

Internet traffic growth: Sources and implications

Andrew M. Odlyzko

University of Minnesota, Minneapolis, MN, USA

ABSTRACT

The high tech bubble was inflated by myths of astronomical Internet traffic growth rates. Yet although these myths were false, Internet traffic was increasing very rapidly, close to doubling each year since 1997. Moreover, it continues growing close to this rate. This rapid growth reflects a poorly understood combination of many feedback loops operating on different time scales. Evidence about past and current growth rates and their sources is presented, together with speculations about the future. The expected rapid but not astronomical growth of Internet traffic is likely to have important implications for networking technologies that are deployed and for industry structure. Backbone transport is likely to remain a commodity and be provided as a single high quality service. It is probable that backbone revenues will stay low, as the complexity, cost, and revenue and profit opportunities continue to migrate towards the edges of the network.

Keywords: Internet traffic growth, network economics, telecom industry structure, QoS

1. INTRODUCTION

The telecom crash and current depression were the result of the “irrational exuberance” of the late 1990s. Technology did meet the demands posed on it by business plans. The problem was that those business plans had been formed in willful ignorance of actual demand. The Internet simply did not grow as fast as had been predicted. Because of a misunderstanding of what customers wanted, the whole industry crashed, in spite of its technical excellence and plentiful capital expenditure.

This paper takes a high level view of the Internet, considering its economics and the needs it serves. Customers have no interest in the separation of TCP from IP, whether the Reno or the Tahoe variant of TCP is used, whether ECN is employed, and so on. Those are all very important technical questions, but users almost universally do not wish to be bothered with them. They care only about applications. The main questions are what kinds of transmissions are they interested in, and how much traffic are they likely to generate. Those are the questions this paper addresses. The answers are tentative, representing lack of precise data on many key factors. Still, they help explain some of the mistakes that have been and continue to be made, and suggest what technologies are likely to be adopted, and how the Internet business structure might evolve.

The discussion in this paper as well as most of the data will be limited largely to the United States. However, much of what is said about traffic growth, relative distribution of costs, or implications for industry structure, apply elsewhere as well.

Internet traffic continues to grow vigorously, approximately doubling each year, as it has done every year since 1997. (By a doubling I mean annual growth between 70 and 150%.) Table 1 presents my estimates for Internet backbone traffic in the U.S. over the last decade. This table extends the estimates in my previous papers with Kerry Coffman⁴⁻⁶ and fits very well the predictions of those papers. Section 2 discusses the methodology used for these predictions, and compares them to other estimates. Traffic growth appears to have declined recently, but not dramatically.

Section 3 is devoted to an exploration of the myth of “Internet traffic doubling every 100 days,” primarily its origins and significance. Section 4 is devoted to the issue of data network utilization. The generally light utilization of data links is partially a reflection of the high growth rates in traffic. For the most part, though, it reflects the desire for low transaction latency that is the main driving force behind deployment of data networks.

Further author information: Email: odlyzko@umn.edu, URL: <http://www.dtc.umn.edu/~odlyzko>, Address: Digital Technology Center, 499 Walter Library, 117 Pleasant St. SE, Minneapolis, MN 55455, USA

Table 1. Traffic on Internet backbones in U.S.. 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	1,500
1997	2,500 - 4,000
1998	5,000 - 8,000
1999	10,000 - 16,000
2000	20,000 - 35,000
2001	40,000 - 70,000
2002	80,000 - 140,000

Section 5 discusses the significance of utilization levels and traffic growth rates for the future of the Internet. They might lead to more volatility in network equipment spending. Should growth slow down, different architectural principles might be appropriate. Section 6 considers Quality of Service (QoS) and how its appropriateness depends on growth rates of traffic. Section 7 is devoted to sources of data traffic growth, and the likely role of “killer apps” in continuing current growth trends.

Section 8 considers the Internet in relation to the entire telecom industry. While the bandwidth of the long distance links in the Internet is far higher than in the other (voice, Frame Relay, ATM, or private line) networks, Internet traffic surpassed voice traffic in volume only recently (most likely in 2002). Further, spending on Internet is still far lower than on voice. In particular, Internet backbones have relatively small revenues. Their costs are also low, and will remain so. Backbone transport is, and will likely remain, a commodity. Carriers will have to strive to increase traffic.

The Internet has accelerated an old trend in telecommunications, in which costs have been decreasing fastest in the core of the network. Section 9 demonstrates how inexpensive Internet backbones are, and how this reinforces many of the predictions from earlier sections about QoS, likely growth rates, and related questions.

Section 10 summarizes the discussion. Internet traffic continues to increase at a healthy rate, but this rate is nowhere near what would have been required to absorb the network capacity that was installed during the bubble. The turmoil in the industry is the result of a combination of gross overcapacity and a restructuring of the industry, in which the core of the network is being hallowed out.

2. SIZE AND GROWTH RATE OF THE INTERNET

Although the telecom industry has gone from boom to bust, Internet traffic growth appears to have slowed down only moderately. During the bubble years, the dominant mantra was of “Internet traffic doubling every three months.” How this myth started and propagated is discussed at some length in the next section. There was plenty of evidence that disproved the myth,^{4-6,30} but it was disregarded. Once the bubble burst, the consensus seemed to shift, and some public figures blamed the crash on an unexpected slowdown in traffic growth. Nortel’s Roth even claimed in mid-2001 that Internet traffic was declining (which would have excused Nortel’s plummeting sales), but quickly had to backtrack in response to vigorous denials from the industry.¹⁷

My estimates of Internet traffic on U.S. backbones are shown in Table 1. They extend the estimates made in the papers with Kerry Coffman⁴⁻⁶ and use the same methodology. We had observed that backbone traffic was approximately doubling each year (something we dubbed “Moore’s Law for data traffic⁵”), by which we meant growth rates between 70% and 150% per year. (Our methodology did not allow for much more accurate

Table 2. Traffic on the link from the University of Waterloo to the Internet. Based on sampling over one week in March of each year. In GB/day. The “other” category includes P2P.

year	total	Web	other	mail
1994	1.6	0.0	0.4	0.2
1995	3.7	0.9	0.7	0.3
1996	11.5	5.4	2.1	0.6
1997	18.1	11.2	3.6	0.7
1998	33.9	21.5	6.8	1.2
1999	65.7	37.2	17.9	1.8
2000	97.9	54.0	32.2	2.6
2001	129.0	65.3	48.9	4.5
2002	237.0	110.0	111.0	3.7
2003	362.0	173.0	175.0	5.3

estimates, for reasons sketched briefly below, and discussed in detail in our papers.⁴⁻⁶) Traffic continues growing at these rates, and recent declines have generally not been dramatic.

The data in Table 1 for 1990 through 1994 is taken from the trustworthy statistics for NSFNET, the original backbone funded by NSF. (It ignores the private backbones, whose share of traffic was growing, but, according to expert opinion, was small through 1994.) Data for later years are based on extrapolations from incomplete data. No government or industry body collected detailed statistics, and carriers were almost uniformly very secretive about their traffic. The studies of the papers⁴⁻⁶ were based primarily on monitoring publicly available traffic statistics for Internet exchanges and especially end users. This data was supplemented with occasional public announcements by some carriers, as well as with data provided under nondisclosure by both carriers and end customers. More recently I have extended this program to follow a greater variety of sources using automated data harvesting tools. However, the basic methodology is the same as in the papers,⁴⁻⁶ and many of the sources are the same. (For example, the University of Waterloo, which had been covered extensively in the earlier studies because of its long history of careful traffic monitoring and careful record keeping, has some data available on <http://www.ist.uwaterloo.ca/cn/#Stats>), which was used to prepare Table 2.) Hence I will not discuss the details here.

Although all governments in the past had taken hands-off attitudes towards Internet traffic, that appears to be changing. In particular, Australia’s government has recently started collecting and publishing statistics.¹ The Australian semi-annual statistics reports show that the volume of data received by business and residential customers in September of 2000, 2001, and 2002 was approximately 350 TB, 430 TB, and 785 TB, respectively. Combined with earlier data⁵ for traffic of Telstra, the dominant Australian carrier, this suggests that Australia experienced several years of regular doubling every year, then a remarkable slowdown in the growth rate, to about 22% in 2001, and then a resumption of almost-doubling, with growth of 83% in 2002.

There are several additional noteworthy aspects to the Australian data.¹ One is that, as a glance at Table 1 shows, the intensity of Internet traffic in Australia is far lower than in the U.S.. Australia has 14 times fewer people than the U.S., yet even when we adjust for this, we find about an order of magnitude less traffic per person than in North America. (In particular, this says that there is still far more voice traffic than Internet traffic in Australia. The same phenomenon is likely to hold in most other countries, if we consider the Internet bandwidth estimates that are available.⁴⁴)

There have been two recent attempts to obtain more systematic and more complete statistics by working primarily with carriers. Remarkably enough, both yielded estimates for volumes of traffic in 2001 and 2002 similar to those of Table 1. Larry Roberts of Caspian Networks announced some estimates (which featured high growth rates) in presentation slides at the Caspian Web site, and at numerous conferences. However, there are serious questions about the reliability of the data he obtained.³⁴

A highly regarded series of reports on Internet traffic is being produced by RHK, Inc., a market research and consulting firm, (<http://www.rhk.com>). Starting in mid-2001, RHK obtained cooperation from some of the

largest ISPs, accounting for more than half of the backbone traffic in the U.S.. By analyzing data provided by the cooperating carriers about their peering with other carriers, RHK could estimate total traffic. RHK's estimate is that North American backbone traffic grew 107% in 2001, 85% in 2002, and will grow 76% in 2003 (with declines to 48% in 2007). Their estimate for year-end 2002 North American backbone traffic was 167,000 TB/month. This estimate is consistent with the 80,000 to 120,000 TB/month of Table 1. The estimates of Table 1, like those of the earlier papers,⁴⁻⁶ attempt to count each byte of traffic just once, as it enters or leaves an end-customer machine. On the other hand, RHK's methodology involves some double counting. They add up the traffic volumes for the various carriers. However, most packets cross over several backbones (although typically no more than two of the Tier-1 carriers participating in RHK's studies). Thus if we allow for the double counting, the volume of real end-user traffic represented in RHK numbers falls in the range of values in Table 1.

The general conclusion is that current growth rates are probably close to RHK's estimate of about 76% per year in North America, and somewhat faster in Europe and Asia. What that means is that the Internet is still growing vigorously, almost as fast as it did from 1997 on. The telecom crash was not caused by a decline in traffic growth rates, but by the "irrational exuberance" that led to gross overinvestment.

As a final remark, the estimates of Table 1 are just for the U.S., which means (given the uncertainty in the numbers) they also apply to all of North America. Other estimates (from RHK, for example) suggest that U.S. Internet backbone traffic is probably close to 50% of the world total.

Future growth rates are uncertain. RHK and some other observers (such as some market research and investment houses¹⁶) predict that they will continue declining, down to 50% or 60% per year by 2006 or 2007. That is possible, but (as will be discussed later) they may continue doubling each year. (That was the prediction from IDC in early 2003, for example.¹⁸) "Moore's laws" are not laws of nature. Furthermore, even when there are extended periods of steady growth at a constant rate in a technology, that growth rate can periodically shift. For semiconductors, the traditional "Moore's Law" has held remarkably well over more than three decades, but only if it was interpreted in a certain way. In other areas, experience has been different. For example, in hard disks, there was steady but slow progress until about 1990 (at a rate of about 30% per year in areal density, say), then steady but much faster progress (at 60-70% per year) in the early 1990s, yet faster progress (around 100% per year) in the late 1990s, and now a reversion to improvements at 60-70% per year that is likely to be sustained for a while. In data transmission, it was pointed out in Ref. 6 that until the arrival of the Internet, data traffic had been growing at something like 30% per year. It is possible that the rapid spurt of growth we have witnessed was an aberration, a catch-up phase as global data connectivity was established. (Prior to the arrival of the commercial Internet, most data communication was within enterprises.) And indeed some large companies do report that their internal data traffic growth rates have subsided down to a 20-40% per year rates. On the other hand, it is possible that growth may continue at current rates, or even accelerate slightly. What is unlikely to happen, though, is growth at the "doubling every 100 days" rate. Technology and economics are almost guaranteed to prevent it. The two-year period of 1995 and 1996 when such growth rates prevailed was anomalous. On the other hand, growth close to a doubling each year for the remainder of this decade appears feasible and even likely.

3. INTERNET GROWTH MYTHS

The most popular and extremely misleading myths of the dot-com and telecom bubbles was that "Internet traffic doubles every 100 days" (or 3 months, or 4 months). It was very widely held. For example, the former FCC Chairman Reed Hundt wrote in his book¹⁵ *You Say You Want a Revolution* that "[i]n 1999, data traffic was doubling every 90 days" This claim was also mentioned (as just one example among many) in two separate articles by two separate authors in the Nov. 27, 2000 issue of *Fortune* magazine. The myth of "Internet traffic doubling every 100 days" was not just a harmless example of the many urban legends. It was often cited by scientists to demonstrate the need for research in transmission (cf. Ref. 13). It was also often a clincher for new venture business plans, the proof that wonderful things were happening on the Internet, that we were living on "Internet time," and that it was imperative to move quickly in order to get in on the new "California gold rush" taking place in cyberspace. And indeed, how else could one justify valuing JDS Uniphase (again, taking just one small example of many) at over \$100 billion, unless spending on telecom infrastructure was about to explode? Bernie Ebbers of WorldCom stated explicitly in a March 6, 2000 presentation at Boston College (cited in a news

story,¹⁴ full transcript available from Boston College) that as a result of surging demand, WorldCom capital spending in 2003 would have to exceed \$100 billion. Stories such as Ebbbers' were widely accepted, and led to huge financial losses and personal and business dislocations of the crash (as well as wealth for a few²⁰).

The myth of "Internet traffic doubling every 100 days" did not come out of thin air, and had a basis in fact. It appears to have originated during the period of abnormally rapid growth at about that rate during 1995 and 1996. For example, a Feb. 19, 1997 WorldCom press release⁴⁵ talks of "traffic over the backbone almost doubling every quarter." This may very well have been correct, as data collected in Ref. 4 do not allow for determination of when this growth rate subsided. (In addition, there is plenty of evidence that the experience of different carriers varied widely.) The memory of this brief period of manic growth (starting when the Internet was tiny) appears to have led to a perception that doubling every quarter was normal for the Internet. For example, the famous 1998 U.S. Department of Commerce report⁸ "The Emerging Digital Economy" stated explicitly that Internet traffic was "doubling every 100 days." As a source, it gave a Nov. 1997 Inktomi white paper (a text version of which can be obtained from the Internet Archive, <http://web.archive.org/>), which in turn cited Mike O'Dell of WorldCom's UUNET. However, the graph showing rapid traffic growth extended only to the end of 1996, suggesting that Inktomi obtained its information from O'Dell around then. The fact that the data was over a year old did not appear to deter the authors of the report⁸ from relying on it.

The Feb. 19, 1997 WorldCom press release⁴⁵ should have made readers cautious about extrapolating the claimed growth rates to infinity. The very same sentence that talked of traffic doubling every quarter also talked of "dial access demand growing at the rate of over 10% every week." Growth by 10% every week corresponds to 14,100% or 142x growth per year, and although it is not entirely clear what "dial access demand" refers to, it should have been clear that this growth rate could not persist for long. (Within a year or so every person would have had to be on the Internet for 24 hours per day to keep this rate of increase on track.) A doubling of traffic every quarter could be sustained longer, but even then not much longer. Yet few people paid attention, and many helped propagate and inflate the myth. Journalists were repeating it (as in the *Fortune* stories mentioned above), and so were financial analysts and industry figures. As one example, during the financial analysts' conference to present the results of the third quarter of 2000, Mike Armstrong, the CEO of AT&T, stated that Internet traffic as a whole was doubling every 100 days.

At its base, most of the support for the myth of astronomical growth rates in Internet traffic was coming from WorldCom. (This was already noted in the papers,^{4,5} but much more detailed evidence has been collected recently, with the collapse of WorldCom.) It was the largest ISP in the world, and at times claimed to be carrying around half of the world's backbone traffic (a figure that appears to have been exaggerated, but not by much). Thus its pronouncements were bound to be paid attention to, as they were likely to reflect the behavior of the entire Internet. (A startup would naturally have an infinite growth rate at the beginning, and would be expected to have high growth rates in its first few months, but that would not say much about the worldwide network.) Moreover, WorldCom was just about the only ISP that was publicly claiming astronomical growth rates for its own network. In his presentation mentioned above, Mike Armstrong did not claim that AT&T's traffic was doubling every 100 days, only that overall Internet traffic was growing that fast. (Later statements by other AT&T officials revealed that its Internet traffic at that time was growing about 300% per year, far faster than the industry average of about 100% per year, but far short of the 1,155% per year that a doubling every 100 days implies.) On the other hand, WorldCom officials such as Bernie Ebbbers, John Sidgmore, Kevin Boyne, and Mike O'Dell frequently talked of the rapid growth of their own network, which they would naturally be expected to know about. That was certainly the case with the quotes in Ref. 2, 14, 21, 41, and they usually talked of annual growth of 8x, or 10x, or "doubling every three month," which corresponds to growth of 16x per year. For example, a September 2000 article² said that

"Over the past five years, Internet usage has doubled every three months. We're seeing an industry that's exploding at exponential rates," said Kevin Boyne, chief operating officer of UUNet, WorldCom Inc.'s Internet networking subsidiary.

A mid-2002 new story⁴⁰ quoted WorldCom sources to the effect that "during recent years" traffic on UUNET's backbone had been growing at "from 70% to 80%." However, that was definitely not the story that was usually attributed to WorldCom. There are still many mysteries about the myth of astronomical Internet growth rates.

For example, most of the WorldCom claims about astronomical growth rates of their network were about network capacity, not traffic. (The Kevin Boyne interview cited above is an exception.) One of the remaining mysteries of this story is how it was that claims about network capacity were universally interpreted and passed on as claims about network traffic. Still, whether the claims were about traffic or capacity, these claims should have aroused suspicion early on. A doubling every three or four months means growing 8x or 16x per year, and if one compounds these rates for even a few years, one comes up with absurd figures.³⁰ Moreover, there were various implausibilities and inconsistencies in the WorldCom claims. As just one example, a report¹⁰ on an April 2000 presentation by Jack Wimmer, vice president of network and technology planning for MCI WorldCom has a variety of statistics that are hard to reconcile. For example, this report says that “To meet the needs of its share of 200 million Internet users, MCI Worldcom’s UUNet division has expanded its backbone 200 times since year-end 1995, Wimmer reports.” Now if we assume he was talking of year-end 1999 (the interview was in April 2000, and the anomalies are even greater if he is talking of that date), we have 200x growth in 4 years, for annual growth rate of 3.8x. If we assume that growth was 10x in 1996 (which is what Kerry Coffman and I estimated for the annual growth rate of Internet traffic during what we feel were the anomalous years 1995 and 1996 of abnormally fast growth from a small base) then we get a growth rate of 20x over 3 years, which comes to 2.7x per year. Either growth rate is a far cry from the 8x, 11x, or 16x that have been claimed at various times for UUNET by various of its spokespeople.

Other glaring inconsistencies were plentiful. In 1998, John Sidgmore was claiming consistent 10x annual growth in UUNET network capacity.⁴¹ At the March 2000 presentation at Boston College referenced above, Bernie Ebberts talked about the bandwidth of his network growing over the previous three years at 8x per year, which he then implicitly claimed in that same sentence was 800% per year (sic!). Perhaps the most instructive example to consider is the May 2000 lecture at a Stanford conference by Mike O’Dell.²¹ The audio part of the lecture talks of 10x annual growth rates, and slide 8 states that growth was by a factor of a million between 1993 and 1999, which does correspond to 10x annual growth exactly. (That slide also predicts growth by between one and 10 million over the next 5 years, which corresponds to annual growth rates of either 16x or 25x.) However, unlike the Ebberts,¹⁴ Boyne,² or Sidgmore⁴¹ references, this one has some actual numbers for network capacity. It states that the UUNET domestic network had capacity of 5,281 OC12-miles in mid-1997, 38,485 in mid-1998, and 268,794 in mid-1999. What is remarkable is that the jump from 5,281 to 38,485 is 7.3x, and from 38,485 to 268,794 is 7.0x, a not-insignificant difference from the 10x claimed. (One could object that O’Dell could have been talking about UUNET global capacity, but if one looks at the data on the slides, it is clear that inclusion of international links could not affect the growth rates much, as the vast majority of network capacity was domestic.) What is most interesting is to take the mid-1997 figure of 5,281 OC12-miles, and combine it with the claim of slide 8 and the verbal part of the presentation (as well as of the Sidgmore paper⁴¹) of 10x annual growth from mid-1993 to mid-1997. Over those 4 years, 10x annual growth compounds to 10,000x growth, which implies that in mid-1993, UUNET must have had only 0.53 OC12-miles in its network. Now 0.53 OC12-miles equals 2,800 DS0 (56 Kb/s) miles, which is about one voice line across the continent! This is certainly absurd, since UUNET had an extensive nationwide network of T1s (often multiple T1s) by that time. Yet somehow this obviously preposterous claim passed unchallenged.

There was one slide in the the O’Dell presentation²¹ that was indisputably correct, namely the last one, “If you aren’t scared, you don’t understand,” but not in the way it was meant to be. It was intended to convince the audience that the Internet was growing so rapidly, that the world was going to be upturned. (For more on the usage of this mantra by WorldCom people, see the book by Malik.²⁰) This, as we have learned, was simply false. However, the slide was correct in that the telecom industry was just then beginning to crash, and anyone involved in it should have been scared.

The dot-com and telecom bubbles and crashes are behind us, but it is instructive to look back to figure out what went wrong. There are several lessons to be drawn from the myth of astronomical Internet traffic growth. One is that almost all people are innumerate, lacking the ability to handle even simple quantitative reasoning, and in particular to appreciate the power of compound interest. Another one is that people are extremely credulous, especially when the message they hear confirms their personal or business dreams (as the Internet growth myth did, by offering the prospects of huge growth in telecom and effortless riches for participants in the game). They are not willing to examine contrary evidence, and overlook glaring implausibilities and inconsistencies in what

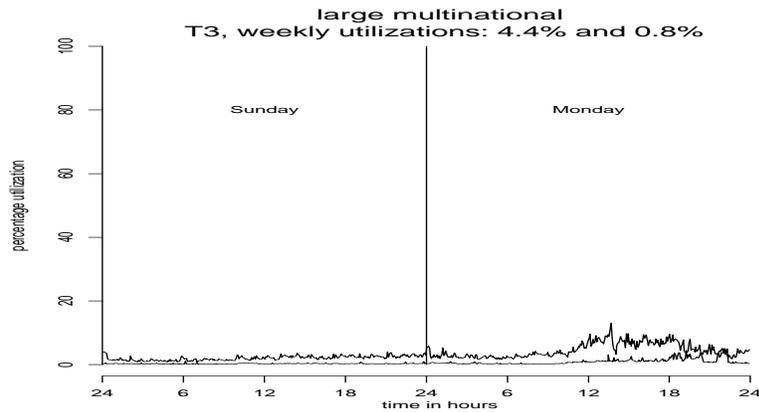


Figure 1. Traffic on a corporate T3 (45 Mb/s) link.

they hear. Finally, myths are very persistent, since respectable financial analysts and reporters were still writing about “Internet traffic doubling every 100 days” as late as 2002.

The moral of this story (and the reason it is covered in so much detail) is that bad ideas are often remarkably difficult to discredit, even when there is extensive evidence against them. Thus it should not be surprising that there are other misleading ideas that are still widely believed, perhaps because they have not proved as destructive as the myth of “Internet traffic doubling every 100 days.” One of them is the myth that “content is king,” and the associated underappreciation of simple connectivity over content.^{29,31} It is leading wireline and wireless service providers to deploy the wrong technologies in the search for “content” revenues from streaming multimedia, and neglecting what are likely to be much more profitable opportunities in seemingly more mundane areas.^{31,33} Another misleading idea is that of metered rates. Amazing enough, even as their employers rush (finally, after long but predictably futile resistance^{29,32}) to offer flat rates for long distance and wireless, or packages of local, long distance, and wireless voice services, some high level managers still argue that healthy development of the Internet requires usage-sensitive pricing. There is also continuing belief in the need for comprehensive measures for quality of service (QoS) on the Internet. Later sections point out how the growth rates observed on the Internet impact some of these ideas.

4. DATA NETWORK UTILIZATION

It is remarkable that little attention has been paid to the issue of network utilization, since it is essential to understanding the past and future of the Internet. The accepted wisdom, that data networks are chronically congested, is simply wrong. But so are some other claims.

Towards the end of 2000, Mike O’Dell of WorldCom, appeared to be saying that while the traffic on UUNET was doubling each year, network capacity, as measured in (gigabits/sec)*miles, “must double every 4 months or so” on this network.²² He also claimed that this followed from “a pretty simple result from graph theory.” Several networking industry consultants and researchers have tried to duplicate this “pretty simple result,” but without much success, and the claim appears more and more questionable as time goes on. No other carrier has reported such phenomena, and while at the end of 2000 it might have seemed possible that UUNET, as the world’s largest ISP, had run into something unusual, by now several carriers are larger than UUNET was then. Furthermore, that claim was not very plausible even then. If UUNET’s network capacity was growing 8x annually while traffic was growing 2x, then average utilization should have been dropping by a factor of about 4 each year. (The precise figure depends on the distance distribution of traffic, which may be getting smaller, but

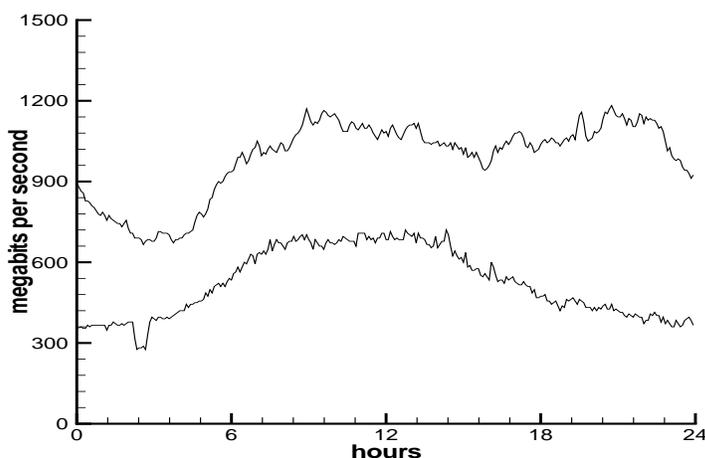


Figure 2. Traffic on AboveNet OC-192 link (9.6 Gb/s) from Washington, DC to New York City, Monday, June 9, 2003.

is not changing too rapidly.) But then the power of compound interest takes over. Even if UUNET’s backbone was operating at 100% of capacity at year-end 1996, if traffic was growing about 2x annually while capacity was growing 8x, it must have been running at most 25% of capacity at year-end 1997, at most 6.25% of capacity at year-end 1998,, and at at most 0.1% of capacity at year-end 2001. That low a level of utilization is hard to believe.

While the O’Dell claims²² suggested very low and rapidly decreasing utilization of Internet backbones, the general opinion has been that they are congested. (For references, see Ref. 23.) This opinion was consistent with one of the key mantras of the telecom bubble, namely of “insatiable demand for bandwidth.” As just one example, Kevin Boyne of UUNET was quoted² in late 2000 as saying “[a]s soon as more capacity becomes available, the Internet community will find interesting, clever ways to use it.” Even more reliable sources, such as the British research network JANET, contributed to the creation of this opinion through misleading press releases about demand instantly saturating increased capacity of data links, releases that were contradicted by the data available on their Web site. (Details are available in Ref. 5.)

The truth is that data networks are relatively lightly utilized, especially when compared to long distance voice links. A brief account is available in Ref. 26, with more details in Ref. 23. Ordinary corporate private lines as well as links to the public Internet tend to have traffic profiles like that of Fig. 1. (Exceptions tend to be lines owned by ISPs, or Web hosting companies.) Utilizations are typically in the 3-5% range (over a full week, say, with peak hour utilizations considerably higher). Even backbone links are not loaded very heavily. Fig. 2 shows the traffic profile of an OC-192 in the AboveNet network, with average utilizations in the two directions of 5.4% and 10.3%. (For the AboveNet network, the only large one in the U.S. that has made detailed information publicly available for several years, average utilizations have been around 10% for the last few years. For the sample of 16 OC-48 interfaces on the Sprint network for which information for April 7, 2003 was made available at <http://ipmon.sprint.com/>, average utilizations were also close to 10%.) In general, the estimate of Ref. 23 was that Internet backbones were running in 1998 at between 10 and 15% of their capacity over a full week, and that estimate still seems to be valid in 2003. Advances in traffic engineering were counteracted by other factors, including the shift away from SONET restoration to mesh architectures with restoration done at IP level. In fact, it is quite possible that average North American Internet backbone utilizations are far lower right now. Many of the new long distance carriers created in the late 1990s have built large backbones, but have hardly any traffic. This is a temporary situation that will disappear with time, but there are other factors that are likely to keep data network utilizations down in the future.

As another illustration of data network utilization, consider residential broadband links. Dial modem sub-

scribers used to download about 60 MB/month (with far smaller uploads), but that figure may have decreased, with many heavy users switching to broadband. Back around 1999, subscribers purchasing DSL or cable modem services tended to download in the range of 300 to 600 MB/month. Today, with the advent of P2P services and the general growth in usage of the Internet, traffic is heavier. There are no publicly available statistics, but a large DSL service provider and a large cable modem provider gave me information about their customers, which showed average downloads of 1 GB/month and 2 GB/month, respectively. (Uploads were half of the downloads in both cases, and there were substantial geographic variations.) If we assume that the cable modem customers all had 1.5 Mb/s connections, we find that a download of 2 GB/month corresponds to a utilization rate (over the full month) of 0.4%. Note that if these cable modem customers only cared about the volume of data, they could download the full 2 GB/month over their regular modem connection, or else they could get it at far lower cost by renting DVD's through Netflix. (This is a general phenomenon.¹⁹ For the same financial cost, it is often possible to get higher data rates through the postal system or through loading up 747s with DVD than by using fiber optic links.)

The reasons for low utilization of data networks are explored in detail in Ref. 23. Two basic factors are involved. At the access link level, utilizations are light because what data networks are for is to provide low transaction latency, making sure that the database update happens quickly, or a Web page shows up on the screen quickly. At the backbone level, that is less important, since there is extensive statistical multiplexing. There, the high growth rates of traffic and the lumpy nature of network capacity are the major culprits.

The implications of low data network utilization are explored extensively in the papers Ref. 24, 27, 28. I will briefly mention some of these implications in later sections. Right now let us note that low utilizations provided obvious disproofs of the myths of "insatiable demand for bandwidth," as well as many basic assumptions as to what data networks were for.

Low utilizations of data networks also have some other interesting side effects. Since there is a less direct connection between network capacity and observed performance than in the voice network, upgrades in data networks can be postponed far more readily. Since upgrades depend on subjective judgements (including judgements of financial officers as to whether investors would respond positively to new capital investments), they are subject to herd instinct behavior, and might become volatile.

5. IMPORTANCE OF TRAFFIC GROWTH RATES

The growth rate in Internet traffic is the most important factor in determining demand for equipment. It also has many other implications for the kind of equipment that is ordered, what services are offered, and even for the basic architectural foundations of the Internet (such as the end-to-end principle). Growth rates at the mythical "doubling every 100 days" rate were feasible in the early days, when the Internet was small. Today they are not, at least not for any extended period of time. A period of 3x or 4x annual growth could be accommodated for several years, but is unlikely for several reasons. One is that it would require huge jumps in spending, as will be discussed later. Another is that historically we have not seen any examples of such large jumps in traffic at any institution that already had extensive data communications. The trend has been for traffic to grow at about 2x per year even in the absence of bandwidth constraints.

A period of approximately 2x annual growth (meaning, as before,⁴⁻⁶ growth between 1.7x and 2.5x per year) appears feasible for the rest of this decade, and might be just about optimal for the industry. It would allow for decent growth in revenues, and might bring some profitability to the industry. There are reasons for thinking that such growth might be obtainable, as I will discuss in Section 7.

Today, many observers (such as RHK, mentioned in Section 2, or the report cited in Ref. 16) appear to be predicting a decline in the growth rate of Internet traffic down to the 50% a year range in the next few years. Such a decline would likely lead to a further squeeze for the long-haul industry, as it would lead to declining revenues, as is discussed in Section 8, and to an even greater degree of carrier consolidation than already seems inevitable.

Even growth rates of 2x per year would imply that there will be no need for new fiber in the long distance networks for the foreseeable future, and some very exciting technologies, such as customer-owned wavelength switching, would see limited applicability in this decade.

6. QoS

There are many reasons why quality of service (QoS) technologies are inappropriate for the core of the Internet, and should be used only sparingly at the edges of the network.^{24, 27-29, 32} The experience of the last few years has only reinforced those arguments (and QoS has indeed not been used widely). At this point let me just remark that many of the arguments against QoS depend on an assumption of continuing rapid growth of the network. Were the Internet to stabilize, the way the old voice network was stable for a long time, some of the arguments against QoS would be weakened. I will illustrate with a simple but illustrative example.

On July 2, 2001, in response to a question on the NANOG (North American Network Operators' Group) mailing list as to what were "the most common causes of performance problems as well as outages," Sean Donelan, an experienced network engineer, responded that they were⁹

In roughly the order

1. Network Engineers (What's this command do?)
2. Power failures (What's this switch do?)
3. Cable cuts (Backhoes, enough said)
4. Hardware failures (What's that smell?)
5. Congestion (More Bandwidth! Captain, I'm giving you all she's got!)
6. Attacks (malicious, you know who you are)
7. Software bugs (Your call is very important to us....)

Note that QoS might help alleviate only one of the problems on this list, namely the fifth one (lack of bandwidth), and would likely worsen the first and seventh ones. Thus introducing QoS in such an environment is likely to increase rather than decrease costs.

A key point about the Donelan list is that most of the problems on it are caused by high growth rates. The complexity of the growing and changing network makes it hard to manage and leads to mistakes by designers and operators. Hardware and software are unreliable, since the emphasis is on getting products and services out. As long as the Internet continues expanding rapidly, this basic problem is not likely to go away. Should the Internet stabilize, though, one can imagine that approaches associated with the old voice network would be more appropriate.

7. SOURCES OF DATA TRAFFIC GROWTH

Very little is known about the observed patterns of Internet traffic growth. Some available evidence is gathered and analyzed in the papers Ref. 4-6. One of the key findings (related to the low utilization of data networks) was that even in the absence of local network congestion, traffic at large institutions with diversified user bodies tended not to increase by more than 2x per year. Table 2 shows statistics for the University of Waterloo, showing one instance of 3x growth, and otherwise growth mostly in the 1.5-2x year. These constrained growth rates are the result of complicated feedback loops operating on different time scales, from individual users deciding how much Web surfing to do, to network managers deciding how much bandwidth to provide, to venture capitalists deciding which innovations to finance. Another finding was that in the absence of strong constraints, traffic often tended to grow close to 2x annually, at least at large and diverse institutions. For more details, see.⁴⁻⁶

A key element of the studies⁴⁻⁶ is that "killer apps" are not required for rapid network traffic growth. The approximately 10x annual growth during 1995 and 1996 visible in Table 1 was associated with the advent of the browser, and the rush of millions of people onto the Web. However, at many institutions that had been on the Internet for a long time, such as the University of Waterloo, the rise of the Web led to only a moderate increase in total traffic growth (cf. Table 2). At those places, the "disruptive innovation" of the browser just reinforced the high growth rates that were already present.

Today, P2P applications are the most prominent contributor to traffic growth. However, they are not the only one. Essentially all other types of traffic are still growing. Table 2 is again an interesting example. (The Waterloo Web site listed in Section 2 from which this data is taken presents a more detailed analysis of their traffic.) P2P is included in the "other" category there, but so are a variety of unknown applications. (For

general information about distribution of traffic by application on a variety of networks, see the reports at.^{3, 39, 42} As one example, the April 7, 2003 data for the Sprint backbone available at (<http://ipmon.sprint.com/>) shows Web as 47% of the traffic, recognized file sharing applications as 21%, and various “other TCP” applications at 23%. It is important to note, though, that identifications are often uncertain, since many applications, especially modern P2P ones, can masquerade as others, for example as Web browsing.)

P2P is often taken as synonymous with sharing of commercially produced music and video files. While such uses may indeed dominate in volume of P2P traffic, that is not all that P2P is used for. There is an increasing volume of enterprise and scientific P2P traffic, as well as residential customers distributing material they generate.

Largely because of their association with music and video file swapping, the spread of P2P applications is inhibited by frequent reluctance of network managers to provide the necessary bandwidth. At many academic institutions, administrators have been reluctant to see too big a fraction of their capacity used by students, and so have imposed various restrictions (such as capping total bandwidth available to dormitories, or limiting bandwidth of individual flows, or even limiting the volume of transmission of individual students). P2P traffic is also being constrained by legal measures by content producers. Still, historical precedents^{29, 31} suggest that the full potential of the Internet will be realized only when it is used extensively for social communication. This would mean swapping of video clips of family picnics and the like. Although that may not be the dominant type of traffic in terms of volumes, it might very well become the most valuable. (The correlation between volume and value in telecommunications is very weak, which is why cell phones attract much higher revenues than the Internet.) However, it will take time for such traffic to become significant, and in the meantime professional music and increasingly video is what is driving P2P traffic. In order to increase traffic growth, carriers should be working with content providers on new business models that would provide for legal distribution of “content.” I have been suggested, only partly in jest, that the telecom industry should pay off the music distributors in order to allow recorded music to be circulated for free.^{31, 36} Content is not king, and recorded music revenues are a tiny fraction (under 5%) of those of the telecom industry, so this might be a good investment, if it were not totally impractical for a variety of political, economic, and legal reasons.

Discussion of the growth of the Internet is often dominated by questions about spread of broadband to residential users. However, such users are not always the main drivers of network development. Unfortunately we simply do not have detailed statistics to determine which sector of the economy produces or consumes the most data in most cases. In Australia, as a result of the recent government statistics program,¹ we know that in Sept. 2000, 43% of traffic was received by business and government customers, whereas in Sept. 2002, this percentage was down to 35%. The situation is different in the U.S.. Although we don’t have detailed information, the estimates cited in Section 4 show that traffic of residential dial modem users is almost negligible, just a few petabytes per month. It is far surpassed by traffic of residential broadband users, but even those send and receive a total of at most 60,000 TB/month, compared to total send and receive volume of between 160,000 and 280,000 TB/month. Thus it appears that business traffic dominates in the U.S., and the role of residential broadband should not be overemphasized in the discussions of Internet development.

The networking industry and networking research community are both preoccupied with the notion of streaming real-time multimedia transmissions. However, it was forecasted a long time ago, apparently first by Negroponte, then by Gilder, and later by St. Arnaud⁴³ (see Ref. 28, 37 for more arguments) that network traffic would be dominated by file transfers. That prediction has proven correct so far. The April 7, 2003 data for parts of the Sprint network cited above show file sharing programs as 21% of the traffic, while streaming is down at 2%. The preponderance of file transfers is likely to hold true into the future, especially if high growth rates of traffic continue.

8. INTERNET COSTS AND REVENUES

In the last few years the Internet has been attracting just about all the attention paid to telecommunications. However, voice is still where the money is. The Internet is not all that large in terms of revenues. (On the other hand, its bandwidth is far larger than that of the voice networks, and the traffic on it has recently exceeded voice traffic, at least in the U.S..) Dedicated Internet access (the kind that businesses, government, and academia buy) brings in around \$15 billion per year. The dial modem customers, still over 40 million, bring in another \$10

billion or so. Finally, residential broadband customers (almost 20 million by year-end 2002) bring in yet another \$10 billion or so, for a grand total of \$35 billion per year. By comparison, the total telecommunications revenues in the U.S. were \$345 billion in 2001,³⁸ with similar spending in 2002, and cellular revenues alone were running at a rate of over \$80 billion per year by year-end 2002.⁷

The implicit message of the “Internet traffic doubling every 100 days” myth was that spending on the Internet was going to explode. Sometimes this message was spelled out explicitly, as in the claim that revenues would grow fast enough to allow WorldCom alone to spend \$100 billion on network capital investments in 2003.¹⁴ Instead, spending on dedicated Internet access in the U.S. has stabilized, at least temporarily. (On the other hand, residential broadband is growing vigorously, about 60% in 2002.) The challenge for the carriers is to induce higher spending from customers. The lesson of history²⁹ is that telecommunications spending goes up, usually somewhat faster than GDP. Still, it typically does not grow too fast. For dedicated access, the best one could probably hope for in terms of revenues is something like 20% to 30% annual growth. This would be similar to what was achieved by Intel in the late 1980s and through the 1990s,¹¹ when it sold increasing number of microprocessors, with each generation more powerful than the previous one, but at a comparable price. The reason one can’t hope to do much better is that we have fierce competition and overcapacity. Furthermore, even if those factors were alleviated by industry consolidation, as they are likely to be, we would still be stuck with the basic fact that core Internet services are not all that expensive to provide. The late greenfield fiber players were able to construct nationwide networks from scratch for under \$10 billion each. However, much of that money went for basic construction of fiber conduits and buildings. That part of the work is done, and is a sunk cost. Since there is a glut of this basic infrastructure, it sells for a small fraction of the construction cost. Hence one can build a complete network very inexpensively. Furthermore, barring some major consolidation that allows a single carrier to monopolize the physical facilities, it will continue to be inexpensive to build new networks, and for a given cost, the network capacity that can be constructed will be climbing rapidly. Advances in DWDM technology mean that there will not be any need for additional fiber or regeneration huts for the rest of this decade. Now one does need to light the fiber and install routers to handle the customer traffic, but the cost of equipment for doing this is dropping rapidly. It is not dropping by 2x annually, but when equipment cost declines are combined with the simplifications in network architectures (getting rid of ATM and SONET, etc.), as well as economies of scale, and improved network management tools, we could have a number of years where the cumulative cost improvement on carrier side might be close to 2x per year. What this is likely to lead to is a number of years in which carriers can double the traffic they carry on their backbones each year with no extra revenues. Another way to look at this is that unless traffic does grow at 2x annually, or close to that rate, revenues of ISPs are likely to start declining. Hence carriers will have to concentrate on increasing usage, and turn away from the attitude of limiting transmissions.¹² It is a tricky situation, in that carriers have to reach for profitability, but unless they ensure that demand keeps up with technology, they could get squeezed out.

Can anyone make money in Internet transport? Some have been doing so from the beginning, and more will likely learn how to do it. There are great economies of scale, and demand is growing. However, basic Internet transport is a commodity, and, as will be discussed in the next section, the backbones are not that expensive, and are unlikely to generate huge revenues. Now commodity businesses are often profitable, and Internet backbones will presumably be profitable eventually for some players. However, the profits will likely be limited, since most of the revenue and profit opportunities are at the edges of the network.

9. THE HOLLOWING OUT OF THE CORE

Although the traffic on the Internet exceeds that on the voice network, it is not all that large when compared to what progress in photonics has made possible. The upper estimate in Table 1 for U.S. Internet backbone traffic at year-end 2002, 140 PB/month, is equivalent to only about 420 Gb/s on average. Even if we make allowances for a considerably higher than usual 1.9 ratio of peak hour to average traffic on backbones, we find that it would take only 800 Gb/s of capacity to transport all the traffic on all U.S. backbones, and that much capacity can now comfortably fit on a single fiber! Of course, there is no interest in sending all the traffic from one point to another, and I am ignoring issues of restoration, etc., but this thought experiment shows graphically why there is a fiber glut right now.

Even for dedicated Internet access, revenues and costs are largely at the edges of the network. It is now (in mid-2003) standard (as can be seen in the quotes at Band-X, <http://www.band-x.com>), for example) to obtain transit service from large ISPs for about \$100 per Mb/s per month in large cities, if one leases high-bandwidth links. But that says that the 800 Gb/s of capacity that would suffice to carry all of U.S. backbone traffic would cost only \$80 million per month, or under \$1 billion per year. Yet total dedicated access revenues in the U.S. are about \$15 billion, showing that most of the money is in aggregating low-bandwidth links at the edges of the network.

That backbone costs are very low can also be seen in various other ways. Consider residential broadband, DSL or cable modem. As was computed in Section 4, even for typical cable modem downloads of 2 GB/month, average traffic downstream is only about 6 Kb/s. Hence even if we allow 20 Kb/s of bandwidth per customer to allow for uneven distribution of traffic, the cost of backbone transit will be only \$2 per month per customer. This compares to typical residential broadband prices in the U.S. that are now in the \$40-50/month range.

The core of the network has been shrinking in importance for a long time, largely because technology has reduced costs there the fastest. Even in voice telephony, by the mid-1990s, the actual network costs of long distance transport were low, under a penny a minute. (This largely explains why voice-over-IP has not been disruptive, as the savings it offered were low.) The Internet has accelerated this trend. It is a wonderful technology, but the backbone transport it offers is a commodity that is inexpensive to provide. The costs at the edges, on the other hand, are very high.^{24,25} This suggests that we are likely to move towards new models, with largely customer-owned edge networks, and with large opportunities in outsourcing.³⁵

10. CONCLUSIONS

Internet traffic continues to grow vigorously, close to 100% per year. On the other hand, revenues from backbone transport are low, and will likely remain so. This is partly a result of the telecom bubble, which led to gross overinvestment that led to cutthroat competition. It is also the result of the natural evolution of the industry, in which the core of the network, although huge in terms of transport capacity, is relatively inexpensive to construct and run. The main challenges (and the opportunities) are at the edges of the network. This is likely to require a restructuring of the telecom industry.³⁵ It should also lead the research community to focus its energies away from QoS and sophisticated pricing schemes and towards managing complexity.

The main imperative for carriers is to increase usage. They do need to obtain enough revenues to cover their costs and make a profit, but they cannot hope to do so unless traffic keeps up with technological progress. They will need to sell data connections on the basis of advantages of speed, and will need to emphasize flat rates (dependent on bandwidth) as a way to stimulate usage.

REFERENCES

1. Australian Bureau of Statistics, "Internet Activity: September 2002," Report 8153.0, dated March 7, 2003. Available for purchase at <http://www.abs.gov.au>.
2. P. Behr, "Technology investors have a new specter to worry about: The bandwidth glut," *Washington Post*, Sept. 24, 2000.
3. Cooperative Association for Internet Data Analysis (CAIDA), report page at <http://www.caida.org/outreach/papers/>.
4. K. G. Coffman and A. M. Odlyzko, "The size and growth rate of the Internet," *First Monday*, **3**, no. 10, Oct. 1998, <http://firstmonday.org/>. Also available at <http://www.dtc.umn.edu/~odlyzko>.
5. K. G. Coffman and A. M. Odlyzko, "Internet growth: Is there a "Moore's Law" for data traffic?," *Handbook of Massive Data Sets*, J. Abello, P. M. Pardalos, and M. G. C. Resende, eds., pp. 47-93, Kluwer, 2002. Available at <http://www.dtc.umn.edu/~odlyzko>.
6. K. G. Coffman and A. M. Odlyzko, "Growth of the Internet," *Optical Fiber Telecommunications IV B: Systems and Impairments*, I. P. Kaminow and T. Li, eds., pp. 17-56, Academic Press, 2002. Available at <http://www.dtc.umn.edu/~odlyzko>.
7. CTIA, "CTIA's semi-annual wireless industry survey results, June 1985 - December 2002," available at <http://www.wow-com.com/industry/stats/surveys/>.

8. U.S. Department of Commerce, "The Emerging Digital Economy," April 1998 report, available at <http://www.esa.doc.gov/TheEmergingDigitalEconomy.cfm>.
9. S. Donelan, posting to NANOG mailing list, July 2, 2001, on the subject of "Common causes of performance problems and outages." Available in online archives at <http://www.nanog.org/maillinglist.html>.
10. Fiber Optics Online, "MCI WorldCom's Wimmer lays out optical roadmap," April 4, 2000 report, available at <http://www.fiberopticonline.com/content/news/article.asp?DocID=31D66203-095B-11D4-8C32-009027DE0829>.
11. P. C. Fishburn and A. M. Odlyzko, "Dynamic behavior of differential pricing and Quality of Service options for the Internet," in *Proc. First Intern. Conf. on Information and Computation Economies (ICE-98)*, pp. 128–139, ACM Press, 1998. Available at <http://www.dtc.umn.edu/~odlyzko>.
12. S. Helms, "Recapturing an expensive resource: Bandwidth," *Cable Datacom News*, Oct. 1, 2002. Available at <http://www.cabledatacomnews.com/oct02/oct02-6.html>.
13. P. Heywood, "Bandwidth bonanza 'won't happen'," *Light Reading*, July 12, 2000. Available at http://http://www.lightreading.com/document.asp?doc_id=1142.
14. P. J. Howe, "MCI chief sees big outlays to handle net traffic: Ebberts estimates \$100B to upgrade network," *Boston Globe*, March 7, 2000.
15. R. Hundt, *You Say You Want a Revolution*, Yale Univ. Press, 2000.
16. P. Kapustka, "Report Heralds IP Doomsday," *Light Reading*, May 16, 2001. Available at http://www.lightreading.com/document.asp?site=lightreading&doc_id=5223.
17. J. Krause, "Has the Net stopped growing?," *The Industry Standard*, July 2, 2001. Available at <http://www.thestandard.com/article/0,1902,27394,00.html>.
18. E. Larson, "Report: Internet traffic boom redux," *Light Reading*, February 28, 2003. Available at http://www.lightreading.com/document.asp?site=lightreading&doc_id=2902.
19. C. Lu, *The Race for Bandwidth: Understanding Data Transmission*, Microsoft Press, 1998.
20. O. Malik, *Broadbandits: Inside the \$750 Billion Telecom Heist*, Wiley, 2003.
21. M. D. O'Dell, "Racing with an exponential or the dangers of linear thinking in an exponential world," lecture at Stanford symposium, *Optical Internet: The Next Generation*, May 16, 2000. Available in audio, video, and slides at <http://stanford-online.stanford.edu/optic/main.html>.
22. M. D. O'Dell, email message with subject line "100% per year, etc.," posted Nov. 30, 2000 on Dave Farber's IP list, available at <http://www.interesting-people.org>.
23. A. M. Odlyzko, "Data networks are lightly utilized, and will stay that way," 1998 manuscript, available at <http://www.dtc.umn.edu/~odlyzko>. Revised version to appear in *Review of Network Economics*, Sept. 2003.
24. A. M. Odlyzko, "The economics of the Internet: Utility, utilization, pricing, and Quality of Service," 1998 unpublished manuscript, available at <http://www.dtc.umn.edu/~odlyzko>.
25. A. M. Odlyzko, "Smart and stupid networks: Why the Internet is like Microsoft," *ACM netWorker* **2**, no. 5, Dec. 1998, pp. 38-46. Available at <http://www.dtc.umn.edu/~odlyzko>.
26. A. M. Odlyzko, "Data networks are mostly empty and for good reason," *IT Professional* **1**, no. 2, March/April 1999, pp. 67-69. Available at <http://www.dtc.umn.edu/~odlyzko>.
27. A. M. Odlyzko, "The current state and likely evolution of the Internet," in *Proc. Globecom '99*, pp. 1869-1875, IEEE, 1999. Available at <http://www.dtc.umn.edu/~odlyzko>.
28. A. M. Odlyzko, "The Internet and other networks: Utilization rates and their implications," *Information Economics & Policy*, **12** (2000), pp. 341-365. (Presented at the 1998 Telecommunications Policy Research Conference.) Also available at <http://www.dtc.umn.edu/~odlyzko>.
29. A. M. Odlyzko, "The history of communications and its implications for the Internet," 2000 unpublished manuscript, available at <http://www.dtc.umn.edu/~odlyzko>.
30. A. M. Odlyzko, "Internet growth: Myth and reality, use and abuse," *iMP: Information Impacts Magazine*, November 2000, http://www.cisp.org/imp/november_2000/odlyzko/11_00odlyzko.htm. Also available at <http://www.dtc.umn.edu/~odlyzko>.
31. A. M. Odlyzko, "Content is not king," *First Monday*, **6**, no. 2, February 2001, http://firstmonday.org/issues/issue6_2/odlyzko/. Also available at <http://www.dtc.umn.edu/~odlyzko>.

32. A.M. Odlyzko, "Internet pricing and the history of communications," *Computer Networks*, **36** (2001), pp. 493-517. Available at <http://www.dtc.umn.edu/~odlyzko>.
33. A. M. Odlyzko, "Talk, Talk, Talk: So who needs streaming video on a phone? The killer app for 3G may turn out to be surprise-voice calls," *Forbes*, August 20, 2001, p. 28. Available at <http://www.dtc.umn.edu/~odlyzko/doc/recent.html>.
34. A. M. Odlyzko, "Comments on the Larry Roberts and Caspian Networks study of Internet traffic growth," *The Cook Report on the Internet*, Dec. 2001, pp. 12-15. Available at <http://www.dtc.umn.edu/~odlyzko>.
35. A. M. Odlyzko, "Roxane Googin's predictions and the telecom world," *The Cook Report on the Internet*, **11**, no. 1-2, April-May 2002, pp. 53-58. Available at <http://www.dtc.umn.edu/~odlyzko>.
36. A. M. Odlyzko, "Comments on 'Solving the broadband paradox' by Adam D. Thierer," letter to the editor, *Issues in Science and Technology*, **18**, no. 4, Summer 2002, pp. 7-8. Available at <http://www.dtc.umn.edu/~odlyzko>.
37. A. M. Odlyzko, "Internet TV: Implications for the long distance network," in *Television over the Internet*, E. Noam, J. Groebel, and D. Gerbarg, eds., Lawrence Erlbaum Associates, in press. (Proceedings of workshop held at Columbia University in Nov. 2000.) Available at <http://www.dtc.umn.edu/~odlyzko>.
38. *OECD Communications Outlook 2003*, electronic and paper copies can be purchased through <http://www.oecd.org/>.
39. D. Plonka, University of Wisconsin in Madison network performance statistics, available at <http://wwwstats.net.wisc.edu/>.
40. J. Rendleman, "WorldCom soldiers on with new service," *Information Week*, Aug. 30, 2002. Available at <http://www.informationweek.com/story/IWK20020830S0010>.
41. J. Sidmore, "The largest ISP and getting larger," *Proc. Vortex 1998*, pp. 157-166.
42. Sprint Advanced Technology Lab, publications Web site at http://ipmon.sprintlabs.com/pubs_trs/webpages/ATLPUB.php.
43. B. St. Arnaud, "The future of the Internet is NOT multimedia," *Network World*, Nov. 1997. Available at <http://www.canarie.ca/~bstarn/publications.html>.
44. "Global Internet Geography 2003," TeleGeography, 2003. Can be purchased through <http://www.telegeography.com>.
45. WorldCom announces \$300 million expansion of UUNET network, WorldCom press release, Feb. 19, 1997. Available at <http://global.mci.com/news/news2.xml?newsid=6221&mode=long&lang=en&width=530&root=/&langlinks=off>.