

Internet TV: Implications for the long distance network

A. M. Odlyzko

AT&T Labs - Research

amo@research.att.com

<http://www.research.att.com/~amo>

Revised version, July 27, 2001

Abstract

The migration of traditional TV to the Internet is likely to have little impact on the long distance network. The main reason is that consumers still take on the order of a decade to embrace new technologies (such as cell phones) or even improved variants of old media (as with CDs replacing vinyl records). Hence we should not expect traditional broadcast TV to change substantially or to migrate to new modes of distribution any time soon. Yet within much less than a decade, progress in photonics will produce an increase in the capacity of Internet backbones far beyond that required to carry all the broadcast TV signals. There will continue to be bottlenecks in the "last mile" that will limit the migration of TV to the Internet (and this will reinforce the natural inertia of the consumer market). However, the backbones are unlikely to be an impediment.

The Internet is likely to have a much larger impact on TV than TV will have on Internet backbones. There is vastly more storage than transmission capacity, and this is likely to continue. Together with the requirements of mobility, and the need to satisfy human desires for convenience and instant gratification, this is likely to induce a migration towards a store-and-replay model, away from the current real-time streaming model of the broadcast world. Further, HDTV may finally get a chance to come into widespread use. The flexibility of the Internet is its biggest advantage, and will allow for continued experimentation with novel services.

Internet TV: Implications for the long distance network

A. M. Odlyzko

AT&T Labs - Research

amo@research.att.com

<http://www.research.att.com/~amo>

1. Introduction

The thesis of this paper is that traditional concerns about the impact of TV on the long distance links in the Internet are unjustified. The convergence of TV and the Internet is likely to be slower and take a path different from the one normally envisioned.

There are various definitions of Internet TV. (See [Egan, Noll4, Noll5, Owen].) The precise one does not matter much for our purposes, though.

Data networks developed rapidly largely because they could use the huge existing infrastructure of the telephone network. Without all the investments made to provide voice services, long distance data transmission would have grown much more slowly. As it is, growth has been fast, although not as fast as is commonly believed. The bandwidth of data networks in the U.S. already exceeds that of voice networks (see Section 2 for more details). Sometime in 2001 or 2002, the volume of data transmitted on data networks will exceed that of voice. (See [CoffmanO3] for the historical growth rates of different types of data and non-data traffic, and of predictions of when data would exceed voice.)

In spite of all the publicity it has attracted in the late 1990s, packetized voice is still a tiny fraction of Internet backbone traffic. This is not likely to change, even as a greater fraction of voice is sent over the Internet. The reason is the far higher growth rate of data traffic than of voice calling, about 100% versus under 10% per year. Even today, to move all current voice traffic to the Internet would require much less than a doubling of the Internet's capacity. Although there are still more bytes of voice traffic than of Internet traffic, packetization of voice naturally lends itself to compression. Hence if voice traffic were to move to the Internet, the volume of packet traffic that would result would be far smaller than current data traffic.

The general conclusion is that the volume of voice calls is not going to overwhelm the Internet. (There are quality issues as well that matter, but we will not deal with those here.) Historically, though, data networks have developed in the shadow of the telephone network. Not only have data networks relied on the infrastructure of voice telephony, but their development was strongly influenced by the

prospect that eventually they would carry voice. Since the telephone network was so much larger than the data networks, the quality requirements for voice transmission played a major role in the planning of data transmission technologies. Right now, it is increasingly realized that voice will not be a large part of the traffic in the future, simply because there is too much data. On the other hand, video is now playing a similar role to the one voice used to play. The volume of TV transmissions is so large that the requirements of real time streaming video dominate planning for the future of the Internet. However, that is also likely to turn out a mistake. By the time TV moves to the Internet, data traffic will likely be so large that streaming video will not dominate it. Moreover, the video traffic on the Internet is likely to be primarily in the form of file transfers, not streaming real time transmission.

The above contrarian predictions are based on a study of rates of change in different fields. Storage, processing, display, and transmission technologies are advancing at rather regular and predictable rates. This is considered in sections 2 and 3. (“Moore’s Law” for semiconductors is only the most famous of the various “laws” that govern progress.) In addition, rates at which new technologies are adopted by society, while not as regular, are almost universally much slower than is commonly supposed. (“Internet time” is a myth.) This is discussed at greater length in Section 4. As a result we can have some confidence in expecting that by the time TV moves to the Internet in a noticeable way, the latter will have huge capacity, at least on the long distance links.

The prediction that the predominant mode of video transmission on the Internet is likely to be through file transfers is justified briefly in Section 5. Section 6, the concluding one, is devoted to what appears to be the most likely impact of the Internet on TV, namely in providing greater flexibility that will encourage exploration of technologies such as HDTV.

2. Network sizes and growth rates

The paper [CoffmanO1] pointed out that already by the end of 1997, the bandwidth of long distance data networks in the U.S. was comparable to that of the voice network, with the public Internet a small fraction of the total. Today, the Internet is by far the largest in terms of bandwidth. However, because bandwidth is hard to measure and changes irregularly, due to the lumpy nature of network capacity as well as the financial climate, it is hard to estimate it precisely. Table 2.1 presents the estimate from [CoffmanO2] of the traffic (in terabytes, units of 10^{12} bytes, per month). The key point, discussed in great detail in [CoffmanO1, CoffmanO2] is that Internet traffic is growing at about 100% per year. That is the growth rate the Internet experienced during the early 1990s. There was then a brief period of two years, 1995 and 1996, when growth was at the “doubling every three or four months” rate that is

usually mentioned. Starting in 1997, though, growth again slowed down to doubling each year. At this rate, by some time in 2001 or 2002, there will be more data than voice traffic in the U.S., as predicted in [Coffman01].

Table 2.1. Traffic on U.S. long distance networks, year-end 2000.

network	traffic (TB/month)
US voice	53,000
Internet	20,000 - 35,000
other public data networks	3,000
private line	6,000 - 11,000

Further, technology advances in transmission and switching appear to offer the prospects of traffic growing at about 100% a year through the year 2010 without astronomical increases in spending. Even if we only have growth by a cumulative factor of 100 over the first decade of the 21st century (as opposed to a factor of 1024 that results from a doubling each year), we will have around 3,000,000 TB per month of traffic by the end of 2010, or around 10 GB per person per month. Now a 90-minute movie, digitized for high resolution at 10 Mb/s, comes to about 7 GB, so we would be able to transmit only about two movies per person (counting all men, women, and children) per month in that format. However, if we lower the resolution to 2 Mb/s, and assume traffic continues doubling each year, we find that by 2010 we could send 100 movies per person per month. Thus the general conclusion is that by 2010 or soon thereafter, the long distance Internet backbone could transmit all the entertainment TV signals that are likely to be demanded.

Another way to consider the problem of the transportation task imposed by TV is by considering capacities of fibers. If we give each of the approximately 300 million inhabitants of the U.S. a 10 Mb/s traffic stream, we find that the total demand is for 3,000 Tb/s of transmission capacity. The DWDM (Dense Wavelength Division Multiplexing) technologies that are widely deployed typically reach about 0.8 Tb/s per fiber strand, but there are good prospects of reaching 10 Tb/s in a few years, and there are even hopes of achieving 100 Tb/s. If we assume conservatively that 10 Tb/s capacity per fiber will be widely deployed by 2010, it would require just 300 strands to provide the 3,000 Tb/s of capacity that the 10 Mb/s traffic stream per person involves. (Actually, we would need double that for two directions of traffic, plus other small multiplier factors to provide for redundancy, etc., but those are not huge factors.) Today, we have several hundred strands of fiber running from coast to coast, and many empty conduits that could be filled with additional fiber. Thus as far as fiber itself is concerned, there will be

plenty of capacity.

Most of the fiber that is in place in the long distance networks is not utilized (“lit” in industry language), and even when it is in use, it is often used at a small fraction of its capacity. The reason is that there is not enough demand to create more usable capacity, certainly not even at the prices of 2001 (which are much lower than they were just a couple of years ago). (The fiber glut we are experiencing resulted from an assumption that there was an insatiable demand for bandwidth. It ignored three key factors: (i) lack of “last mile” connectivity, (ii) the cost to provide usable bandwidth, as opposed to raw fiber, and, perhaps most important, (iii) that traffic demand is growing at only about 100% per year, even in the absence of bandwidth constraints, gated more by rate of adoption of new applications than anything else.)

The general conclusion is that there already is enough fiber to allow for transmission of individual TV signals over the long distance Internet backbones, and that sometime around the year 2010, transmission and switching technologies are likely to allow for this to be done economically. The question is, will we want to do that? The volume of unique TV content is simply not all that large, as is shown in [Lesk, LymanV]. Given the trends in storage capacity mentioned below, it is feasible to store copies of all the non-real time material (which is the overwhelming bulk of what TV transmits) on multiple local servers, and avoid burdening the backbones with it.

3. Moore’s laws (technology trends)

In the previous section, there was an implicit assumption, namely that the highest resolution video signals that would be typical by 2010 would be no more than 10 Mb/s. Today, on digital cable TV systems, typical transmission rates are around 2 Mb/s, and HDTV signals tend to be compressed to somewhat below 10 Mb/s. We can certainly expect increases in resolution of video signals. (Movies are filmed at over 1 Gb/s, and stored as such.) However, these increases are likely to be modest. (Note that TV resolution has not changed in over 50 years, and HDTV and other forms of enhanced display technologies have been making slight progress, a point to be considered further later.)

In general, technological prognostications have a miserable track record. The one area where they have been outstandingly successful, though, has been in forecasting continuation of various types of laws similar to the “Moore’s Law” of semiconductors, which says that the number of transistors on a chip doubles every 18 months. (See [Schaller] for the history and fuller description. The basic law is often reported as stating that processor power doubling every 18 months, which is not quite right, but reasonably close.) The key point, discussed at greater length in [CoffmanO2], is that the different

Moore's laws for different areas operate at different speeds. Display resolution is improving slowly (and battery capacity even more slowly), while transmission and magnetic storage capacity are growing even faster than processor power. Table 3.1, taken from [CoffmanO2], shows the growth in the volume of hard disk storage that is shipped each year. It is about doubling annually, comparable to the rate at which transmission capacity is growing.

Table 3.1. Worldwide hard disk drive market. (Based on Sept. 1998 and Aug. 2000 IDC reports.)

year	revenues (billions)	storage capacity (terabytes)
1995	\$21.593	76,243
1996	24.655	147,200
1997	27.339	334,791
1998	26.969	695,140
1999	29.143	1,463,109
2000	32.519	3,222,153
2001	36.219	7,239,972
2002	40.683	15,424,824
2003		30,239,756
2004		56,558,700

The rapid growth of storage capacity is significant, since it makes non-streaming modes of operation much more attractive. Back in the 1980s and 1990s, disk storage available on PCs in households was so small that streaming real time delivery of video was the only feasible alternative. Today, local storage is becoming viable even for high resolution movies. (Note the estimate of 7 GB for a single HDTV movie, versus a capacity of 80 GB that often comes with high-end PCs in mid-2001, and the likelihood that this will reach 1 TB around the year 2005). As time goes on, and the disk capacity grows rapidly, while digital movie sizes grow slowly, the attractions of local storage will only increase.

4. Rates of change, technological and sociological

We hear constantly how we live on "Internet time," and how the Internet changes everything. Yet "Internet time" is a myth. The pace at which new products and services are adopted is not notably faster than it used to be in the past. This contrarian view is considered in greater detail in [Odlyzko2]. Since it is so contrarian, though, I devote some space here to justifying it (and presenting more examples).

There are frequently cited graphs showing faster diffusion of new technologies today than a century ago, say, such as those in [CoxA]. However, those comparisons have to be treated with caution. Yes,

the telephone, the automobile, and electricity did spread slowly, but then each had to build its own extensive infrastructure, and each one was very expensive in its first few decades. The Internet could take advantage of the existing telephone network to grow, and yet even the Internet did not really grow on “Internet time,” since its origins go back to the Arpanet, which was put into operation in 1969. For successful new consumer products or services that do not require large investments, a decade appears to be about the length of time it takes for wide penetration. This has been noted a long time ago.

A modern maxim says: “People tend to overestimate what can be done in one year and to underestimate what can be done in five or ten years.”

(footnote on p. 17 of [Licklider])

Arthur C. Clarke, the science fiction writer, is said to have similarly claimed that people tend to overestimate the short term impact of new technologies and to underestimate the long term impact.

Color TV took about a decade to reach 75% of the households in the U.S. It is not much different today. The paper [Odlyzko2] presents statistics on sales of recorded music in the U.S. by format. Music CDs are much better than vinyl LPs (at least for 99% of the population, as it has to be admitted that there is a small but influential segment that insists on the superiority of the older medium), yet it took them around a decade to attain dominance. Cell phones are all the rage, but they have been around since the middle 1980s, and yet by the end of 2000 were used by just about 40% of the population of the U.S.

The standard example of how things do move on “Internet time” is the browser. It did attain dominance in providing online access in well under two years. But that is just about the only such example of rapid change! Even on the Internet, technologies such as IPv6 and HTTP1.1 have been talked about as the “next big thing” for about half a dozen years, and are not yet dominant. Amazon.com did revolutionize retailing. However, it took quite a while, since it was established in November 1994, and 6 years later it had not yet taken even 10% of the U.S. book market. (Whether Amazon.com is viable in the long run or just an outstanding example of the “irrational exuberance” of the financial markets is another story.)

Much of the dot-com bubble appears to have been due to the expectations that the world was changing on Internet time. For example, in the middle of 2001, just before Webvan closed down, its new CEO was quoted as saying “We made the assumption that capital was endless, and demand was endless.” The idea of deliveries to the home may yet find a market and lead to financial success. However, Webvan was acting under the assumption that they had to build a giant distribution network

in a year or two, or else somebody else would. Instead, when demand was slow to materialize, they went bankrupt.

The entertainment area is full of examples of slow changes. The paper [Galbi] provides interesting statistics on a variety of subjects. Some of the most relevant for our purpose have to do with the slow rate at which people reallocate their time. For example, reading went from 4 hours per week to 3 hours, but it took from 1965 to 1995 to accomplish this.

More examples of slow consumer adoption rates are appearing all the time. For example, personal video recorders, such as TiVo and ReplayTV, have so far failed to take off, even though their users praise them highly [Hamilton].

The general conclusion is that we should not expect to see much change in consumer behavior as far as entertainment is concerned, at least not in less than 10 years. In particular, TV is likely to retain its format, and be delivered through TV sets, not PCs. In the meantime, the backbones of the Internet will be growing, to the stage where they will be capable of delivering all the TV content as separate streams for individual users even from a single central location. Since that mode of delivery is irrationally inefficient, it is unlikely to be employed, and so TV signals will not fill much of the Internet pipelines.

5. Streaming media versus store-and-replay

If Internet traffic continues doubling each year, where will the increases come from? There are some speculations in [CoffmanO2]. Video is likely to play an increasing role, taking over as a major driver of traffic growth from music (which got a large boost from Napster). However, this video is likely to be in the form of file transfers, not streaming real time traffic. There are more detailed arguments in [CoffmanO2], but the basic argument is that video will follow the example of Napster (or MP3, to be more precise), which is delivered primarily as files for local storage and replay, and not in streaming form. This local storage and replay model been known as a possibility for a long time, cf. [Owen]. It has several advantages. It can be deployed easily (no need to wait for the whole Internet to be upgraded to provide high quality transmission). It also allows for faster than real time transmission when networks acquire sufficient bandwidth. (This will allow for sampling and for easy transfer to portable storage units.)

The prediction that streaming multimedia traffic will not dominate the Internet has been made before, in [Odlyzko4, StArnaud]. It fits in well with the abundance of local storage we are increasingly experiencing.

6. Conclusions

The general conclusion is that the long distance Internet backbones are not going to be affected much by TV. Local “last mile” bottlenecks in data networks, as well as the slow adoption rates of new technologies by consumers, will ensure that by the time true convergence takes place between the Internet and entertainment TV, something on the order of a decade will have gone by. By that point, the backbones will have more than enough capacity to handle TV transmission. Even though it may be wasteful, it may then very well be less expensive to handle everything over the Internet, to avoid having several separate networks.

The Internet may very well have a larger impact on TV than TV will have on the Internet. The main advantage of the Internet has always been its flexibility, not its low cost. (See the discussions in [Coffman03, Odlyzko4].) The broadcast model, in which people have to adjust their schedules to fit those set by network executives was an unnatural one, forced by the limitations of the available technology. The popularity of video tape rentals showed that people preferred flexibility. Similarly, when cable TV operators chose to offer more channels as opposed to higher resolution channels, they were presumably responding to what they saw as their customers’ desires for variety.

The Internet will offer even more flexibility, but its impact is unlikely to be very rapid. Its main effect may be on high resolution video. HDTV has made practically no inroads because of the usual chicken-and-egg syndrome. Sets are expensive since there is no mass market, people do not buy sets since they are expensive and there is nothing novel to watch, stations do not carry HDTV programming since there is no audience, and so on. Internet allows for marketing to small groups. Studios already are making high resolution digital version of movies, and over the Internet will be able to reach the initially small groups of fans willing to pay extra for them. (This too will take time, not least because of fears of piracy.) Experiments with novel modes of presentation will also get a boost.

References

- [CoffmanO1] K. G. Coffman and A. M. Odlyzko, The size and growth rate of the Internet. *First Monday*, Oct. 1998, (<http://firstmonday.org/>). Also available at (<http://www.research.att.com/~amo>).
- [CoffmanO2] 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., Kluwer, 2001. To appear. Available at (<http://www.research.att.com/~amo>).
- [CoffmanO3] K. G. Coffman and A. M. Odlyzko, The growth of the Internet, *Optical Fiber Telecommunications IV*, I. P. Kaminow and T. Li, eds., Academic Press, 2001. To appear. Available at (<http://www.research.att.com/~amo>).
- [CoxA] W. M. Cox and R. Alm, Exhibit D (p. 14) in The economy at light speed: Technology and growth in the Information Age and beyond, Federal Reserve Bank of Dallas, 1996 Annual Report, available at (<http://www.dallasfed.org/hm/pubs/annual.html>).
- [Egan] B. L. Egan, *Information Superhighway Revisited: The Economics of Multimedia*, Artech House, 1996. Updates available at (<http://began.com/broad3.htm>).
- [Galbi] D. Galbi, Some economics of personal activity and implications for the digital economy, *First Monday*, vol. 6, no. 7, July 2001, (http://firstmonday.org/issues/issue6_7/galbi/).
- [Hamilton] D. P. Hamilton, TiVo, ReplayTV fail to take off despite big fans, *Wall Street Journal*, Feb. 7, 2001.
- [Lesk] M. Lesk, How much information is there in the world?, 1997 unpublished paper, available at (<http://www.lesk.com/mlesk/diglib.html>).
- [Licklider] J. C. R. Licklider, *Libraries of the Future*, MIT Press, 1965.
- [LymanV] P. Lyman and H. R. Varian, How much information?, report available at (<http://www.sims.berkeley.edu/how-much-info/>).
- [Noll2] A. M. Noll, *Highway of Dreams: A Critical Appraisal of the Communications Superhighway*, Lawrence Erlbaum Associates, 1997.

- [Noll4] A. M. Noll, Internet TV: Definition and prospects, to appear.
- [Noll5] A. M. Noll, TV over the Internet: Technological challenges, to appear.
- [Odlyzko2] A. M. Odlyzko, The slow evolution of electronic publishing, pp. 4–18 in *Electronic Publishing - New Models and Opportunities*, A. J. Meadows and F. Rowland, eds., ICC Press, 1997. Available at <http://www.research.att.com/~amo>.
- [Odlyzko4] A. M. Odlyzko, The Internet and other networks: Utilization rates and their implications, *Information Economics & Policy*, vol. 12 (2000), pp. 341-365. (Presented at the 1998 Telecommunications Policy Research Conference.) Also available at <http://www.research.att.com/~amo>.
- [Odlyzko5] A. M. Odlyzko, Content is not king, *First Monday*, vol. 6, no. 2, February 2001, http://firstmonday.org/issues/issue6_2/odlyzko/. Also available at <http://www.research.att.com/~amo>.
- [Odlyzko6] A. M. Odlyzko, The history of communications and its implications for the Internet, available at <http://www.research.att.com/~amo>.
- [Owen] B. M. Owen, *The Internet Challenge to Television*, Harvard Univ. Press, 1999.
- [Schaller] R. R. Schaller, Moore's law: Past, present, and future, *IEEE Spectrum*, vol. 34, no. 6, June 1997, pp. 52-59. Available through Spectrum online search at <http://www.spectrum.ieee.org>.
- [StArnaud] B. St. Arnaud, The future of the Internet is NOT multimedia, *Network World*, Nov. 1997. Available at <http://www.canarie.ca/~bstarn/publications.html>.
- [USDOC] U.S. Department of Commerce, Bureau of the Census, *Statistical Abstract of the United States*. Available online at <http://www.census.gov/prod/www/statistical-abstract-us.html>.
- [Vogel] H. L. Vogel, *Entertainment Industry Economics: A Guide for Financial Analysis*, 4-th ed., Cambridge Univ. Press, 1998.