

DIMACS Technical Report 99-11
February 1999

The Size and Growth Rate of the Internet

by

K. G. Coffman and A. M. Odlyzko ¹

AT&T Labs - Research

kgc@research.att.com, amo@research.att.com

¹Permanent Member

DIMACS is a partnership of Rutgers University, Princeton University, AT&T Labs-Research, Bell Labs, Bellcore and NEC Research Institute.

DIMACS is an NSF Science and Technology Center, funded under contract STC-91-19999; and also receives support from the New Jersey Commission on Science and Technology.

ABSTRACT

The public Internet is currently far smaller, in both capacity and traffic, than the switched voice network. The private line networks are considerably larger in aggregate capacity than the Internet. They are about as large as the voice network in the U.S., but carry less traffic. On the other hand, the growth rate of traffic on the public Internet, while lower than is often cited, is still about 100% per year, much higher than for traffic on other networks. Hence, if present growth trends continue, data traffic in the U.S. will overtake voice traffic around the year 2002 and will be dominated by the Internet.

1. Introduction

There are many predictions of when data traffic will overtake voice. It either happened yesterday, or will happen today, tomorrow, next week, or perhaps only in 2007. There are also wildly differing estimates for the growth rate of the Internet. The number of Internet users is variously given as increasing at 20 or 50 percent per year, and the traffic on the Internet is sometimes reported as doubling every three months, even in sober government reports such as [DOC]. Often the same source contains seemingly contradictory information. For example, John Sidgmore, the chief operating officer for WorldCom, and the person in charge of all its Internet activities, gave an interview (apparently held in early 1998) that was published in [Sidgmore]. He stated that revenues from Internet operations at WorldCom were about doubling each year. Later in the interview, he said that the bandwidth of UUNet's Internet links was increasing 10-fold each year. Since the prices that UUNet charges have not decreased recently, certainly not by a factor of 5, both of these claims can be correct only if something unusual is happening to the WorldCom network.

Since there is no comprehensive source for information on the size and growth rate of the Internet, it seemed worthwhile to do as careful an analysis as the fragmentary publicly available data allows. This is especially important because, as we point out below, the growth rate of the Internet has not been stable.

In this study we focus on the sizes of networks (measured by their transmission capacity) and the traffic they carry (measured in bytes). We find that the public Internet (the part of the Internet that is not restricted to users from any single organization) is still small, whether measured in capacity of links or in traffic, when compared to the voice telephone network. (This may not be true for all routes. There are frequent reports, for example, that there is more Internet traffic than voice traffic between the U.S. and Scandinavia.) In contrast, the private parts of the Internet (largely corporate private line networks) already have capacity close to that of the voice network. On the other hand, traffic on private line networks is still much smaller than that on the voice network, possibly not much bigger than traffic on the public Internet.

In modeling the transition to a network dominated by data it appears important to recognize three distinct growth rates:

(i) Interstate voice traffic (which carries some fax and modem data) has been growing recently about 8% a year when measured in minutes of use. This is an acceleration from the 4% rates of the early 1990s, but not as fast as the record 23% increase in 1984, or the average of 10.3% per year that was observed during the 1980s. (See Table 12.1 in [FCC2] or [FCC3] for precise numbers. Revenue growth has been far lower as prices have dropped, which was undoubtedly the major reason for the increase in calls.)

(ii) Capacity and presumably traffic on private line networks have been growing 15% to 20% per year in the last few years [VS].

(iii) Traffic and capacity of the public Internet grew at rates of about 100% per year in the early 1990s. There was then a brief period of explosive growth in 1995 and 1996. During those two years, traffic grew by a factor of about 100, which is about 1,000% per year. In 1997, it appears that traffic growth has slowed down to about 100% per year.

Traffic on the Frame Relay semi-public data networks is also growing about 100% per year. (We refer to Frame Relay and ATM networks as semi-public, since they carry traffic from many sources, but almost universally from a source within an organization to a destination within the

same organization. In contrast, public networks like the voice network and the Internet allow connections of any source to any destination.)

Reports, such as [DOC], which claim 1,000% growth rates for the Internet, appear to be inaccurate today, since they are based on a brief period of anomalously rapid growth a short while ago. (The claim in [DOC] is based on the report [Inktomi] that uses data no more recent than the end of 1996, the end of the period of abnormally fast growth.) Still, even a doubling each year is fantastically fast by the standards of the communications industry.

If traffic on the Internet continues to double each year, data should exceed voice on U.S. long distance networks around the year 2002. (We are using data here to refer to packet traffic, and voice to the circuit switched traffic, whether those are used to carry voice or data.) There are obvious uncertainties in making such projections, but in Section 8 we discuss the long historical trend of consistently rapid growth of the Internet, and the reasons we expect it to continue.

Our estimates of the transition date from domination by voice to domination by data differ from many published ones. Most of the claims (such as that of [Thyfault], which says that the bandwidth of data networks will equal that devoted to voice in 2000) are not substantiated by detailed analysis and appear incorrect, as that transition is occurring about now, it appears. The analysis of [MutooniT] appears to go astray by assuming utilization rates of data and voice networks are about the same. However, as is shown in the companion paper [Odlyzko1] and discussed at greater length in Section 7 below, data networks typically are used much less intensively than voice networks. Thus network capacities do not represent the amount of traffic those networks carry. Finally, there are claims (such as in some of the presentations of Vint Cerf [Cerf]) which may agree with our prediction of a transition around 2002, but they are not accompanied by detailed arguments.

2. What to measure and how

Many studies of the Internet look at the number of users (cf. [Nua]). This is the most relevant measure for some purposes, although it is inadequate for others, as it does not say anything about the intensity of usage. Other studies measure the Internet's size by the number of computers connected to it (cf. [ITU] and Table 7 in this paper, both based on data from [NW]). In this study we focus on the sizes of networks and the traffic they carry.

Since the Internet is a loose collection of networks, it is hard to decide what to include in estimating its size. See Fig. 1 for a sketch of the universe of data networks and the role that the public Internet plays in it. A key point is that what is commonly thought of as the Internet, namely the public backbones connecting all the local networks, is only a small part of the data networking universe. Most data traffic, just like most voice phone traffic, is local. Also, most of the cost of data and voice communications is associated with local facilities. For example, universities tend to devote between 10% and 20% of their network budget (counting the cost of people as well as equipment and services obtained from carriers) to Internet connections. In voice telephony, of the approximately \$200 billion that is spent in the USA each year, only about \$80 billion is for long-distance (inter-LATA) calls. Moreover, of that \$80 billion, about \$30 billion is paid to the local carriers in access charges, so the true long-distance component of the cost to the public is around \$50 billion, only about a quarter of the total for the voice

system.

Although most of the cost of telecommunications networks is for local facilities, we consider long-distance transport only. The economics and patterns of usage of local facilities are different (see [Odlyzko2] for more extensive discussion of this point). Studies of data as well as voice communications have historically concentrated on long haul circuits. They are the most heavily utilized and also the most difficult or expensive to upgrade, so sophisticated engineering and pricing approaches are most appropriate there. We will follow this precedent, and will not consider the LANs and WANs. Similarly, we will not consider local access links, such as the phone lines used by residential Internet users to connect to their ISPs or to make local voice calls. While we cannot avoid comparing apples to oranges, we try not to compare apples to orange trees.

Unlike most studies of the Internet, we consider not only the public Internet, but also the long-distance private line networks. These networks are estimated to mostly use IP (Internet Protocol), and, as we show, are far larger in both cost and bandwidth (but not necessarily in traffic carried) than the public Internet. The evolution of the Internet in the next few years is likely to be determined by those private networks, especially by the rate at which they are replaced by VPNs (Virtual Private Networks) running over the public Internet. Thus it is important to understand how large they are and how they behave.

In this study, we consider only U.S. networks. These networks still account for between 60% and 70% of users and host computers in the world [ITU, Nua], and almost surely a higher fraction of capacity and traffic, since transmission costs much less in the U.S. than in most of the world [ITU]. It is also easier to obtain data for North America. Further, most countries are striving to lower their telecommunications costs, and so the patterns we observe in the USA are likely to be replicated elsewhere in the next few year. The geographical restriction does mean, however, that the equality of data and voice traffic we predict for 2002 applies only to the U.S., and the crossover is likely to be somewhat later in most other countries.

Even in the U.S., we exclude government networks from consideration. We also do not consider research networks such as vBNS, since they carry little traffic, although they do have substantial capacity (as we will mention).

We consider only lines that are used for carrying voice and data traffic. The capacities of the underlying fiber networks (cf. [FCC1]) are far higher than those we will be listing. It would take us too far afield to try to explain the disparity in the estimates, but they have to do with differences between air distances and fiber route distances, presence of dark fiber, restoration capacity, and other factors.

3. Costs and prices and the decline of distance

We study capacities of networks and the traffic they carry. We measure these in Gbps (gigabits per second, 10^9 bits per second) and TB/month (terabytes per month, 10^{12} bytes per month). However, it seems intuitive that a terabyte carried between Baltimore and Philadelphia is not equivalent to a terabyte between Baltimore and San Francisco. Thus a more complete description of communications traffic should incorporate a measure of how far that traffic travels. Distance does play an important role in the evolution of networks. For example, an important reason cited for the migration of corporate data traffic to the public Frame Relay networks is that

charging for Frame Relay is insensitive to distance, making it much less expensive for long distance communications [Cavanagh].

While distance does play a role in telecommunications, it is a decreasing role. (See the book [Cairncross] for general discussions of this phenomenon, and its likely effects.) The monthly tariffs for interstate T1 leased lines from MCI (quoted from [Leida]) consist of a fixed fee of \$3,234 and \$3.87 for each mile. (The corresponding figures for a T3 are \$22,236 and \$52.07, respectively, so the distance dependence is stronger for higher capacity links, a phenomenon we expect to continue.) For the typical 300-mile leased circuit distance, the fixed fee is 73.6% of the cost of a T1, and 58.7% for a T3, and a more representative cost estimate would show much greater share for fixed costs, since it would include local connections and lease of terminating equipment. (The figures in Table 3 are for all-inclusive costs of private lines of various speeds, and for a 56 Kbps line include 57% for local access costs. It should be noted, as is detailed in [Leida], for example, that customers typically lower their costs by up to 50% through long-term contracts and bulk purchase discounts.) The historical trend in pricing of a T1 connection is shown in Fig. 2, which displays the tariffed rates for a T1 line of 700 miles, broken down into fixed and mileage-sensitive components, from the time this service was first offered until the end of 1997. Even at such a large distance, the distance-sensitive part of the price has decreased rapidly.

The decreasing role of distance can also be seen in the voice telephony price structure. As is shown in Table 14.2 of [FCC2], in the early 1980s there were many rates for interstate calls, depending on both distance and time of day. Today there is still some variation depending on time of day, even though that is also decreasing (see [Odlyzko2] for a discussion), but calls are priced independently of distance (within the U.S.). This reflects the decreasing fraction of the price that goes to cover network costs. It is estimated that of the 12 cents a minute that carriers collect on average for a voice call within the U.S., only about 1.5 cents is needed to pay for the network. (By far the largest component of cost is the approximately 5 cents a minute of access charges paid to local carriers largely to subsidize local service.)

The discussion above was intended to show that while our analysis does ignore important factors of distance, it is reasonable to do so as a first approximation. In further defense of our approach, let us mention that most communication is still local. This was apparently first noted in the 19th century by Carey and others (see Chapter 11 of [Isard]), but is best known from the work of Zipf [Zipf], who collected a variety of statistics on communication and transportation patterns. Zipf observed that whether one measured phone calls, car travel, or mail usage, the interaction between two cities with populations A and B at distance D appeared to be proportional to $A * B / D^\alpha$, with $\alpha = 1$. Other investigators since that time have found better fit for other values of α , typically with $1 \leq \alpha \leq 2$. Even services with distance-insensitive fees, such as mail, appear to be closely tied to social interactions, and are mostly local. That is certainly the case for usage of the telephone network. Interstate voice calls on average go over 500 air miles, while private lines are on average about 300 air miles in length.

A fascinating topic for further research is whether Zipf's observations will apply to the Internet. In voice telephony, the last 20 years have seen a growth in interstate and intrastate toll calls from 8 minutes per line per day to 14 minutes, while local calling has stayed about constant at 40 minutes (Table 12.2 of [FCC2]). This was presumably caused by the decline in long-distance prices and the greater mobility of the population. Still, most calls are local, and even in the interstate calling case, the intensity of calls drops off with distance, as Zipf ob-

served. The Internet is a world-wide network, and much traffic comes from downloading from popular Web servers, many of which appear to be located in California. On the other hand, a disproportionate share of Internet traffic is within California in any event (as is seen by examining the backbone maps in [Boardwatch]), so it could be that most traffic is local even on the Internet. Even if that is not the case now, the penetration of the Internet into everyday life may mean that Internet traffic will again follow patterns of our everyday social and economic interactions, and be largely local in the future. Further, even if there is no trend towards local information sources, the spread of caching may mean that most packets will be transported over short distances. Additional investigation is clearly desirable, especially since the speed with which it is economical to deploy some novel transport technologies depends on the distances over which traffic is to be carried.

4. Units of measurement

It will be convenient to state some conversion factors between different units and between the bandwidth of a connection and the traffic carried by that connection. Since there is substantial uncertainty about many estimates, we will not attempt to achieve precision, and will often not worry about 10% differences.

Voice on phone networks is carried in digitized form at 64,000 bits per second. Each voice call occupies two channels, one in each direction, so takes up 128,000 bps of network bandwidth. Thus one minute of a voice call takes $60 \times 128 \times 1000$ bits, or 937.5 KB (kilobytes, units of $1024 = 2^{10}$ bytes). Rounding this off, we get

1 minute of switched voice traffic \approx 1 MB.

(There is a discrepancy between the meaning of the “*k*” or “*K*” prefixes, which commonly denote 1000 in communication and $1024 = 2^{10}$ in computing. Given the lack of precision in most of the estimates we will be dealing with, this difference will be immaterial and will be ignored.) Compression can reduce that to a much smaller figure, and is used to some extent on high-cost international circuits, as well as on some corporate private line networks. As far as the network is concerned, though, it is carrying almost 1 MB of digital data for each minute of a voice call. Further, most data traffic can also be compressed, so we will ignore this factor.

A T3 (or DS3) line operates at 45 Mbps (actually, closer to 43 Mbps, but again we won’t worry about this discrepancy) in each direction, so that if it were fully loaded, it would carry 90 Mbps. Over a full month of 30 days, that comes to 29 TB (terabytes, which are 10^{12} bytes). We will say that

full capacity of a T3 link \approx 30 TB/month.

A T1 line (1.5 Mbps) is 1/28-th of a T3, and we will say that

full capacity of a T1 link \approx 1 TB/month.

5. Voice networks

The FCC collects and publishes comprehensive statistics on long-distance switched voice networks. They show that at the end of 1997, U.S. carriers had about 40 billion voice minutes

per month of interstate traffic on their public networks (Table 11.1 of [FCC2]), which is about 40,000 TB/month. This number has recently been growing at about 8% per year (Table 12.1 of [FCC2] and [FCC3]). Including local and intrastate toll traffic boosts the estimate to about 275,000 TB/month. (Table 12.1 of [FCC2] also shows that since 1980, intrastate and interstate toll calls have grown from 7% and 8%, respectively, of switched voice minutes, to 11% and 15%, another sign of the declining role of distance.)

The 40,000 TB/month figure for long distance switched network traffic includes a large but unknown fraction of fax and modem calls, which carry data. However, since they appear on the network as switched calls, they will be counted as voice.

Defining long distance traffic is easy compared to defining what is meant by long distance network capacity. There are various special connections for operator services, 800 number services, and the like. We consider just the long distance lines between large switches. Then known distributions of traffic over a week, achievable busy hour utilizations, and reserve capacity rules of thumb (all described in the literature, for example in Figure 1.13, Table 1.8, and discussion on p. 49 of [Ash]) show that the average utilization of such links is around 33%. Combined with the traffic estimates above, this shows that all the switched voice networks in the U.S. had capacity of around 350 Gbps at the end of 1997.

6. The public Internet

There are many estimates of the size and growth rate of the Internet that are either implausible, or inconsistent, or even clearly wrong. We already cited the Sidgmore interview [Sidgmore] as an example. In a somewhat different vein, the article [Gareiss] claimed that one third of Internet traffic went through the MAE East peering point. While it is probably true that at the time [Gareiss] was written, about one third of the traffic that went through public peering points went through MAE East, this traffic was only a part of total Internet backbone traffic.

The major reason for the uncertainties in measuring the Internet is that carriers do not release detailed information about their networks. As a result, any estimates made from publicly available data will necessarily have a large error margin.

As a first step, to provide a "sanity check" on other estimates, we consider the traffic generated by residential users accessing the Internet with a modem. There are about 20 million of them (or, more precisely, there are about 20 million active accounts) and according to the latest information from America Online and other services, on average an account is connected about 25 hours each month. These users download data at a rate of about 5 Kbps when they are online (with considerably smaller average upload rates), which generates traffic load of just about 1,000 TB/month. (To the voice phone network, which dedicates 128 Kbps for each connection, the load appears as 26,000 TB/month, about 10% of the total load of voice calls, local and long distance.) Since there are more PCs in corporate environments than at home, we should expect to see total Internet traffic of at least twice that, or 2,000 TB/month.

We next consider traffic through the public peering points. Statistics for them are available through the links at [CAIDA, NLANR], often going back a year or more. The 5 largest ones are shown in Table 1. The traffic estimates are for the early part of December 1997 (to avoid the Christmas and New Year holiday effects). The other public peering points are much smaller. Total traffic through all the public peering points is dominated by that through the 5 points

in Table 1, and comes to about 4 Gbps, or 1,200 TB/month. For comparison, in mid-1996, traffic through these points was about 1.6 Gbps, or 500 TB/month. Growth has been uneven, with especially rapid increase in traffic at the Chicago NAP. That peering point had traffic of only around 0.2 Gbps as late as October 1997, but by April 1998 was carrying about 0.7 Gbps. Overall, though, aggregate traffic through the NAPs and MAEs appears to have been growing at about 100% per year from late 1996 through April 1998. This agrees with the 100% growth rates for 1997 for MCI and an unnamed ISP that are cited in [Schulzrinne].

We assume that little traffic goes through more than a single peering point. That assumption appears reasonable, especially in view of the congestion at the NAPs and MAEs. What we do not know how to evaluate with high confidence is the fraction of Internet traffic that goes through the public peering points. A substantial part of backbone traffic, which we estimate to be about 50%, stays within a single ISP. (This estimate could easily be too high, and is almost surely too high for traffic from residential users. On the other hand, corporations are increasingly moving their internal traffic to the Internet, and they appear to try to stay with a single provider. Also, many hosting services have connections to several ISPs, and so much of their traffic does not have to go through exchange points.) Even of the traffic crossing ISP boundaries, some fraction, which we estimate to be 40%, goes through a private peering point. With these assumptions, we find that about 30% of all backbone traffic goes through the NAPs and MAEs. This leads to an estimate of backbone traffic of 13 Gbps, or 4,000 TB/month at the end of 1997. By comparison, in mid-1996, the estimate was 5.7 Gbps, or 1,800 TB/month.

Another way to measure the Internet is to look at particular carriers. MCI is regarded as carrying between 20% and 30% of the backbone traffic. In mid-1996, MCI press releases claimed their network was carrying 250 TB/month. Around November 1997, Vint Cerf stated that the MCI network was carrying 140 TB/week (and was growing about 100% per year) [Schulzrinne]. In December 1997, the MCI "white paper" [MCI] said that MCI was carrying 170 TB/week. A weekly load of 170 TB/week corresponds to about 740 TB/month, and depending on whether we assume the 20% or 30% figure for the MCI share of backbone traffic, we obtain an estimate of between 2,500 and 3,700 TB/month for all the backbones.

To determine the size of the Internet, we examined all of the major NSPs (National Service Providers, the large carriers with backbones extending across the U.S.) and obtained a bandwidth capacity for each of these carriers. Typically, the actual networks are composed of combinations of T1s, T3s, OC3s, and OC12s. For consistency we chose to express the bandwidth capacity in terms of equivalent T3s. We realize that a more appropriate metric would be in terms of circuit miles. As mentioned earlier in this document a shorter 100 mile T3 link is much different than a 700 mile T3. In fact several of the smaller carriers have a large number of short T3 links (as compared to say MCI and the other large NSPs). In many cases, it appears that the average lengths of the T3s in the smaller networks are between 1/3 and 1/5 of that for the larger NSPs. Nevertheless, we simply counted the number of equivalent T3s, and compared.

The data for Internet backbones was obtained from a variety of sources. There are extensive listings for many NSPs at [CAIDA], and we relied on those to a large extent. However, we often had to adjust the data there. For example, the listing for UUNet in the [CAIDA] files in April 1998 (when we completed our study) listed the equivalent of about 380 T3s. A count of the UUNet backbone links on the UUNet map at [UUNet] showed about 480 T3s at that time, and if the map were current, that would represent the state of the UUNet network in April 1998.

Since the 4 months between December 1997 and April 1998 represents about 25% growth when traffic doubles each year (as it appears to be doing currently), we assumed that UUNet had the equivalent of about 400 T3s in their network at the end of 1997.

In mid-1996 we estimated the total number of equivalent backbone T3s to be around 400. (This may have been an underestimate, and the value was probably closer to 500.) At that time we obtained detailed information on the key carriers, specifically: AGIS, BBN, PSINET, UUNet, MCI, Sprint, ANS, and AT&T. These eight accounted for about 75% of the total number. For December 1997, we examined 35 commercial backbone providers (such as internetMCI and UUNet), and 8 research networks (such as MAGIC and vBNS). We found about 2,100 equivalent T3s in the commercial networks and about 500 equivalent T3s in the research networks, giving a total of 2,600 equivalent DS3s. This was the estimate based on data available in April 1998. Much of that information was several months old. On the other hand, we were trying to measure the Internet at the end of 1997, at which point it was probably some 20% smaller. Thus we assume that the two sources of error cancel each other out, and that the commercial Internet backbones had the equivalent of about 2,100 T3s in December 1997.

Table 2 gives some of equivalent T3 counts for several of the major NSPs (both commercial and R&D). We note that 2,100 T3s represents a 100-fold increase from the 20 or so in the NSFNet backbone at the end of 1994. Traffic of 2,500 to 4,000 TB/month at the end of 1997 represents more than a 100-fold increase from the 15 TB/month carried by NSFNet at its peak.

Our estimates show considerably greater growth in the sizes of backbones than in traffic through the NAPs and MAEs between mid-1996 and end of 1997. This could be caused by more traffic bypassing the public peering points. It could also be that the rapid growth in traffic during 1995 and 1996 led NSPs to project faster growth than materialized in 1997. (It takes close to a year to obtain a T3 private line, since capacity is short at present.) Another possibility is that the NSPs were responding to complaints about congestion and decreasing utilization rates of their networks. There are also stories that they may have put in more capacity, especially of very fast links, than was absolutely required in order to win battles for business customers, for whom ability to burst to high speeds and low utilization rates are attractions.

A network of 2,100 T3s has a bandwidth of 190 Gbps. However, that is the total bandwidth of all the links, and could only be utilized fully if every packet went from the node where it enters the backbone to an adjacent node and exited the network right at that point. There are studies (see [Schulzrinne] for links) which show that the average number of hops that a packet makes on the Internet is around 15. However, those studies measure all the hops a packet makes, and most of them are in the access part of the network. We are looking only at the backbones. There do not appear to be any comprehensive studies of how many backbone hops a packet makes. Statistics for the NSFNet backbone (available at [NSFNet]) show that towards the end of its existence, in late 1994, its T3 links were running around 5% of capacity. Next, NSFNet data show transport of about 15 TB/month at that time. Since there were 19 T3s in service, a 5% utilization rate should have led to total traffic of 28.5 TB/month on all the T3s. This is consistent with NSFNET moving 15 TB/month if each packet on average traveled over two T3s.

For the current public Internet, an average of 2.5 hops on the backbone links per packet appears reasonable. Experts we have consulted thought it was about right, and it is also the figure used in [Leida]. Further, it agrees with some recent routing data provided by Ramon Caceres (private communication) for various Internet connections in the U.S.. If we assume

this estimate of 2.5 hops per packet, the effective bandwidth of the public Internet becomes 75 Gbps.

7. Private line networks

Many private lines are leased by one carrier from another to carry voice. Others carry Internet traffic or the traffic on the semi-public Frame Relay and ATM networks, or else government traffic. If we consider just the so-called retail market, in which lines are leased to private non-carrier organizations, then industry statistics (collected by the Vertical Systems Group, [VS]) give the estimates of Table 4. There are several noteworthy features of these numbers. One is that the bulk of the bandwidth will soon be in T3 lines, yet these lines bring in only 7% of the industry revenue. This explains how bandwidth has been exploding while leased line revenues have been increasing at modest rates.

Adding up the bandwidths in Table 4, we obtain an estimate of total bandwidth of the private line networks of 630 Gbps, substantially more than the bandwidth of the voice network. However, as for the Internet, we have to consider the effective bandwidth. Users care just about getting a message or a file from point A to point B, and not how it gets there. What we don't know is how many private line links a typical message in a corporate networks traverses. In the early days of data networking, most corporate networks appear to have been star-shaped, with branch locations communicating with a central facility. In those networks, usually just a single hop is involved in a message. However, recent years have seen development of mesh networks, especially among the large corporations that are the primary customers for the T1 and T3 lines that contain the bulk of the bandwidth. For those networks, it appears reasonable to assume that on average a message will make two hops. (Even in star networks, some fraction of the traffic is between points on the periphery, which again requires two hops.) With that two-hop assumption (justified by data from two large corporate networks) the effective bandwidth of the private line networks reduces to about 330 Gbps.

Just as the voice network carries substantial but unknown proportion of data calls, private line networks are not all dedicated to data. Large organizations often use them to transmit voice calls, especially on international links. However, the general opinion seems to be that although at one point this was the main use of private lines, today it is a minor factor. We will therefore ignore it.

Traffic on private line networks is much harder to estimate than their capacity. The key point of the companion paper [Odlyzko1] is that conventional capacity utilization estimates, such as those of [Leida], are almost an order of magnitude too high. It is impossible to obtain precise estimates, since no measurements are taken on many lines, and even when there are statistics, those are not released. However, both direct and circumstantial evidence is presented in [Odlyzko1] for the claim that private lines are used at 3% to 5% of their capacity, when averaged over a full week. These estimates produce estimates of traffic on private line networks between 3,000 and 5,000 TB/month.

Finally, we should mention the semi-public Frame Relay and ATM networks. No firm figures are available, but industry estimates (partially based on data in [VS]) suggest the capacity and traffic estimates in Tables 5 and 6.

8. Rates and sources of growth

Voice traffic is currently growing about 8% per year. Private line capacity, and presumably also traffic, is growing at about 15% to 20% per year. The semi-public Frame Relay and ATM networks are growing about 100% a year. (It is noteworthy that although Frame Relay is sometimes claimed to be cannibalizing the leased line business, it is not doing it fast enough yet to affect the growth of the latter. Most of the growth in Frame Relay appears to be coming from new applications.) Internet traffic on the NSFNet backbone was doubling each year between early 1991 and the end of 1994, when the impending phaseout of NSF support led traffic to shift to private backbones. (The growth in NSFNet traffic is shown in Table 8. The Dec. 1990 entry in that table is extrapolated from data starting in March 1991. Otherwise all the numbers through 1994 are taken from [NSFNet].)

We do not have precise traffic statistics for NSFNet before 1991. However, the project overview (available at [NSFNet]) does mention that the number of packets transmitted increased by a factor of 62.5 in the 30 months between July 1988 and January 1992, for an annual growth rate 400%. Internet host counts (see Table 7, based on statistics at [NW]) show slower growth, with regular doubling each year throughout the 1980s and 1990s. Host measurements are unreliable and hard to interpret, but in general we might expect their growth rates to be at least slightly indicative of those of network traffic. An extensive study of data traffic through 1993 by Paxson [Paxson] found many volume measures showing about a doubling each year.

The rapid growth spurt in traffic in 1995 and 1996 was presumably caused by several closely related phenomena. One was that the Internet caught the public's attention, with millions of new users signing up for home accounts or getting access at work. The other was that user-friendly Web browsers made the Internet more valuable even than it would have been otherwise. (General growth of the Internet can be ascribed to Metcalfe's Law, which says that the value of a network is proportional to the square of the number of users.) Also, the graphical user interface of the popular Web browsers led to the creation and use of illustrations, which consume far more bandwidth than text. (Some projections for amount of data traffic that in retrospect can be seen to have been far too conservative, such as that on pp. 171-175 of [Noll], erred primarily by not anticipating the growth of rich graphical content.) What is remarkable, though, is that in some networks, the Web did not appear to have a noticeable effect on the growth rate of traffic. For example, the Swiss SWITCH network for academic and research and research institutions has seen growth in IP traffic by a factor of about 2.5 each year in the 1990s ([Harms] and private communication with J. Harms covering the period since 1994).

As we explain in Section 6, most indications are that after the anomalous period of 1995 and 1996, when Internet backbone traffic increased by a factor of 100, it appears that it has slowed down to a growth rate of about 100% a year. The big question is whether this growth rate can be sustained for long. We feel that there is no reason to expect a slowdown in the next decade, and there could even be periods of more rapid growth.

The 100% annual growth rates for the Internet were a result of increase in number of users as well as increased traffic from existing users. While there will surely continue to be growth in users of the Internet, it will not be by a factor of 100 or even close to it over the next decade, since that would require more people than live on the Earth. However, rapid growth has been observed even from small communities. The Swiss SWITCH network was mentioned above. The University of Waterloo [Waterloo] also appears to have experienced roughly a doubling

of traffic to and from the Internet each year for the last five years. The statistics for their traffic are shown in Table 9. (Data is available only for the month in each school term that had the greatest traffic during that term.) The Library of Congress statistics (Table 10 below, based on data at [LOC]) also show more than a doubling each year. Nortel has seen 80% growth in its IP traffic volume for the last three years (private communication from Terry Curtis of Nortel). Press releases of presentations by Lew Platt of HP in 1996 and 1997 show that HP's IP traffic doubled during that year. Thus it appears that organizations find IP sufficiently attractive that they double their traffic each year, although the sources of demand may change.

Where will future growth come from? Unlike voice phone traffic, for which there is a natural limit, since people are not willing to spend all their time talking to others, there is no obvious bound on data traffic demand. For residential customers, there is a serious current limiting factor in the modems they have. However, even without deployment of new technologies, substantial further growth of the Internet can come from that source. America Online reports that its customers have tripled the time they stay connected, to 45 minutes per day, in the year and a half that their flat-rate pricing plan has been in effect. There is room for further expansion in this area. About half the households in the U.S. have PCs, and of these only half have residential Internet service. Further, 45 minutes per day is a small fraction of the time that American families watch TV. As more people join the network, they create more content, and make it more attractive for others to create content for the Internet, which draws in more users, and so on, as predicted by Metcalfe's Law.

Once broadband technologies such as cable modems, various DSL technologies, or wireless data links are deployed, residential customers will be able to substantially increase their traffic per household. In the very near future, though, rapid growth is most likely to come from institutions such as corporations, which have the broadband communication infrastructure to increase increased traffic. Costs of providing local access at T1 speeds to the Internet from corporate WANs are dropping rapidly with the deployment of HDSL (although that is not reflected in prices yet), and fiber will be increasingly feasible for higher speeds. Even without novel applications, such as packet telephony and videoconferencing, ordinary Intranet and Extranet applications could lead to continued growth at historic rates. We could also see spurts of growth even faster than 100% per year if some of the novel applications start growing. The public Internet could also grow at high rates if more of the growth in internal private line networks shifts to it.

It is worth remarking that packet telephony may cause a spurt in Internet traffic. However, this is likely to be just one of many spurts powering the growth of the Internet, just as streaming audio, for example, is doing today. The reason not to expect packet telephony to be a gigantic influence is that there is not that much voice traffic to lead to a major change in Internet traffic statistics. This may seem paradoxical in view of the evidence we have presented that there is much more voice than data traffic. However, packetization of voice offers natural opportunities for compression. The switched voice network devotes 128 Kbps to each conversation, whereas decent quality can be obtained with rates of 8 Kbps. Even if we do not pursue the most aggressive compression schemes, say in order to keep latencies low, and digitize voice calls at 32 Kbps, the 40,000 TB/month of voice traffic becomes 10,000 TB/month, a level that at present rates of growth the Internet is likely to reach in less than two years. Thus the fears about lack of bandwidth for packet voice transmission, such as those in [GMLCOBRS], appear to be unwarranted.

9. Conclusions

We have shown that in the U.S., traffic on the public Internet is under 10% and bandwidth is around 20% of the switched long distance network. On the other hand, the bandwidth of the private line networks is already comparable to that of the voice network, although traffic on them is probably not much higher than on the public Internet. Also, the Internet appears to be growing at 100% per year, compared to 15-20% for private line networks and under 10% for the voice network. Thus if current trends continue, and there seems to be no reason they should not, data traffic will overtake voice traffic around the year 2002, and will be going primarily over the public Internet.

The 100% annual growth rate of the Internet forces new ways of thinking about telecommunications. As one simple example, the estimates we made are primarily for December 1997. However, this paper was written four months later, at which point all the estimates for the Internet in the tables already have to be increased by about 25%.

Acknowledgements: We thank Vijay Bhagavath, Ehud Gelblum, Jacek Kowalski, Clem McCalla, Roger Watt, and Bill Woodcock for comments and enlightening information.

References

- [Ash] G. R. Ash, *Dynamic Routing in Telecommunications Networks*, McGraw Hill, 1998.
- [Boardwatch] *Boardwatch* magazine, <http://www.boardwatch.com>.
- [CAIDA] Cooperative Association for Internet Data Analysis (CAIDA), <http://www.caida.org/>.
- [Cairncross] F. Cairncross, *The Death of Distance: How the Communications Revolution will Change our Lives*, Harvard Business School Press, 1997.
- [Cavanagh] J. P. Cavanagh, *Frame Relay Applications: Business and Technical Case Studies*, Morgan Kaufman, 1998.
- [Cerf] V. Cerf, PowerPoint slides of presentations, available at <http://www.mci.com/aboutyou/interests/technology/ontech/powerpoint.shtml>.
- [DOC] U.S. Department of Commerce, *The Emerging Digital Economy*, April 1998. Available from <http://www.ecommerce.gov/emerging.htm#>.
- [FCC1] J. M. Kraushaar, *Fiber deployment update: End of year 1996*, U.S. Federal Communications Commission, 1997. Available from <http://www.fcc.gov/ccb/stats>.
- [FCC2] U.S. Federal Communications Commission, *Trends in Telephone Service*, Feb. 1998. Available from <http://www.fcc.gov/ccb/stats>.
- [FCC3] J. Zolnierok, K. Rangos, and J. Eisner, *Long Distance Market Shares; First Quarter 1998*, U.S. Federal Communications Commission, June 1998. Available from http://www.fcc.gov/Bureaus/Common_Carrier/Reports/FCC-State_Link/comp.html.
- [Gareiss] R. Gareiss, Is the Internet in Trouble?, *Data Communications*, Sept. 21, 1997. Available at <http://www.data.com/roundups/trouble.html>.
- [GMLCOBRS] V. Granger, C. McFadden, M. Lambert, S. Carrington, J. Oliver, N. Barton, D. Reingold, and K. Still, Net benefits: The Internet - A real or virtual threat, Merrill Lynch report, March 4, 1998.
- [Harms] J. Harms, From SWITCH to SWITCH* - extrapolating from a case study, *Proc. INET'94*, pp. 341-1 to 341-6, available at <http://info.isoc.org/isoc/whatis/conferences/inet/94/papers/index.html>.
- [Inktomi] Inktomi white paper, *Traffic server: Large scale network caches provide more bandwidth for your money*, 1997. Available at <http://www.inktomi.com/Tech/EconOfLargeScaleCache.html>.
- [Isard] W. Isard, *Methods of Regional Analysis*, MIT Press, 1960.
- [ITU] International Telecommunication Union, *Challenges to the Network: Telecommunications and the Internet*, Sept. 1997. Can be purchased through <http://www.itu.ch>.

- [Leida] B. Leida, A cost model of Internet service providers: Implications for Internet telephony and yield management, M.S. thesis, department of Electr. Eng. and Comp. Sci. and Technology and Policy Program, MIT, 1998. Available at <http://www.nmis.org/AboutNMIS/Team/BrettL/contents.html>.
- [LOC] Library of Congress, data traffic statistics, available at <http://lcweb.loc.gov/stats/>.
- [MCI] *Quality counts*, MCI white paper, available at <ftp://ftp.mci.net/pub/Quality.pdf>.
- [MutooniT] P. Mutooni and D. Tennenhouse, Modeling the communication network's transition to a data-centric model, available at <http://www.sds.lcs.mit.edu/~mutooni/telecomms/>.
- [NLANR] National Laboratory for Applied Network Research, <http://www.nlanr.net/>.
- [Noll] A. Michael Noll, *Introduction to Telephones and Telephone Traffic*, 2nd ed., Artech House, 1991.
- [NSFNet] Historical data for NSFNet, available at <http://www.merit.edu/nsfnet/>.
- [NW] Network Wizards, <http://www.nw.com>.
- [Nua] Nua, Ltd., resources available at <http://www.nua.ie/>.
- [Odlyzko1] A. M. Odlyzko, Data networks are lightly utilized, and will stay that way. Available at <http://www.research.att.com/~amo>.
- [Odlyzko2] A. M. Odlyzko, The economics of the Internet: Utility, utilization, pricing, and Quality of Service. Available at <http://www.research.att.com/~amo>.
- [Paxson] V. Paxson, Growth trends in wide-area TCP connections, *IEEE Network*, 8 (no. 4) (July 1994), pp. 8-17. Available at <http://www-nrg.ee.lbl.gov/nrg-papers.html>.
- [Schulzrinne] H. Schulzrinne, traffic statistics collected at <http://www.cs.columbia.edu/~hgs/internet/traffic.html>.
- [Sidgmore] John Sidgmore, interviewed by R. L. Brandt, *Upside* 10 (no. 5) (May 1998), pp. 78ff. Available at <http://www.upside.com/tehis/mvm/story?id=3512c4520>.
- [Thyfault] M. E. Thyfault, Resurgence of convergence, *Information Week*, April 13, 1998, pp. 50ff.
- [UUNet] UUNet backbone map, available at <http://www.us.uu.net/>.
- [VS] Vertical Systems Group, *ATM & Frame Relay Industry Update*, 1997.
- [Waterloo] University of Waterloo network statistics, available at <http://www.ist.uwaterloo.ca/cn/#Stats>.
- [Zipf] G. K. Zipf, Some determinants of the circulation of information, *Amer. J. Psychology* 59 (1946), 401-421.

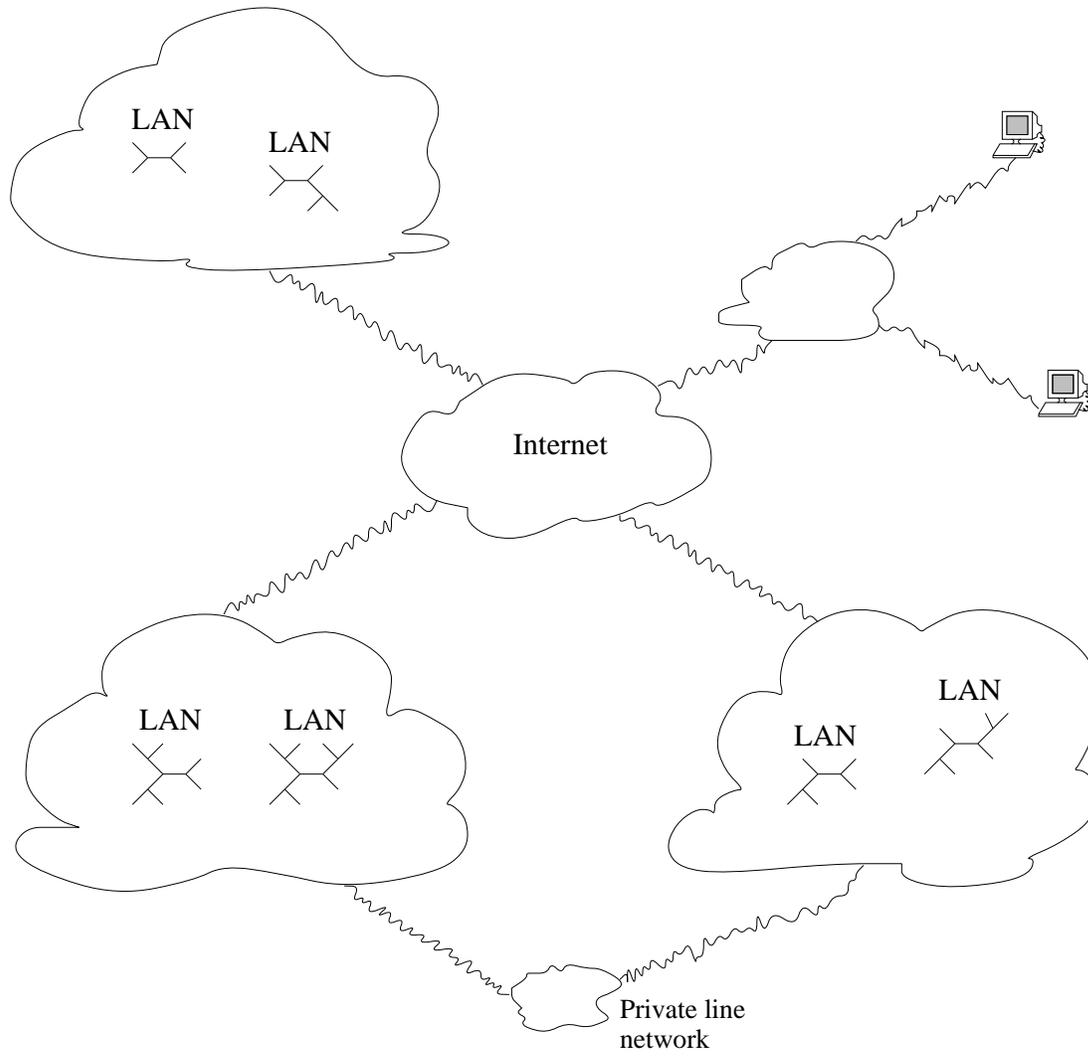


Figure 1: The Internet and private line networks. Artwork by Sue Pope and Thelma Pickell.

monthly lease of a 700-mile T1 line

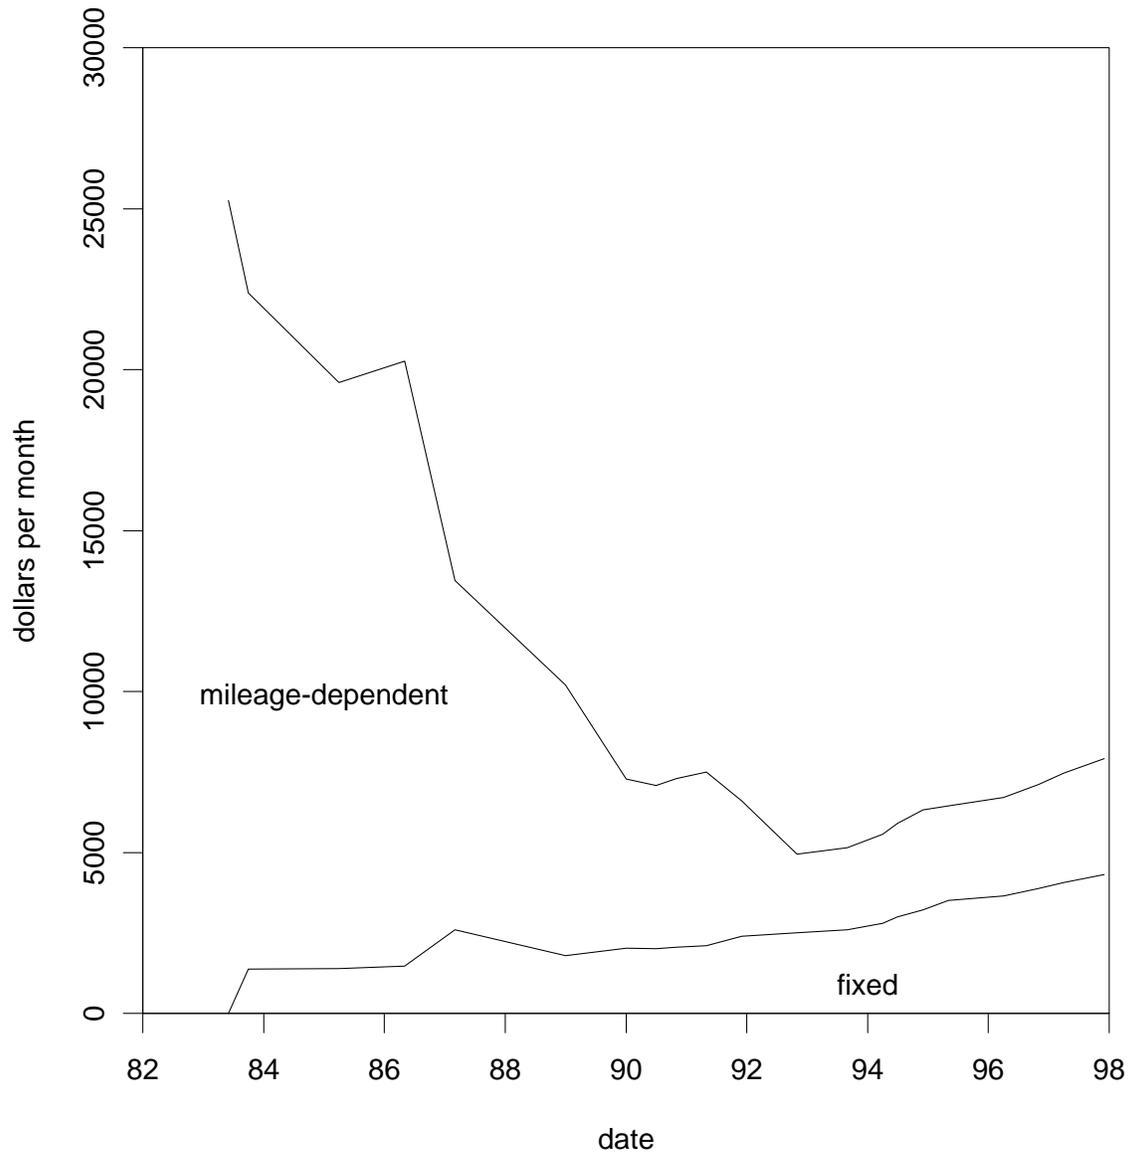


Figure 2: The decreasing role of distance. Top curve shows the total cost of the monthly lease on a T1 line, the bottom curve the distance-insensitive part of the cost (all in current dollars, without adjusting for inflation).

Table 1: Major public exchange point traffic, end of year 1997.

<u>peering point</u>	<u>traffic in Gbps</u>
Sprint NAP (New York City)	0.5
Ameritech NAP (Chicago)	0.7
Pac Bell NAP (San Francisco)	0.5
MAE East (Washington, DC)	1.1
MAE West (San Jose)	1.1

Table 2: National Internet backbones, in T3 equivalents.

<u>network</u>	<u>mid-1996</u>	<u>year-end 1997</u>
MCI	75	400
UUNet	< 50	400
BBN	30	52
AGIS	35	61
PSINET	20	51
Sprint	50-70	137
MAGIC		86
vBNS		255

Table 3: Leased line prices (300 miles long distance, 5 miles local).

<u>speed</u>	<u>price per month</u>
9.6 Kbps	\$1,150
56 Kbps	\$1,300
1.5 Mbps (T1)	\$7,000
45 Mbps (T3)	\$66,000

Table 4: Retail leased line market in the U.S., end of year 1997. Bandwidth in Gbps, revenue in billions of dollars.

line speed	no. lines	bandwidth Gbps	projected 98 growth	revenue \$ B
56/64 & lower	447,530	57	-1%	4.87
fractional T1	19,880	10	2%	0.26
T1	98,850	304	7%	4.58
T3 & higher	3,010	259	34%	0.72

Table 5: Effective bandwidth of long distance networks, year-end 1997.

network	bandwidth (Gbps)
US voice	350
Internet	75
other public data networks	40
private line	330

Table 6: Traffic on long distance networks, year-end 1997.

network	traffic (TB/month)
US voice	40,000
Internet	2,500 - 4,000
other public data networks	500
private line	3,000 - 5,000

Table 7: Growth in number of Internet hosts.

Date	Hosts
08/81	213
05/82	235
08/83	562
10/84	1,024
10/85	1,961
11/86	5,089
12/87	28,174
10/88	56,000
10/89	159,000
10/90	313,000
10/91	617,000
10/92	1,136,000
10/93	2,056,000
10/94	3,864,000
01/96	9,472,000
01/97	16,146,000
01/98	29,670,000

Table 8: Growth in traffic on Internet backbones. For each year, shows estimated traffic in terabytes during December of that year.

year	TB/month
1990	1.0
1991	2.0
1992	4.4
1993	8.3
1994	16.3
1995	?
1996	1500
1997	3000

Table 9: Growth in traffic between the University of Waterloo and the Internet. For each school terms, shows volume of data transmitted during the month of highest traffic. By permission of University of Waterloo.

month	MB/day in	MB/day out
93-03	453	227
93-07	541	271
93-10	648	324
94-04	977	543
94-08	1,377	915
94-11	2,098	1,426
95-04	2,285	1,730
95-07	3,229	2,588
95-11	6,000	3,450
96-03	7,229	4,275
96-07	7,823	4,572
96-11	10,962	5,984
97-03	11,463	6,235
97-07	12,087	7,223
97-11	24,621	10,572
98-03	24,676	9,502

Table 10: Growth in data traffic at Library of Congress. For each year, shows total traffic in gigabytes during February of that year.

year	GB/month	increase from previous year
1995	14.0	
1996	31.2	123%
1997	109.4	251%
1998	282.0	158%