Technical Report CIT-CDS 04-004

Technological and Economic Drivers and Constraints in the Internet's "Last Mile"

David L. Alderson*

Engineering and Applied Science, MS 107-81

California Institute of Technology

Pasadena, CA, USA, 91125

March 2004

Abstract

This paper investigates the physical topology of the Internet at the edge of the network, known as the "last mile", and considers the technological and economic features that drive and constrain its ongoing deployment and operation. In particular, by considering in detail the various technologies used in the delivery of network bandwidth to end-users, it is shown how the need to aggregate traffic is a dominant design objective in the construction of edge networks and furthermore that the large-scale statistics of network topologies as a whole, including features such as overall node degree distribution, are dominated by the structural features at the network edge.

Keywords: Internet topology, Internet Service Provider, network design, last mile.

*Email: alderd@cds.caltech.edu

1 The Ubiquitous Internet

Over the last two decades, the Internet has been transformed from a research network connecting university and government scientists, to a critical component of our economic and social fabric. Today, the Internet spans the globe and connects millions of people across all sectors of life[1]. It is a personal communication tool, allowing family and friends to engage in various forms of dialogue via email, instant messenger, and chat forums. It is a data storage and dissemination mechanism, allowing everyone from students to the general public to find a wealth of information on just about any topic imaginable. It is a worker productivity tool, enabling coworkers and business partners to share information and collaborate with unprecedented efficiency. And it is a transaction processing tool, facilitating broad participation in everything from electronic commerce to electronic voting. Underlying this growing importance of the Internet as a critical infrastructure is its ubiquity, that is, nearly universal access to this data and communication platform from our schools, our places of work, and our homes. This ubiquity is facilitated by the Internet's ability to support tremendous heterogeneity, in the constantly evolving applications that enable this interaction as well as in the physical technologies that constitute the network.

This paper investigates the design, structure, and function of the Internet in its "last mile" of connectivity to users at the network edge and addresses the technological and economic features that drive and constrain its ongoing deployment and operation. In particular, by considering in detail the demand for bandwidth by end users, the various technologies used in the delivery of that bandwidth, and the economics for network service providers who deploy and operate these networks, we provide insight into the drivers and evolution of last-mile connectivity. The main contribution is to show how the technological and economic considerations can dominate deployment at the network edge. And since a large portion of the total infrastructure is at the network edge (e.g., most computers connected to the Internet are at the edge), the design principles and resulting structural features can be expected to provide significant insight into the network as a whole. However, before addressing the particular issues related to the last mile, we review the fundamental design principles and basic architecture underlying the Internet.

2 Internet Architecture and Design

Much of the Internet's success can be attributed to a set of fundamental network design principles (e.g., layering, end-to-end, transparency) that were adopted by its early builders and that has yielded a layered architecture consisting of a suite of protocols (i.e., the "TCP/IP protocol stack")[2]. Layering is particularly important since it provides modularity by separating complex tasks, such as a Web transfer between two end hosts, into subtasks each of which is relatively simple and can be implemented independently in different layers. Each layer relies on the next lower layer to execute more primitive functions and provides services

[1] Recent estimates suggest that as of March 2004 there were 729 Million worldwide users of the Internet who were accessing it using dozens of languages [13].

[2] A thorough review of the Internet architecture is available from [30]; here, we comment on the points most salient to the discussion of topology at the network edge.

to the next higher layer. Two hosts with the same layering architecture communicate with one another by having their corresponding layers talk to one another through the use of formatted blocks of data that obey a set of rules or conventions known as a *protocol*.

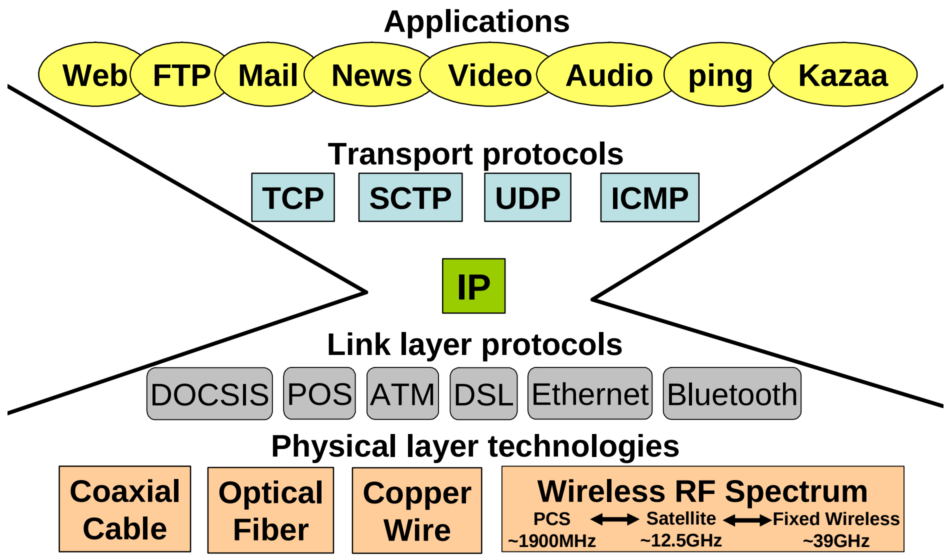

Figure 1: Internet hourglass architecture. The physical layer and link layer below the waist deal with the wide variety of existing transmission and link technologies and provide the protocols for running IP over whatever bit-carrying network infrastructure is in place ("IP over everything"). Above the waist is where the transport layer and application layer provide the protocols that enhance IP (e.g., at the transport layer, TCP ensures reliable transmission) and greatly simplify the process of writing applications (e.g., WWW) through which users actually interact with the Internet ("everything over IP"). Thus, the architecture of the Internet protocol stack uses selective homogeneity (in the form of IP as a type of "common currency") in order to support extreme heterogeneity in the applications at the top and the transmission media at the bottom.

This *vertical decomposition* that results from layering also gives rise to the "hourglass" metaphor for the TCP/IP protocol stack [22] (see Figure 1). In this perspective, the Internet protocol (IP) constitutes the waist-layer of the hourglass and provides a simple abstraction of a generic, yet unreliable data delivery service. This simple data delivery service can be implemented on top of a variety of physical technologies (e.g. optical fiber, copper wires, coaxial cable, radiowave spectrum) a feature commonly known as "IP over everything". The IP layer in turn provides a fundamental building block for an endless variety of applications (e.g. email, WWW, IP telephony, instant messenger) a feature similarly known as "everything over IP". In this manner, homogeneity at the data delivery layer in the form of IP enables heterogeneity in the physical networks that underlie the Internet and also provides heterogeneity in the applications that are supported. The ability of the Internet protocol stack to support and be supported by such a diversity of technologies has yielded incredible evolvability and has ultimately allowed the infrastructure as a whole to outlast its original design objectives. The fact that this hourglass design predated some of the most popular communication technologies, services, and applications in use today—and that within today's Internet, both new and old technologies and services can coexist and evolve—attests to the vision of the early architects of the Internet when deciding in favor of this layered design.

To increase the reliability and robustness of the network, each layer of this architecture is built using decentralized components having fully distributed provisioning, control, and management mechanisms[3] implemented through multiple layers of feedback regulation between different end systems. As a result of this additional *horizontal decomposition*, there exists a "network" of components at each layer of the architecture. Thus, there is no single notion of *network topology*, since one can consider the interconnection of a myriad of components at different layers of the architecture. For example, three distinct topologies that have received significant attention are the router-level graph (representing a type of machine-level connectivity), the WWW graph (representing the relationship between documents on the World Wide Web), and the autonomous system (AS) graph (representing organizational interconnectivity between subnetworks). We review each briefly in turn.

2.1 Router-level topology

The router-level graph reflects one-hop connectivity between routers running the Internet Protocol (IP). Information about the connectivity between routers can be inferred from *traceroute* experiments which record successive IP-hops along paths between selected network host computers (for example, the Mercator [14], Skitter [4], and Rocketfuel [25] projects). Ongoing research continues to reveal more and more idiosyncracies of traceroute data and shows that their interpretation requires great care and diligent mining of other available data sources. On the one hand, researchers of the Rocketfuel project [25] have used traceroute experiments to obtain a detailed and sophisticated mapping of a number of router-level ISP topologies, and their data represents the current state-of-the-art in router-level connectivity (see however [28] for further ambiguities concerning the Rocketfuel data). On the other hand, a recent study [18] pointed out another problem inherent to traceroute experiments that can give rise to a potential bias when inferring node degrees.

Unfortunately, a potential source of confusion for non-specialists is the dependence between the layer-2 (physical) and layer-3 (IP) issues associated with router-level connectivity, and how this somewhat subtle interaction can give the appearance of high connectivity at the IP-level. For example, at the physical layer the use of Ethernet at the network edge or Asynchronous Transfer Mode (ATM) technology in the network core can give the appearance of high IP-connectivity since the physical topologies associated with these technologies may not be seen by IP. In such cases, machines that are connected to the same Ethernet or ATM network may have the illusion of direct connectivity from the perspective of IP, even though they are separated by an entire network (potentially spanning dozens of machines or hundreds of miles) at the physical level. In an entirely different fashion, the use of Multiprotocol Label Switching (MPLS) at higher levels of the protocol stack can give the illusion of one-hop connectivity at the IP-layer. In both cases, it is the explicit and intended design of these technologies to hide the real network connectivity from IP.

[3] One noticeable exception is a centrally administered universal logical addressing scheme.

2.2 WWW topology

When representing the structure of the World Wide Web, network "nodes" are documents encoded with the Hyper Text Markup Language (HTML), and "links" are references within one document that point to another. Web documents are accessed using the Hyper Text Transfer Protocol (HTTP), which is the protocol implemented within web servers and web browsers at the application level of the protocol stack. Thus, the relationship between documents on the Internet can be mapped by applications that iteratively "crawl" the web by accessing one document, then accessing all of the documents that are linked from the first document, and so on. The resulting graph structure provides insight into the information content available on the web and perhaps even information about the use of web browsing applications, but it holds little meaning for networks that exist at lower layers of the protocol stack. Thus, the WWW graph is *virtual* in the sense that it does not correspond to any physical components in the network, and its connectivity is not constrained by any physical or geographic limitations. In fact, the WWW is an example of how the aformentioned architecture facilitates an endless number of virtual topologies on top of the physical infrastructure. Other, vastly different virtual topologies are the result of other applications. Two popular examples are *peer-to-peer (P2P)* file sharing programs that construct and utilize virtual networks of users who are willing to share data and *electronic mail* programs that implicitly map out the social relationships between individuals through their address books.

2.3 Autonomous System (AS) topology

In models of the AS topology, each node represents an autonomous system (AS) while a link between two nodes indicates the presence of a "peering relationship" between the two ASes, reflecting a mutual willingness to carry or exchange traffic. In this representation, a single "node" (e.g., AT&T) represents potentially hundreds or thousands of routers as well as their interconnections. Although most large ASes have have several peering points with other networks (e.g., AT&T has dozens of peering points with other ISPs), representing the network in this manner means that one is collapsing possibly thousands of router-level connections into a single link between two ASes. In this sense, the AS graph is yet another virtual graph representing peering (i.e., business) relationships among ASes (i.e., businesses). It is expressively not a representation of any physical aspect of the Internet structure, but is sometimes confused by nonexperts with representing some kind of physical Internet connectivity.

Directly measuring the AS connectivity is infeasible. The measurements that form the basis for inferring AS level connectivity consist of BGP routing table snapshots collected, for example, by the University of Oregon Route Views Project [24]. Here the *Border Gateway Protocol (BGP)* is the de facto standard inter-domain routing protocol deployed in today's Internet [26]. To illustrate the degree of ambiguity in the inferred AS connectivity data, note for example that due to the way BGP routing works, snapshots of BGP routing tables taken at a few vantage points on the Internet over time are unlikely to uncover and capture all existing connections between ASes. Indeed, [5] report that AS graphs inferred from the Route Views data typically misses between 20-50% or even more of the existing AS connections. Other ambiguities that are of concern in this context have to do with the dynamic nature of AS level connectivity, whereby new ASes can enter and existing ASes can leave, merge, or split at any time.

2.4 The Importance of Topology

A detailed understanding of the many facets of the Internet's topological features is critical for evaluating the performance of networking protocols, for developing improved designs for resource provisioning, or for understanding system-wide robustness properties that result from its structure. Recent attention on the large-scale topological structure of the Internet has been heavily focused on statistical properties of graphs modeling the *connectivity* of network components, whether they be machines in the router-level graph [14, 3] or entire subnetworks (Autonomous Systems) in the AS-level graph [11, 5]. However, a fundamental challenge in using large-scale statistical features to characterize something as complex as the topology of the Internet is that it is difficult to understand the extent to which any particular observed feature is "fundamental" to its structure. The remainder of this paper deals with *router-level* models of Internet topologies at the network edge. The focus is on design and deployment of network topologies in the "last mile", that is the portion of the physical infrastructure that is closest to the end user. By focusing on the technological and economic forces shaping the physical deployment at the network edge, the intent is to provide a complementary approach to these abstract studies, one in which we examine the practical considerations related to the manner in which real networks are designed and built.

3 Drivers of Network Topology

As originally outlined in [1], the premise for this investigation is that any explanatory framework for router-level Internet topology modeling should incorporate both the *economic* and *technological* factors faced by Internet Service Providers (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 network that it can support. At the same time, in designing the topology that best supports its business, the ISP is constrained by the technologies currently available to it. While a complete review of these issues is beyond the scope of this paper, these drivers in their simplest form can be understood in terms of *link costs* and *router technology.*

3.1 Link Costs

On a fundamental level, the purpose of a network like the Internet is to transcend distance by facilitating communication between parties that are located in different geographic locations. To do so requires the installation, operation, and maintenance of communication links that span great distances. For the purposes here, we use the term "link" to mean both the network cable and the equipment used to send traffic along that cable[4]. While a signficant portion of the link cost is often associated with obtaining the "right of way" to install the network cables, there is generally an even greater cost associated with the installation and maintenance of the

[4]At the national level, the cables are usually fiber optic and the equipment consists of transmitter/receivers at the end points and signal repeaters along the way.

equipment used to send the traffic across these cables. Both the installation and maintenance costs tend to increase with link distance. Thus, one of the biggest infrastructure costs facing a network provider is the cost associated with the deployment and maintenance of its links.

National ISPs are one type of network provider for which link costs are significant. However, their challenge in providing network connectivity to millions of users spread over large geographic distances is made somewhat easier by the fact that most users tend to be concentrated in metropolitan areas. Thus, there is a natural separation of the connectivity problem into providing connectivity within a metropolitan region and providing connectivity between these regions[5]. In considering the costs associated with providing connectivity between metropolitan regions, the ISP has strong economic incentive to spread the cost of an intercity link over as many customers as possible. This is the basic motivation for *multiplexing*—a fundamental concept in networking by which a link is shared by many individual traffic streams.

Multiplexing is one of the most basic design principles in networking and has tremendous impact on the types of topologies chosen by network architects. In it simplest form, it states that the only type of design that makes sense from an economic perspective is one that aggregates as much traffic on the fewest number of long distance links. This principle applies at all levels of network design, including the local and regional levels, and not just the national backbone. Thus, for economic reasons alone one would expect that individual links at the edge of the network have relatively low bandwidth and that the link capacities increase as one moves to the network core. Although link costs are not the only factors affecting network design, an inspection of real networks reveals a common theme, namely that the topology at the network edge is designed to aggregate traffic within a local region, while the topology within the core of the network is designed to transport aggregated traffic between geographically disparate regions[6].

3.2 Router Technology

Another major constraint affecting the types of topologies available to network designers is related to the technologies used in routing equipment. Based on the technology used in the cross-connection fabric of the router itself, a router has a maximum number of packets that can 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 model type. This limitation creates an "efficient frontier" of possible bandwidth-degree combinations available 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 the router must obey a form of *flow conservation* in the traffic that it can handle. While it is always possible to configure the router so that it falls below the efficient frontier (thereby underutilizing the router capacity), it is not possible to exceed this frontier (for example, by having an ever

[5]Within the ISP industry, this distinction often separates service offerings into two lines of business known as "metro service" and "long-haul service" respectively.

[6]The telephone network is subject to the same economics associated with link costs and also exhibits the same type of network design in which local traffic is aggregated along "trunks" which interconnect local regions. Given a history in which the modern data networks grew out of traditional phone networks, the reuse of commonly accepted and successful design principles is not surprising.

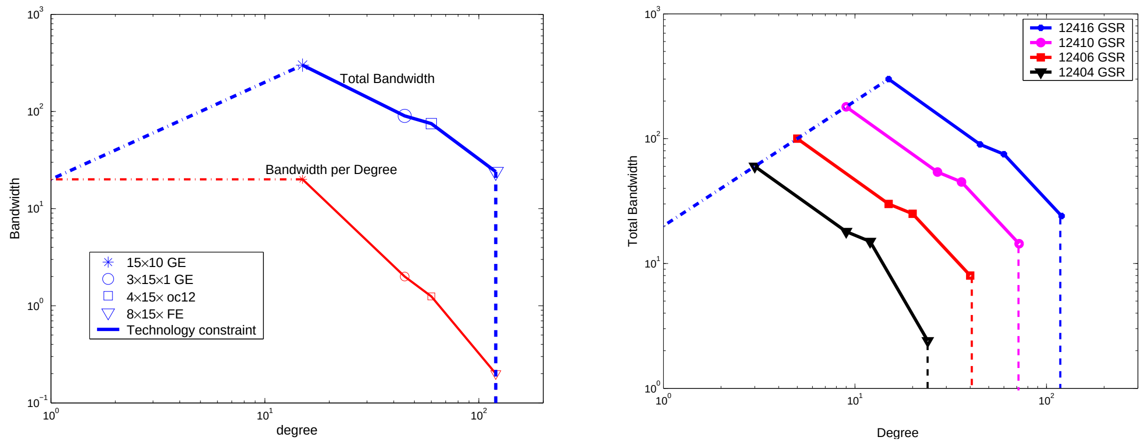

Figure 2: **(a) Technology constraint for Cisco 12416 Gigabit Switch Router (GSR): Degree vs. maximum throughput.** The router can achieve any combination of maximum throughput and degree below the technology constraint line. Each point on the plot corresponds to routers with different interfaces and the price as shown in the enclosed inbox. **(b) Technology constraints for GSR Models 12404, 12406, 12410, 12416.** Each line represents one type of router.

increasing number of high bandwidth connections). For any particular router model, there will be a frontier representing the possible combinations that are available. Router models with greater capacity are generally more expensive.

Consider as an example the Cisco Gigabit Switch Routers (GSRs), which are one of the most widely deployed routers within the Internet backbone[7]. In Figure 2(a), we show an example of the technology constraint of the Cisco 12416 GSR. When the number of ports is less than 15, the throughput of each port is limited by the maximum speed of supported line cards (10 GE) and the router's maximum throughput increases with the number of ports. When the number of ports is greater than 15, the maximum router throughput *decreases* as the total number of ports increases. The reason for this decrease is related to an increased routing overhead for handling traffic over a greater number of ports. In particular, router line cards receiving incoming traffic process it based on its final destination address in order to assign it to an outgoing port that sends the traffic to the next hop of its journey. Since the number of possible input-output traffic assignment combinations grows as n^2, where n is the number of ports, there is greater processing overhead for a router as its total number of ports increases. Figure 2(b) illustrates the efficient frontiers of several Cisco GSR routers taken from a recent product catalog [6]. Although engineers are constantly increasing the frontier with the development of new routing technologies, 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.

As noted in Figure 2, these high capacity core routers can have node degree on the order of only 100 direct connections. As a result, observed IP measurements for nodes having

[7] As reported in [8], Cisco's share of the worldwide market for service provider edge and core routers was approximately 70% during 2002.

thousands of connections cannot correspond to physical connections between these routers. Until new technology shifts the frontier, the only way to create throughput beyond the frontier is to build networks of routers[8]. In making this claim, 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 of these products will enable their wide deployment within the Internet.

4 Connectivity at the Network Edge

There is tremendous diversity in the technologies and topologies used at the network edge. These include campus networks, DSL networks, cable networks, dial-up networks, and wireless networks. A common theme in the design of these networks is the need to aggregate traffic while leveraging existing wiring so as to minimimze the wiring cost. Since these technologies are all implemented at layer-2 (physical layer), the details of their connectivity are hidden from IP, and the relatively high "fan-out" of their connectivity can give the appearance of very high IP connectivity. We consider the technical features for each type of network in turn.

4.1 Campus Networks

Campus networks constitute a significant portion of the total edge connections in the Internet. For the purposes here, we use the term "campus" to mean any local collection of buildings (this could be an office park or a university campus) that share a high-bandwidth Internet connectivity.

As seems to be the dominant trend in campus network design, we assume that connectivity within the Local Area Network (LAN) is achieved using Ethernet. In this setup, a computer on a campus is directly connected to an Ethernet switch based on its local geography, since the Ethernet protocol specifies a maximum connection distance for each supported speed (i.e., the maximum connection distance for a twisted-pair 100Mb connection is 100 meters). Most commercial Ethernet switches can support approximately 24 direct connections, so multiple switches must be used to aggregate greater numbers of machines. Within a large office building, this aggregation often occurs on each floor, where one or more ethernet switches may reside in an Intermediate Distribution Facility (IDF), usually a network "closet". The IDFs on different floors are connected in turn to a MDF (Main Distribution Facility) in the same building. In this manner, connectivity within an individual building takes the form of a tree, with constraints on the "fan out" for each level of hierarchy, and often without any redundancy in the connections.

Connectivity from each building is provided by one or more connections from the MDF in the building to a campus "backbone". The size and topology of this backbone will vary with the size and available budget of the campus, but it generally provides for some level of

[8]A few companies such as Avici Systems (www.avici.com) have started to offer scalable routing devices built from "stacks" of routers, with some recent success in the marketplace [27].

redunancy in its connectivity. For example, the campus backbone at Caltech logically divides its more than 100 buildings into primary, secondary, and third level sites. Most buildings on campus are third level sites, which are connected to two other primary or secondary sites depending on their geographical location (i.e., distance). There are about a dozen secondary sites and each is connected to two primary sites. The primary sites at Caltech are connected to each other by two connections (one is the backup). See Figure 3 for a schematic representation of the campus network at Caltech. In this manner, the Campus Network is roughly

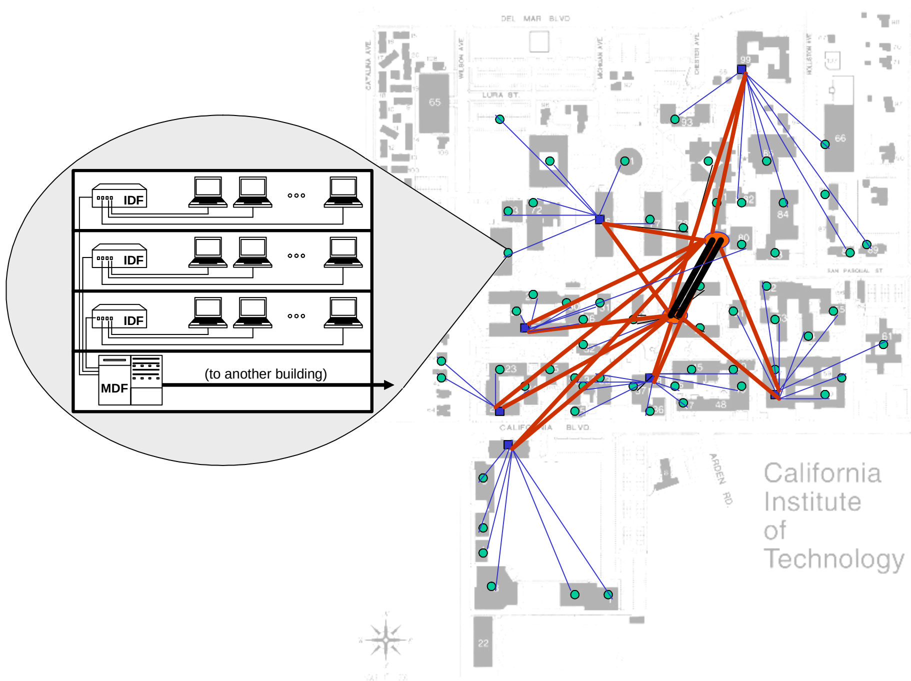

Figure 3: Qualitative design of a wired campus network. (a) Sample layout of machines within a single building. (b) Each node represents a building and is connected to other buildings with some redundancy. Within Caltech, each node connects to two other nodes at a higher level within the tree.

a tree structure with a minimal amount of redundancy at its upper branches. The campus backbone connects to the outside world through one or more regional providers. For example, Caltech is connected to the CENIC educational network and the LosNettos commercial network.

4.2 DSL Networks

Digital subscriber line (DSL) technology allows for high-bandwidth network connectivity over standard (copper) telephone wires. The technology works through the use of a DSL modem at the customer site which connects to an aggregator modem at the Central Office (CO) of the local phone company (see Figure 4). Most standard DSL service offerings provide for guaranteed download speeds at either 768kbps or 1.5Mbps, however upstream speeds are usually significantly less. The available bandwidth for DSL will vary with connection distance (the closer the better), with a hard limit that each connection must be less than 18,000 feet from the CO. As a result, DSL service is not available in all areas, although its growing popularity have made ongoing deployment a priority for local ISPs.

Aggregation of individual users occurs through the use of a DSL access multiplexer (DSLAM) which aggregates individual lines into larger bandwidths before connecting them to a gateway router. For example, the Cisco 6260 Access Concentrator can manage 200+ simultaneous DSL connections (each at 768kbps) into a single 155 Mbps (OC-3) upstream connection. In other words, through the use of DSL technology, a single OC-3 connection at a gateway router can support more than 200 logical connections to end hosts. High speed routers that support multiple OC-3 connections can increase the total connectivity at the CO significantly. For example, a gateway router with a single 10Gbps connection to the network core could support more than 10,000 end hosts, each of which would appear from the perspective of IP to have a direct connection to the gateway router. However, the distance restriction on DSL technology would require that all 10,000 users be located in a relatively small geographic area, and this constraint by itself severely limits such deployments to high population areas. Current deployment statistics (see Figure 1) suggest that the average numbers are significantly lower.

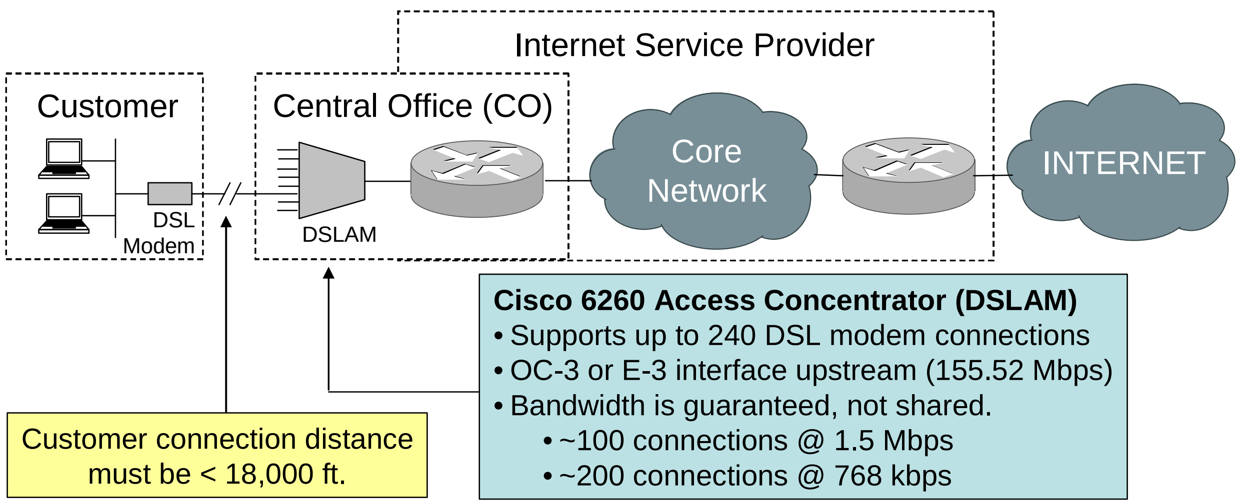

Figure 4: Schematic representation of a DSL Network. Individual users connect to aggregation equipment (DSLAMs) over copper telephone wires. The DSLAMs give the appearance of very high connectivity at the gateway routers.

4.3 Cable Networks

Broadband cable service provides another source of high speed network connectivity for residential customers. This service leverages unused capacity over coaxial cable to transmit data—downstream data travels over a single shared 6MHz channel while upstream data travels over a single shared 2MHz channel (see Figure 5). Available bandwidth does not depend on connection distance, however the 6MHz channel can only carry 30-40 Mbps, and

	Population*	Telephone Lines*	Central Office Switches**	DSL Lines*	DSL / 100 population	DSL / CO Switch
USA	275,130,000	192,519,000	17,140	4,363,900	1.59	254
World	4,135,211,000	813,664,600	n/a	17,211,700	0.42	n/a
	*Source: Point Topic Ltd, April 2002 [23]					
	**Source: Federal Communication Commission, December 2002 [10]					

Table 1: *Statistics on the number of DSL subscribers for 2002.* Using the number of CO switches as proxy for the number of COs is not entirely correct, since some switches are "remote" and may not be in a CO, while in other cases a CO may have more than one switch. In addition, DSL is not currently available in every CO. Thus, the estimate of 254 DSL lines per CO switch underestimates the average DSL lines per CO, but it indicates that the number of lines is *at least* in the hundreds.

this bandwidth is shared among all active users. Thus, a single fiber node in a residential neighborhood that supports 100 active users is able to deliver approximately 300 kbps to each user. Although there is usually not a hard technical limit on the number of users that can be supported from a single fiber node, statistics suggest that the average number of broadband cable users per headend is approximately 1,100 (see Table 2).

End hosts use a cable modem to connect over the coaxial network to a Cable Modem Termination System (CMTS) which acts as a gateway providing connectivity to the ISP network backbone. The density of customers connected to the CMTS (i.e., degree) will depend on the equipment in use, however several Cisco products (e.g., the uBR7200 CMTS and uBR10012 CMTS) advertise the ability to support thousands of end hosts while guaranteeing minimal service quality. Again, since the underlying cable technology is hidden to IP, each of these end hosts would appear to IP as having direct connectivity to the gateway router (CMTS) at the network edge.

Total Households* (Nov. 2002)	106,712,845
Cable Headends* (May 2003)	10,273
Basic Cable Customers* (Nov. 2002)	73,525,150
Customers Per Cable Headend (Estimated)	7,157
Broadband Cable Subscribers** (Dec. 2002)	11,265,200
Broadband Customer Frequency (Estimated)	$\sim$ 15%
Broadband Customers Per Cable Headend (Estimated)	$\sim$ 1100
*Source: Point Topic Ltd. (http://www.point-topic.com)	
**Source: National Cable & Telecommunications Association (http://www.ncta.com/)	

Table 2: *Estimate of the number of broadband cable customers per fiber headend,* based of statistics on the number of television and cable households.

4.4 Dial-Up Networks

While traditional dialup connections provide inferior connection speeds to the aforementioned broadband services, they still constitute a significant portion of the total number of connections at the network edge. Dialup connections leverage existing telephone (copper)

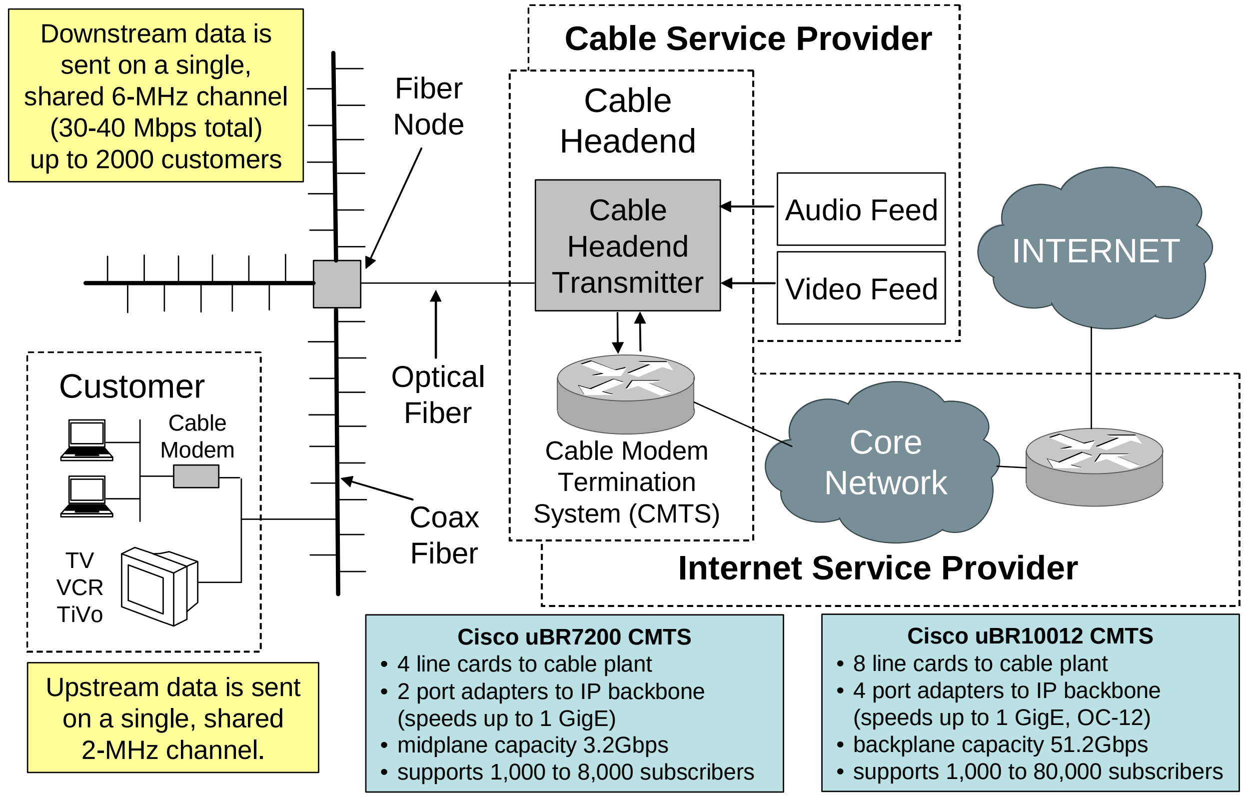

Figure 5: Schematic representation of a cable network. Individual users connect to a traffic aggregator (CMTS) over coaxial cable lines. As a gateway router, the CMTS has the appearance of very high connectivity. The bandwidth performance seen by individual users depends on the total number of active users.

wiring to provide connections speeds typically at 56kbps. The end host uses a modem to connect over the public telephone network to aggregation equipment at the ISP (see Figure 6).

The aggregation of dial-up traffic from individual end hosts to the ISP backbone is similar to that for DSL or cable, except that the small bandwidths of each connection requires significantly more aggregation to achieve high bandwidth flows. As a result, this aggregation often occurs in a hierarchical manner. The bandwidth capacity of typical gateway routers allows them easily to handle thousands or tens of thousands of connections. Thus, a gateway router supporting dial-up traffic may have thousands of "direct" IP connections to the end hosts.

4.5 Wireless Networks

Wireless networks are becoming increasingly popular in office buildings, on campuses, in public places such as coffee shops or airports, and even within the home. Most wireless networks run a version of Ethernet, and their edge connectivity is similar to that of wired Ethernet except that the transfer of data occurs over a wireless channel. Because the wireless signal degrades with distance, each wireless "base station" can support only users within close geographic proximity. Almost without exception, these wireless base stations act as

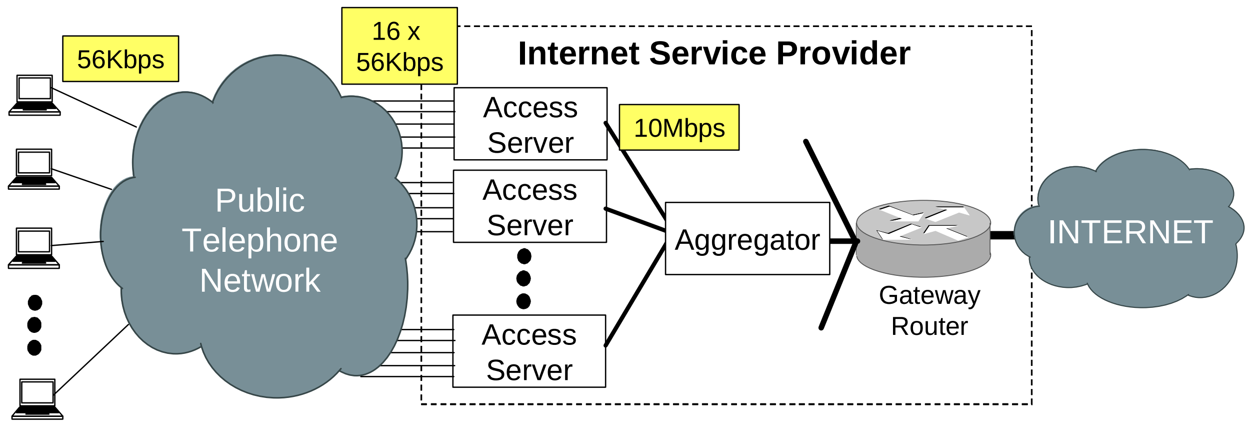

Figure 6: Schematic representation of a dialup network. Individual 56kbps lines are aggregated (possibly in stages) to gateway routers, which can "see" the IP addresses of thousands of end hosts.

gateways which connect directly into the wired network. Thus, wireless networks serve as simply another method of aggregating traffic at the network edge.

5 Current Internet Usage

It is clear that the selection of end-user technologies by an ISP will determine in part the bandwidths delivered to its customers. However, in order to get a realistic sense of how the use of these technologies translates to the overall degree distribution for connectivity at the network edge, one must consider current Internet usage. In general, accurate market statistics about Internet usage are hard to obtain and even published reports should be viewed with a certain amount of skepticism. Nonetheless, in this section, we provide a first-order estimate of these numbers for the United States.

According to Internet Software Consortium, the total number of unique hosts belonging to commerical, educational, and government domains in January 2003 is more than 113 Million [16]. Of course, the assumption that all of these hosts reside within the U.S. is not correct, but since all non-U.S. countries have their own domains (e.g., `.uk` for the United Kingdom, `.jp` for Japan), this approach is reasonable as a first approximation. A partial breakdown of unique hosts among individual domains is listed in Table 3.

We assume that each of these hosts represented above either connects to a campus network or uses a residential (i.e.,non-campus) connection. A leading on-line resource for ISPs estimates that the total number of non-campus users within the U.S. was approximately 75.3 Million in 2003 [17]. The corresponding breakdown of the top ten ISPs and their market share is shown in Table 4.

Next, we assume that non-campus ISP subscribers can be categorized uniquely according to the type of edge network to which they belong, and furthermore, we assume that we need to consider only broadband DSL, broadband cable, and dialup connections. A leading market research firm estimates that there were more than 20 million total broadband customers within the U.S. in 2003, and of those approximately 7.5 million of them were DSL subscribers while approximately 13.4 million of them were cable subscribers [23]. A comparison of these figures with global statistics is presented in Table 5.

Domain	Unique Hosts
.net	61,945,611
.com	40,555,072
.edu	7,459,219
.us	1,735,734
.org	1,116,311
.gov	607,514
Total	113,419,461
World Total	171,638,297

Table 3: *Distribution of Top-Level Domain Names by Host Count, January 2003.* Source: Internet Software Consortium (http://www.isc.org/).

Rank	ISP	Millions of Subscribers	Market Share
1	America Online	25.3	27.5%
2	MSN	8.7	9.4%
3	Netzero & Juno	5.2	5.6%
4	EarthLink	5.0	5.5%
5	Comcast	4.4	4.8%
6	SBC	2.8	3.0%
7	Verizon	1.9	2.1%
8	Cox	1.7	1.8%
9	Charter	1.3	1.5%
10	BellSouth	1.2	1.3%
	Total Among Top 10 ISPs	57.5	62.5%
	Total Among All ISPs	75.3	100%

Table 4: *Top U.S. ISPs by Subscriber: Q2 2003.* These numbers do not include: subscribers at work, at universities, or in government. It only includes residential consumer accounts. Source: ISP-Planet (http://www.isp-planet.com/research/rankings/usa.html).

Finally, these assumptions and statistics allow us to compile a composite view of the relative frequency of the various edge network technologies, shown in Table 6. We estimate that dial-up users account for nearly one-half of all edge connections, while campus users account for nearly one-third of all connections. Broadband DSL and cable account for the remainder with cable nearly twice as popular as DSL. If this perspective is correct and nearly one half of the current connections to the Internet are supported by dial-up access, then the aggregation requirements at the network edge are tremendous.

6 Modeling Implications

Recently, there has been considerable debate about the prevalence of heavy-tailed distributions in node *degree* (e.g., number of connections) and whether or not these heavy-tailed

	Thousands of lines		
	DSL	Cable	Total
USA	7,576	13,367	20,943
World	46,683	30,358	77,041

Table 5: *DSL lines, cable modems and total broadband lines in major countries at 30 June, 2003.* Source: Point Topic Ltd. "World Broadband Statistics: Q2 2003", September 2003 (http://www.point-topic.com).

Type of Edge Connection	Typical Connection Speed	Approximate Number of Connections (Millions)	Relative Frequency
Campus Users	1.544Mbps(T-1) – 10Gbps(OC-192)	38.1	33.6%
Broadband DSL	512kbps – 6Mbps	7.6	6.7%
Broadband Cable	300kbps – 30Mbps	13.4	11.8%
Dial-Up	56kbps	54.4	47.9%
Total		113.4	100%

Table 6: *Estimated distribution of end host connection types in the United States for 2003.* It is important to note that the bandwidth performance seen by an individual user may be less than the total connection speed if the user's network interface card (NIC) is relatively slow. For example, a user on a university campus with a Fast Ethernet (100Mb) card will never achieve more than 100Mbps even if the university has a 10Gbps connection to the Internet.

distributions conform to power laws [9, 21, 7, 20]. In the context of router-level graphs, the presence of heavy-tailed node degree distributions means that while most nodes have relatively few connections, a few nodes have orders of magnitude more connections. The picture presented here is consistent with that general finding, and it argues that the need for traffic aggregation at the network edge can be expected to result at times in high connectivity in the topology of the Internet's last mile. Given the tradeoffs between bandwidth and connectivity for routing equipment, the wide variability in the connection speeds of end users additionally suggests that one should expect wide variability in the measured degree of nodes at the network edge. Since it is generally accepted that most of the computers in the network are at its edge, it is reasonable to expect that the overall connectivity statistics of the network are dominated by the degrees of the nodes near its edge. Given that population densities themselves range widely by geography, it should not be surprising that connectivity at the network edge also display wide variability. For example, according to 1990 U.S. Census data[29], the most densely populated county was New York County in New York State with a density of 20,239 people per square kilometer, while the least densely populated county was Lake and Peninsula Borough, Alaska with a density of 0.27 people per square kilometer. The overall distribution is heavy-tailed and follows what is commonly known as a "power law" (see Figure 7).

These considerations suggest a natural starting point for a "fictitious, yet realistic" model of a national ISP in which network topology is constructed to meet the demands of geographically distributed users. The design of these networks can be accomplished by separating the network design problem into local and long-haul subproblems. Using the type of

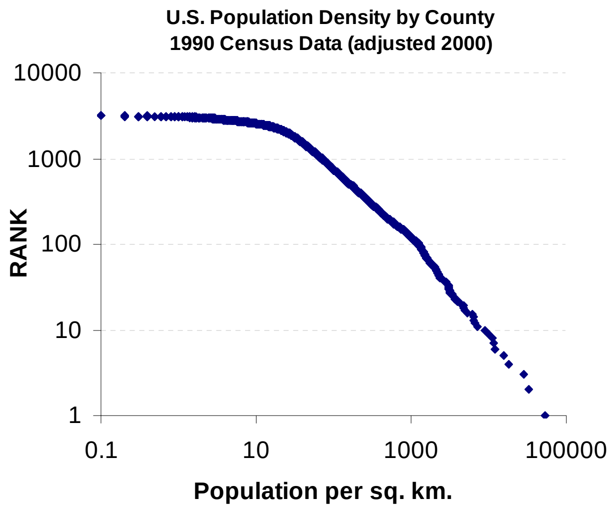

Figure 7: Population density of the United States by county in 1990. Most counties are sparsely populated, but a few counties have extremely high densities.

optimization-based framework discussed in [1], each subproblem can be explicitly formulated and solved. For example, local edge networks can be designed for the specific technological considerations outlined in this report and should be consistent with the usage exhibited by end users. In parallel, the design of backbone networks supporting the aggregate demand of local geographies can be formulated using router-level technology constraints identified in this report and solved using traditional optimization techniques [12]. While significant work remains before realistic generative models of network topology are available, this approach is already providing new insight into the design and structure of router-level topologies [19, 2].

Acknowledgments

This report has resulted from conversations with many individuals, and represents part of an ongoing collaboration with John Doyle, Walter Willinger, and Lun Li. Specifically, Lun Li assisted in the preparation of Figures 2 and 3. Additional help in the preparation of this material was provided by John Dundas at Caltech, Heather Sherman at CENIC, and Stanislav Shalunov at Internet2.

References

[1] 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*, January 2003.

[2] D. Alderson, W. Willinger, L. Li, and J. Doyle. The Role of Optimization in Understanding and Modeling Internet Topology. Submitted, 2004.

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

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

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

[6] Cisco Catalog. Master Contract for the State of Washington. http://techmall.dis.wa.gov/master_contracts/intranet/routers_switches.asp

[7] 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)

[8] J. Duffy. "Cisco's loss is Juniper's gain." *Network World Fusion* February 18, 2003. Available electronically at http://www.nwfusion.com/edge/news/2003/0218mktshare.html.

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

[10] Federal Communications Commission. Statistics of Communications Common Carriers 2002. Available from http://www.fcc.gov/web/stats/.

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

[12] B. Gendron, T.G. Crainic and A. Frangioni. Multicommodity Capacitated Network Design. In B. Sansó and P. Soriano (eds), Telecommunications Network Planning, pages 1–19. Kluwer, Norwell, MA. 1998.

[13] Global Reach, "Global Internet Statistics (by Language)". Available on-line at http://global-reach.biz/globstats/.

[14] R. Govindan and H. Tangmunarunkit. Heuristics for Internet Map Discovery, *Proceeding of IEEE INFOCOM* (2000)

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

[16] Internet Software Consortium. "Distribution of Top-Level Domain Names by Host Count, January 2004." Available on-line at http://www.isc.org/.

[17] ISP Planet. "Top U.S. ISPs by Subscriber: Q2 2003." Available on-line at http://www.isp-planet.com/research/rankings/usa.html.

[18] A. Lakhina, J.W. Byers, M. Crovella, and P. Xie Sampling Biases in IP topology Measurements, *Proc. IEEE INFOCOM* (2003)

[19] L. Li, D. Alderson, J. Doyle, and W. Willinger. A First-Principles Approach to Understanding the Internet's Router-level Topology. *Proc. ACM SIGCOMM*, Portland, Oregon, Aug. 30-Sept. 2, 2004.

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

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

[22] Computer Science and Telecommunications Board (CSTB), National Research Council. *The Internet's Coming of Age.* National Academy Press, Washington, D.C., 2001.

[23] Point Topic Ltd. "World Broadband Statistics: Q2 2003", September 2003. Available on-line at http://www.point-topic.com.

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

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

[26] J. W. Stewart III. *BGP4: Inter-Domain Routing in the Internet.* Addison-Wesley, Boston, 1999.

[27] SURFnet Press Release. "SURFnet Builds Advanced Research Network Based on Solutions, Services from Nortel Networks, Avici Systems, Telindus". Utrecht, 24 March 2004. Available at http://www.gigaport.nl/publicaties/pers/en_pers240304.html.

[28] R. Teixeira, K. Marzullo, S. Savage, and G.M. Voelker. In Search of Path Diversity in ISP Networks, In *Proc. IMC'03*, Miami Beach, Florida. October 27-29, 2003.

[29] United States Census Bureau. Census Data for 1990, Available electronically from http://www.census.gov/main/www/cen1990.html.

[30] W. Willinger and J. C. Doyle. Robustness and the Internet: Design and Evolution. In *Robust design: A Repertoire of Biological, Ecological, and Engineering Case Studies*, E. Jen, Editor, Oxford University Press (to appear).