic flow through each router, which we term *router utilization*. Since routers are constrained by the feasible region for bandwidth and degree, the topology of the network and the set of maximum flows will uniquely locate each router within the feasible region. Routers located near the frontier are used more efficiently, and a router on the frontier is saturated by the traffic passing through it. For real ISPs, the objective is clearly not to maximize throughput but to provide some service level guarantees (e.g. reliability), and modeling typical traffic patterns would require additional considerations (such as network overprovisioning) that are not addressed here. Our intent is not to reproduce real traffic, but to evaluate the raw carrying capacity of selected topologies under reasonable traffic patterns and technology constraints.

End User Bandwidth Distribution. In addition to the router utilization, each set of maximum flows also results in a set of bandwidths that are delivered to the end users of the network. While not a strict measure of performance, we consider as a secondary measure the ability of a network to support "realistic" end user demands.

4.3 Likelihood-Related Metric

To differentiate between graphs g having the same vertex set V and the same degree distribution or, equivalently, the same node degree sequence $\omega = (\omega_1, \cdots \omega_n)$, where ω_k denotes the degree of node k, consider the metric

$$L(g) = \sum_{(i,j) \in E(g)} \omega_i \omega_j, \tag{1}$$

where $E(g)$ represents the set of edges (with $(i, j) \in E(g)$ if there is an edge between vertices i and j). This (deterministic) metric is closely related to previous work on assortativity [33], however for the purpose of this paper, we require a renormalization that is appropriate to the study of all simple connected graphs with vertex set V and having the same node degree sequence ω. To this end, we define the normalized metric

$$l(g) = (L(g) - L_{\min})/(L_{\max} - L_{\min}), \tag{2}$$

where L_{max} and L_{min} are the maximum and minimum values of $L(g)$ among all simple connected graphs g with vertex set V and one and the same node degree sequence ω. Note that, for example, the L_{max} graph is easily generated according to the following heuristic: sort nodes from highest degree to lowest degree, and connect the highest degree node successively to other high degree nodes in decreasing order until it satisfies its degree requirement. By performing this process repeatedly to nodes in descending degree, one obtains a graph with vertex set V that has the largest possible $L(g)$-value among all graphs with node degree sequence ω. A formal proof that this intuitive construction yields an L_{max} graph employs the Rearrangement Inequality [46]. It follows that graphs g with high $L(g)$-values are those with high-degree nodes connected to other high-degree nodes and low-degree nodes connected to low-degree nodes. Conversely, graphs g with high-degree nodes connected to low-degree nodes have necessarily lower $L(g)$-values. Thus, there is an explicit relationship between graphs with high $L(g)$-values and graphs having a "scale-free" topology in the

---

product of their degrees as in [25], and qualitatively similar performance values are obtained but with different router utilization.

sense of exhibiting a "hub-like" core; that is, high connectivity nodes form a cluster in the center of the network.

The $L(g)$ and $l(g)$ metrics also allow for a more traditional interpretation as likelihood and relative likelihood, respectively, associated with the *general model of random graphs (GRG) with a given expected degree sequence* considered in [17]. The GRG model is concerned with random graphs with given expected node degree sequence $\omega = (\omega_1, \cdots \omega_n)$ for vertices $1, \cdots, n$. The edge between vertices i and j is chosen independently with probability p_{ij}, with p_{ij} proportional to the product $\omega_i \omega_j$. This construction is general in that it can generate graphs with a power law node degree distribution if the given expected degree sequence ω conforms to a power law, or it can generate the classic Erdös-Rényi random graphs [21] by taking the expected degree sequence ω to be $(pn, pn, \cdots, pn)$. As a result of choosing each edge $(i, j) \in E(g)$ with a probability that is proportional to $\omega_i \omega_j$, in the GRG model, different graphs are assigned different probabilities. In fact, denoting by $G(\omega)$ the set of all graphs generated by the GRG method with given expected degree sequence ω, and defining the *likelihood* of a graph $g \in G(\omega)$ as the logarithm of the probability of that graph, conditioned on the actual degree sequence being equal to the expected degree sequence ω, the latter can be shown to be proportional to $L(g)$, which in turn justifies our interpretation below of the $l(g)$ metric as *relative likelihood* of $g \in G(\omega)$. However, for the purpose of this paper, we simply use the $l(g)$ metric to differentiate between networks having one and the same degree distribution, and a detailed account of how this metric relates to notions such as graph self-similarity, likelihood, assortativity, and "scale-free" will appear elsewhere.

5. COMPARING TOPOLOGIES

In this section, we compare and contrast the features of several different network graphs using the metrics described previously. Our purpose is to show that networks having the same (power-law) node degree distribution can (1) have vastly different features, and (2) appear deceivingly similar from a view that considers only graph theoretic properties.

5.1 A First Example

Our first comparison is made between five networks resulting from preferential attachment (PA), the GRG method with given expected node degree sequence, a generic heuristic optimal design, an Abilene-inspired heuristic design, and a heuristic sub-optimal design. In all cases, the networks presented *have the same power-law degree distribution*. While some of the methods do not allow for direct construction of a selected degree distribution, we are able to use degree preserving rewiring as an effective (if somewhat artificial) method for obtaining the given topology. In particular, we generate the PA network first, then rearrange routers and links to get heuristically designed networks while keeping the same degree distribution. Lastly, we generate an additional topology according to the GRG method. What is more important here are the topologies and their different features, not the process or the particular algorithm that generated them.

Preferential Attachment (PA). The PA network is generated by following process: begin with 3 fully connected nodes, then in successive steps add one new node to the graph, such that this new node is connected to the existing nodes with probability proportional to the current node degree. Eventually we generate a network with 1000 nodes and 1000 links. Notice that this initial structure is essentially a tree. We augment this tree by successively adding additional links according to [3]. That is, in each step, we choose a node randomly and connect it to the other nodes with probability proportional to the current node degree. The resulting PA topol-

ogy is shown in in Figure 6(b) and has an approximate power-law degree distribution shown in Figure 6(a).

General Random Graph (GRG) method. We use the degree sequence of the PA network as the expected degree to generate another topology using the GRG method. Notice that this topology generator is not guaranteed to yield a connected graph, so we pick the giant component of the resulting structure and ignore the self-loops as in [42]. To ensure the proper degree distribution, we then add degree one edge routers to this giant component. Since the total number of links in the giant component is generally greater than the number of links in an equivalent PA graph having the same number of nodes, the number of the edge routers we can add is smaller than in the original graph. The resulting topology is shown in Figure 6(c), and while difficult to visualize all network details, a key feature to observe is the presence of highly connected central nodes.

Heuristically Optimal Topology (HOT). We obtain our HOT graph using a heuristic, nonrandom, degree-preserving rewiring of the links and routers in the PA graph. We choose 50 of the lower-degree nodes at the center to serve as core routers, and also choose the other higher-degree nodes hanging from each core as gateway routers. We adjust the connections among gateway routers such that their aggregate bandwidth to a core node is almost equally distributed. The number of edge routers placed at the edge of the network follows according to the degree of each gateway. The resulting topology is shown in Figure 6(d). In this model, there are three levels of router hierarchy, each of which loosely correspond (starting at the center of the network and moving out toward the edges) to backbone, regional/local gateways, edge routers. Of course, several other "designs" are possible with different features. For example, we could have rearranged the network so as to have a different number of "core routers", provided that we maintained our heuristic approach in using low-degree (and high bandwidth) routers in building the network core.

Abilene-inspired Topology. We claim that the backbone design of Abilene is heuristically optimal. To illustrate this, we construct a simplified version of Abilene in which we replace each of the edge network clouds in Figure 4 with a single gateway router supporting a number of end hosts. We assign end hosts to gateway routers in a manner that yields the same approximate power law in overall node degree distribution. The resulting topology with this node degree distribution is illustrated in Figure 6(d).

Sub-optimal Topology. For the purposes of comparison, we include a heuristically designed network that has not been optimized for performance (Figure 6(f)). This network has a chain-like core of routers, yet again has the same overall degree distribution.

Performance. For each of these networks, we impose the same router technological constraint on the non-edge routers. In particular, and to accomodate these simple networks, we use a fictitious router based on the Cisco GSR 12410, but modified so that the maximum number of ports it can handle coincides with the maximum degree generated above (see the dot-line in Figure 7(b-f)). Thus, each of these networks has the same number of non-edge nodes and links, as well as the same degree distribution among non-edge nodes. Collectively, these assumptions guarantee the same total "cost" (measured in routers) for each network. Using the performance index defined in Section 4.2, we compute the performance of these five networks. Among the heuristically designed networks, the HOT model achieves 1130 Gbps and the Abilene-inpsired network achieves 395 Gbps, while the sub-optimal network achieves only 18.6 Gbps. For the randomly generated graphs, the PA and GRG achieve only 11.9 Gbps and 16.4 Gbps respectively, roughly 100 times worse than the HOT network. The main reason for PA and GRG models to have such terrible performance
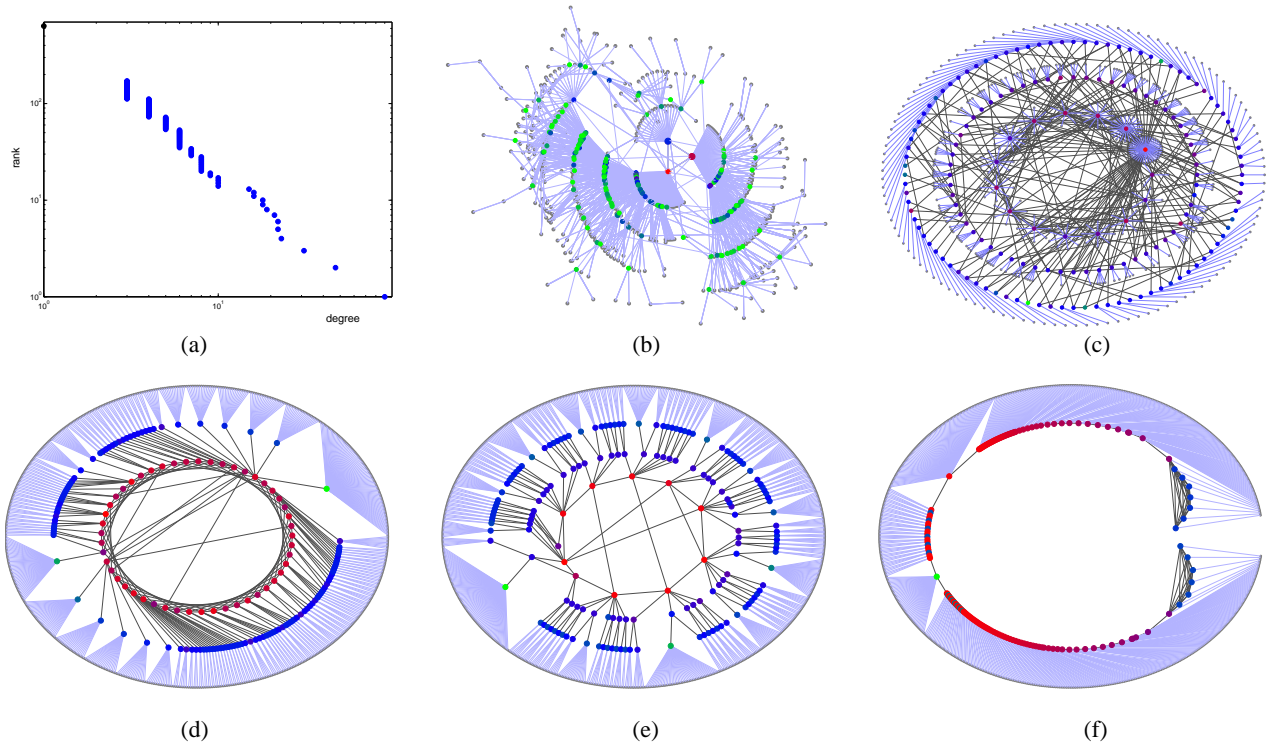
Figure 6: Five networks having the same node degree distribution. (a) Common node degree distribution (degree versus rank on log-log scale); (b) Network resulting from preferential attachment; (c) Network resulting from the GRG method; (d) Heuristically optimal topology; (e) Abilene-inspired topology; (f) Sub-optimally designed topology.
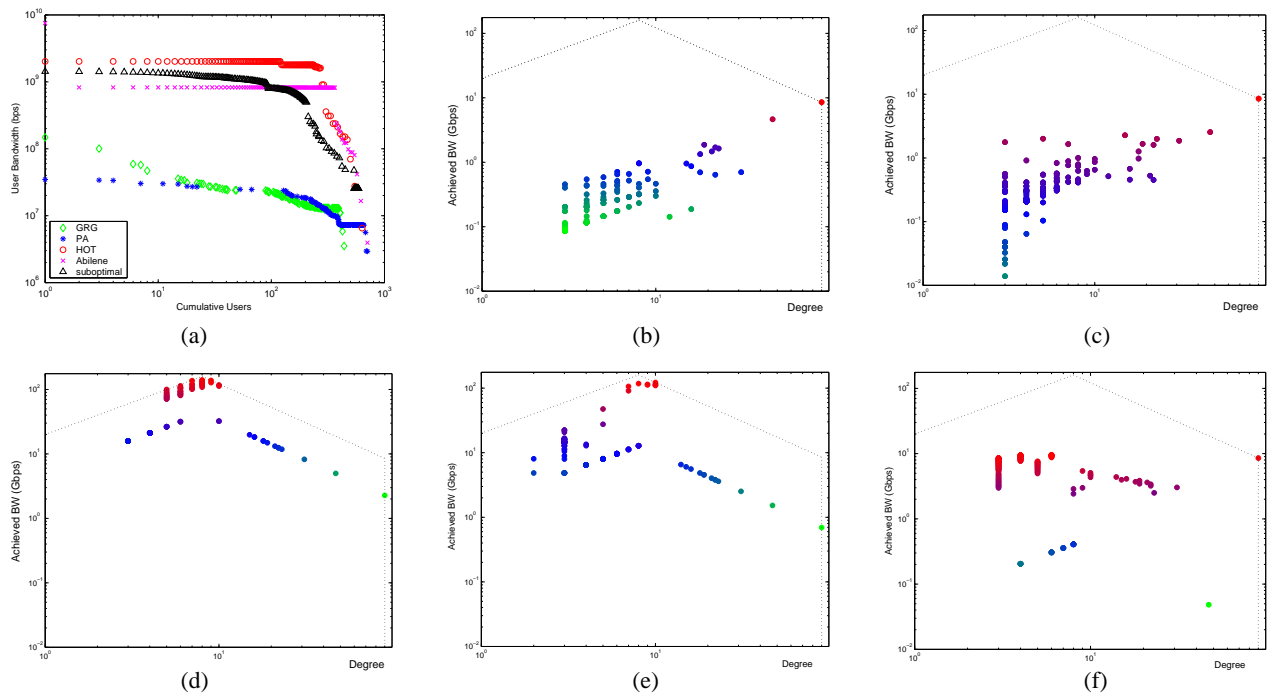

Figure 7: (a) Distribution of end user bandwidths; (b) Router utilization for PA network; (c) Router utilization for GRG network; (d) Router utilization for HOT topology; (e) Router utilization for Abilene-inspired topology; (f) Router utilization for sub-optimal network design. The colorscale of a router on each plot differentiates its bandwidth which is consistent with the routers in Figure 6.
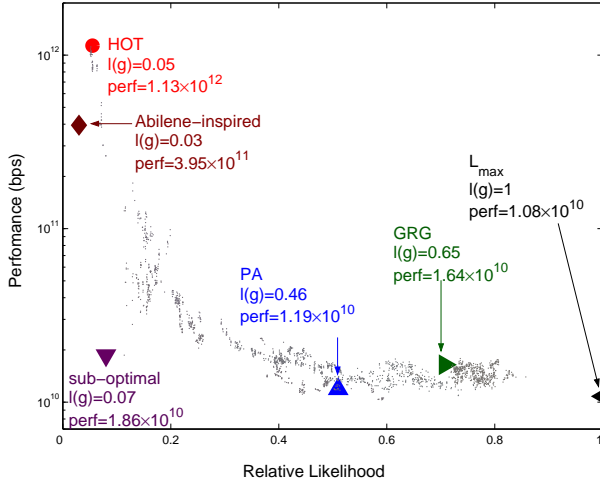
Figure 8: Performance vs. Likelihood for each topology, plus other networks having the same node degree distribution obtained by pairwise random rewiring of links.

is exactly the presence of the highly connected "hubs" that create low-bandwidth bottlenecks. The HOT model's mesh-like core, like the real Internet, aggregates traffic and disperses it across multiple high-bandwidth routers. We calculate the distribution of end user bandwidths and router utilization when each network achieves its best performance. Figure 7 (a) shows that the HOT network can support users with a wide range of bandwidth requirements, however the PA and GRG models cannot. Figure 7(d) shows that routers achieve high utilization in the HOT network, whereas, when the high degree "hubs" saturate in the PA and GRG networks, all the other routers are left under-utilized (Figure 7(b)(c)). The networks generated by these two degree-based probabilistic methods are essentially the same in terms of their performance.

Performance vs. Likelihood. A striking contrast is observed by simultaneously plotting performance versus likelihood for all five models in Figure 8. The HOT network has high performance and low likelihood while the PA and GRG networks have high likelihood but low performance. The interpretation of this picture is that a careful design process explicitly incorporating technological constraints can yield high-performance topologies, but these are extremely rare from a probabilistic graph point of view. In contrast, equivalent power-law degree distribution networks constructed by generic degree-based probabilistic constructions result in more likely, but poor-performing topologies. The "most likely" L_{max} network (also plotted in Figure 8) has poor performance.

This viewpoint is augmented if one considers the process of pairwise random degree-preserving rewiring as a means to explore the space of graphs having the same overall degree distribution. In Figure 8, each point represents a different network obtained by random rewiring. Despite the fact that all of these graphs have the same overall degree distribution, we observe that a large number of these networks have relatively high likelihood and low performance. All of these graphs, including the PA and GRG networks, are consistent with the so-called "scale-free" models in the sense that they contain highly connected central hubs. The fact that there are very few high performance graphs in this space is an indication that it would be "hard" to find a relatively good design using random rewiring. We also notice that low likelihood itself does not guarantee a high performance network, as the network in Figure 6(f) shows that it is possible to identify probabilistically rare and poorly performing networks. However, based on current evidence, it does appear to be

the case that it is impossible using existing technology to construct a network that is both high performance and high likelihood.

5.2 A Second Example

Figure 6 shows that graphs having the same node degree distribution can be very different in their structure, particularly when it comes to the engineering details. What is also true is that the same core network design can support many different end-user bandwidth distributions and that by and large, the variability in end-user bandwidth demands determines the variability of the node degrees in the resulting network. To illustrate, consider the simple example presented in Figure 9, where the same network core supports different types of variability in end user bandwidths at the edge (and thus yields different overall node degree distributions). The network in Figure 9(a) provides uniformly high bandwidth to end users; the network in Figure 9(b) supports end user bandwidth demands that are highly variable; and the network in Figure 9(c) provides uniformly low bandwidth to end users. Thus, from an engineering perspective, not only is there not necessarily any implied relationship between a network degree distribution and its core structure, there is also no implied relationship between a network's core structure and its overall degree distribution.

6. DISCUSSION

The examples discussed in this paper provide new insight into the space of all possible graphs that are of a certain size and are constrained by common macroscopic statistics, such as a given (power law) node degree distribution. On the one hand, when viewed in terms of the (relative) likelihood metric, we observe a dense region that avoids the extreme ends of the likelihood axis and is populated by graphs resulting from random generation processes, such as PA and GRG. Although it is possible to point out details that are specific to each of these "generic" or "likely" configurations, when viewed under the lens provided by the majority of the currently considered macroscopic statistics, they all look very similar and are difficult to discern. Their network cores contain high connectivity hubs that provide a relatively easy way to generate the desired power law degree distribution. Given this insight, it is not surprising that theorists who consider probabilistic methods to generate graphs with power-law node degree distributions and rely on statistical descriptions of global graph properties "discover" structures that are hallmarks of the degree-based models.

However, the story changes drastically when we consider network performance as a second dimension and represent the graphs as points in the likelihood-performance plane. The "generic" or "likely" graphs that make up much of the total configuration space have such bad performance as to make it completely unrealistic that they could reasonably represent a highly engineered system like an ISP or the Internet as a whole. In contrast, we observe that even simple heuristically designed and optimized models that reconcile the tradeoffs between link costs, router constraints, and user traffic demand result in configurations that have high performance and efficiency. At the same time, these designs are highly "non-generic" and "extremely unlikely" to be obtained by any random graph generation method. However, they are also "fragile" in the sense that even a small amount of random rewiring destroys their highly designed features and results in poor performance and loss in efficiency. Clearly, this is not surprising—one should not expect to be able to randomly rewire the Internet's router-level connectivity graph and maintain a high performance network!

One important feature of network design that has not been addressed here is *robustness* of the network to the failure of nodes or links. Although previous discussions of robustness have featured
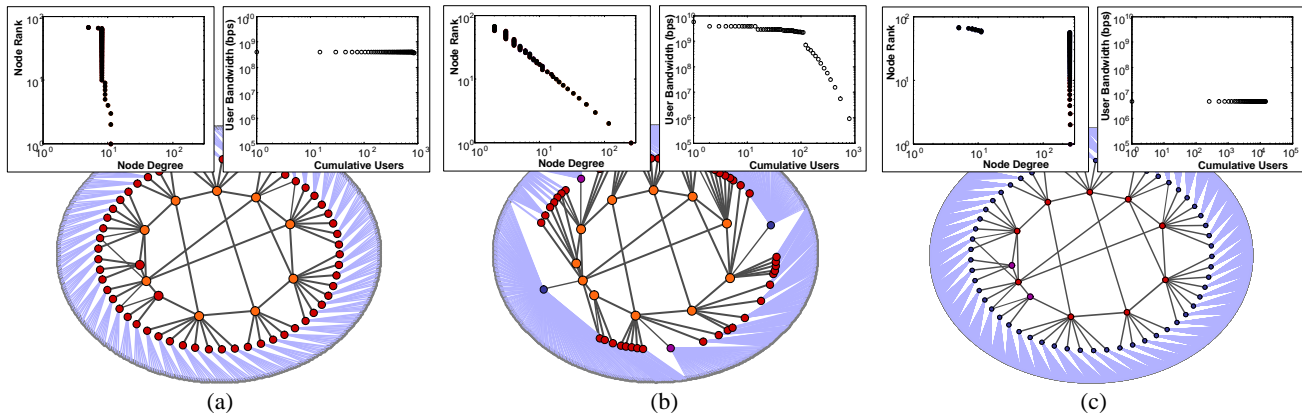
Figure 9: Distribution of node degree and end-user bandwidths for several topologies having the same core structure: (a) uniformly high bandwidth end users, (b) highly variable bandwidth end users, (c) uniformly low bandwidth end users.

prominently in the literature [4, 42], we have chosen to focus on the story related to performance and likelihood, which we believe is both simpler and more revealing. While there is nothing about our first-principles approach that precludes the incorporation of robustness, doing so would require carefully addressing the network-specific issues related to the design of the Internet. For example, robustness should be defined in terms of impact on network performance, it should be consistent with the various economic and technological constraints at work, and it should explicitly include the network-specific features that yield robustness in the real Internet (e.g., component redundancy and feedback control in IP routing). Simplistic graph theoretic notions of connected clusters [4] or resilience [42], while perhaps interesting, are inadequate in addressing the features that matter for the real network.

These findings seem to suggest that the proposed first-principles approach together with its implications is so immediate, especially from a networking perspective, that it is not worth documenting. But why then is the networking literature on generating, validating, and understanding network designs dominated by generative models that favor randomness over design and "discover" structures that should be fully expected to arise from these probabilistic models in the first place, requiring no special explanation? We believe the answer to this question lies in the absence of a concrete methodological approach for understanding and evaluating structures like the Internet's router-level topology. Building on [12, 48], this work presents such an approach and illustrates it with alternate models that represent a clear paradigm shift in terms of identifying and explaining the cause-effect relationships present in large-scale, engineered graph structures.

Another criticism that can be leveled against the approach presented in this paper is the almost exclusive use of toy models and only a very limited reliance on actual router-level graphs (e.g., based on, say, Mercator-, Skitter-, or Rocketfuel-derived data). However, as illustrated, our toy models are sufficiently rich to bring out some of the key aspects of our first-principles approach. Despite their cartoon nature, they support a very clear message, namely that efforts to develop better degree-based network generators are suspect, mainly because of their inherent inability to populate the upper-left corner in the likelihood-performance plane, where Internet-like router-level models have to reside in order to achieve an acceptable level of performance. At the same time, the considered toy models are sufficiently simple to visually depict their "non-generic" design, enable a direct comparison with their random counterparts, and explain the all-important tradeoff between likelihood and performance. While experimenting with actual router-level graphs will

be an important aspect of future work, inferring accurate router-level graphs and annotating them with actual link and node capacities defines a research topic in itself, despite the significant progress that has recently been made in this area by projects such as Rocketfuel, Skitter, or Mercator.

Any work on Internet topology generation and evaluation runs the danger of being viewed as incomplete and/or too preliminary if it does not deliver the "ultimate" product, i.e., a topology generator. In this respect, our work is not different, but for a good reason. As a methodology paper, it opens up a new line of research in identifying causal forces that are either currently at work in shaping large-scale network properties or could play a critical role in determining the lay-out of future networks. This aspect of the work requires close collaboration with and feedback from network engineers, for whom the whole approach seems obvious. At the same time, the paper outlines an approach that is largely orthogonal to the existing literature and can only benefit from constructive feedback from the research community. In either case, we hope it forms the basis for a fruitful dialogue between networking researchers and practitioners, after which the development of a radically different topology generator looms as an important open research problem.

Finally, we do not claim that the results obtained for the router-level topology of (parts of) the Internet pertain to logical or virtual networks defined on top of the physical infrastructure at higher layers of the protocol stack where physical constraints tend to play less of a role, or no role at all (e.g., AS graph, Web graph, P2P networks). Nor do we suggest that they apply directly to networks constructed from fundamentally different technologies (e.g., sensor networks). However, even for these cases, we believe that methodologies that explicitly account for relevant technological, economic, or other key aspects can provide similar insight into what matters when designing, understanding, or evaluating the corresponding topologies.

Acknowledgments

7. REFERENCES

[1] Abilene Network. Detailed information about the objectives, organization, and development of the Abilene network are available from http://www.internet2.edu/abilene.

[2] W. Aiello, F. Chung, and L. Lu. A Random Graph Model for Massive Graphs. *Proc. STOC 2000.*

[3] R. Albert, and A.-L. Barabási. Statistical Mechanics of Complex Networks, *Rev. of Modern Physics* (74), 2002.

[4] R. Albert, H. Jeong, and A.-L. Barabási. Attack and error tolerance of complex networks, *Nature* 406, 378-382 (2000)

[5] D. Alderson. Technological and Economic Drivers and Constraints in the Internet's "Last Mile", Tech Report CIT-CDS-04-004, California Institute of Technology (2004)

[6] D. Alderson, J. Doyle, R. Govindan, and W. Willinger. Toward an Optimization-Driven Framework for Designing and Generating Realistic Internet Topologies. In *ACM SIGCOMM Computer Communications Review*, (2003)

[7] A.-L. Barabási and R. Albert. Emergence of scaling in random networks, *Science* 286, 509-512 (1999)

[8] B. Bollobás and O. Riordan. Robustness and vulnerability of scale-free random graphs, *Internet Math.* 1, pp. 1–35, 2003.

[9] L. Briesemeister and P. Lincoln and P. Porras, Epidemic Profiles and Defense of ScaleFree Networks, *ACM Workshop on Rapid Malcode (WORM)* October, 2003.

[10] A. Broido and k. Claffy. Internet Topology: Connectivity of IP Graphs, *Proceeding of SPIE ITCom WWW Conf.* (2001)

[11] T. Bu and D. Towsley. On distinguishing Between Internet Power Law Topology Generators, *IEEE INFOCOM 2002.*

[12] K.L. Calvert, M. Doar, and E. Zegura., Modeling Internet topology, *IEEE Communications Magazine*, June (1997).

[13] J.M. Carlson and J.Doyle. Complexity and Robustness *PNAS*, 99, Suppl. 1, 2539-2545 (2002)

[14] H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. Towards Capturing Representative AS-Level Internet Topologies *Proc. of ACM SIGMETRICS 2002* (2002)

[15] H. Chang, S. Jamin, and W. Willinger. Internet Connectivity at the AS-level: An Optimization-Driven Modeling Approach *Proc. of MoMeTools 2003* (Extended version, Tech Report UM-CSE-475-03, 2003).

[16] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. The Origin of Power Laws in Internet Topologies Revisited, *Proc. IEEE INFOCOM 2002.*

[17] F. Chung and L. Lu. The average distance in a random graph with given expected degrees, *Internet Mathematics*, 1, 91–113, 2003.

[18] Corporation for Education Network Intitiatives in California (CENIC). Available at http://www.cenic.org.

[19] Cooperative Association for Internet Data Analysis (CAIDA), Skitter. Available at http://www.caida.org/tools/measurement/skitter/.

[20] M. B. Doar. A Better Model for Generating Test Networks. In *Proc. of GLOBECOM 1996.*

[21] P. Erdos and A. Renyi. On random graphs I *Publ. Math. (Debrecen) 9* (1959), 290-297.

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

[23] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On Power-Law Relationships of the Internet Topology, *SIGCOMM 1999.*

[24] L. Gao. On inferring autonomous system relationships in the Internet, in *Proc. IEEE Global Internet Symposium*, 2000.

[25] C. Gkantsidis, M. Mihail, A. Saberi. Conductance and congestion in power law graphs, *ACM Sigmetrics 2003.*

[26] R. Govindan and H. Tangmunarunkit. Heuristics for Internet Map Discovery, *Proc. IEEE INFOCOM 2000.*

[27] Internet2 Consortium. Internet2 NetFlow: Weekly Reports, Available at http://netflow.internet2.edu/weekly/.

[28] C. Jin, Q. Chen, and S. Jamin. Inet: Internet Topology Generator. Technical Report CSE-TR443-00, Department of EECS, University of Michigan, 2000.

[29] B.B. Mandelbrot. *Fractals and Scaling in Finance: Discontinuity, Concentration, Risk*. Springer-Verlag, 1997.

[30] A. Medina, A. Lakhina, I. Matta, and J. Byers. BRITE: An Approach to Universal Topology Generation, in *Proceedings of MASCOTS '01*, August 2001.

[31] A. Medina, I. Matta, and J. Byers. On the Origin of Power Laws in Internet Topologies, *ACM SIGCOMM Computer Communications Review* 30(2), 2000.

[32] M. Mitzenmacher. A Brief History of Generative Models for Power Law and Lognormal Distributions, *Internet Mathematics*. To appear. (2003)

[33] M.E.J. Newman Assortative Mixing in Networks, *Phys. Rev. Lett.* 89, 208701 (2002)

[34] M.E.J. Newman. The Structure and Function of Complex Networks, *SIAM Review* 45, 167-256 (2003)

[35] A. M. Odlyzko. Internet traffic growth: Sources and implications, in *Optical Transmission Systems and Equipment for WDM Networking II*, B. B. Dingel, W. Weiershausen, A. K. Dutta, and K.-I. Sato, eds., Proc. SPIE, vol. 5247, 2003, pp. 1-15.

[36] C. R. Palmer and J. G. Steffan. Generating network topologies that obey power laws. *Proc. GLOBECOM 2000.*

[37] R. Pastor-Satorras and A. Vespignani. Epidemic spreading in scale-free networks *Physical Review Letters*, 86(14), pp. 3200–3203, April 2, 2001.

[38] M. Roughan, A. Greenberg, C. Kalmanek, M. Rumsewicz, J. Yates and Y. Zhang. Experience in Measuring Backbone Traffic Variability: Models, Metrics, Measurements and Meaning *International Teletraffic Congress (ITC)* 18, 2003.

[39] Route Views, *University of Oregon Route Views Project*, Available at http://www.antc.uoregon.edu/route-views/.

[40] N. Spring, R. Mahajan, and D. Wetherall. Measuring ISP Topologies with Rocketfuel, *Proc. ACM SIGCOMM 2002.*

[41] L. Subramanian, S. Agarwal, J. Rexford, and R. Katz. Characterizing the Internet Hierarchy from Multiple Vantage Points, *Proc. IEEE INFOCOM 2002.*

[42] H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. Network Topology Generators: Degree-Based vs. Structural, In *Proc. ACM SIGCOMM 2002.*

[43] State of Washington Master Contract for Cisco Products, June 2002. Available from http://techmall.dis.wa.gov/ master_contracts/intranet/routers_switches.asp

[44] B.M. Waxman. Routing of multipoint connections, *IEEE Jour. of Selected Areas in Comm.*, Vol. 6, No. 9, 1988.

[45] 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)

[46] K.Wu and A. Liu. The Rearrangement Inequality http://matholymp.com/ TUTORIALS/Rear.pdf

[47] S.-H. Yook, H. Jeong, and A.-L. Barabási. Modeling the Internets large-scale topology, *PNAS* 99, 13382-86 (2002)

[48] E. Zegura, K.L. Calvert, and M.J. Donahoo. A quantitative comparison of graph-based models for Internet topology. *IEEE/ACM Transactions on Networking*, Vol. 5, No. 6, 1997.