

The delusions of net neutrality

Andrew Odlyzko

School of Mathematics, University of Minnesota

Minneapolis, MN 55455, USA

odlyzko@umn.edu

<http://www.dtc.umn.edu/~odlyzko>

Revised version, August 17, 2008

Abstract. Service providers argue that if net neutrality is not enforced, they will have sufficient incentives to build special high-quality channels that will take the Internet to the next level of its evolution. But what if they do get their wish, net neutrality is consigned to the dustbin, and they do build their new services, but nobody uses them? If the networks that are built are the ones that are publicly discussed, that is a likely prospect.

What service providers publicly promise to do, if they are given complete control of their networks, is to build special facilities for streaming movies. But there are two fatal defects to that promise. One is that movies are unlikely to offer all that much revenue. The other is that delivering movies in real-time streaming mode is the wrong solution, expensive and unnecessary.

If service providers are to derive significant revenues and profits by exploiting freedom from net neutrality limitations, they will need to engage in much more intrusive control of traffic than just provision of special channels for streaming movies.

1 Introduction

What if you build it and they don't come? That is what happened with the landline and underwater cables of the telecom bubble of a decade ago, and many other seemingly promising technologies. And that is almost bound to happen if net neutrality is blocked, and service providers do what they have been promising, namely build special facilities into their networks for streaming movies. The huge investments that supposedly can only be justified if non-neutral network management policies are allowed, are going to be wasted.

The public stance of the service providers, a stance that appears to be accepted as valid by the press, research community, and decision makers in government and industry, is based on two delusions. Both delusions are neatly captured in a single sentence by Jim Cicconi, one of AT&T's senior executives, made at the TelecomNext conference in March 2006 [19]. He said that net neutrality "is about streaming movies." The first delusion here is that movies are the most important material to be transmitted over the Internet, and will determine the future of data networking. But video, and more generally content (defined as material prepared by professionals for wide distribution, such as movies, music, newscasts, and so on), is not king, and has never been king. While content has frequently

dominated in terms of volume of traffic, connectivity has almost universally been valued much more highly and brought much higher revenues. Movies cannot be counted on to bring in anywhere near as much in revenues as voice services do today. This is discussed in a bit more detail in Section 2, but briefly, since this topic has been covered in detail elsewhere.

Even if we allow video the dominant role in shaping the future of the Internet, we have to cope with the second delusion captured in Cicconi’s quote, namely that movies should be streamed. This is an extremely widely shared assumption, even among networking researchers, as is discussed in Section 4. However, there is an argument that except for a very small fraction of traffic (primarily phone calls and videoconferencing), multimedia should be delivered as faster-than-real-time progressive downloads (transfer of segments of files, each segment sent faster-than-real-time, with potential pauses between segments). That is what is used by many P2P services, as well as YouTube. This approach leads to far simpler and less expensive networks than real-time streaming. And there is a noticeable minority of the technical community that regards this approach as the only sensible one. A truly astonishing phenomenon is that this group and the far larger streaming advocacy group do not seem to talk to each other, or even be aware that the other alternative is to be taken seriously. This is discussed in Section 3.

Section 4 outlines why the faster-than-real-time transmissions of video (or music) are the best solution, and require a far less expensive network that requires neither any fancy new technologies, nor any fancy new network management policies. Section 5 sketches out a particular view of the history and present state of data networks, which suggests a scenario of future evolution that supports the vision of faster-than-real-time multimedia transfers in preference to the streaming mode of the evolution of data networks

The general conclusion is that the story presented by service providers, that they need to block net neutrality in order to be able to afford to construct special features in their networks for streaming movies, is simply not credible. If lack of net neutrality requirements is to be exploited, it will have to be done through other, much more intrusive means. This is discussed in the conclusions section.

2 Content versus connectivity

The dogma of streaming video is a very damaging one, but it is certainly not the only damaging false myth in telecom. There are many others [14]. It is not even the most damaging. That position surely belongs to the “content is king” dogma. It stretches back for centuries, and has been consistently wrong for centuries, as is obvious to anyone who cares to look at the evidence, as is shown in [11, 12], or the more recent report [1]. Given all the details in those papers, I won’t devote much space to it, except to recapitulate some of the main points.

Although content has traditionally (almost invariably) been accorded special care by policy makers, people have always been willing to pay far more for connectivity. That video already dominates in terms of the volume of traffic on the Internet is not a counterargument. Almost two centuries ago, newspapers (the main “content” of the day) also dominated the

traffic carried by postal services, accounting for about 95% of the weight. But at the same time, newspapers provided only about 15% of the postal revenues (p. 38 of [5]). What people really cared about, and were willing to pay top dollar for, was connectivity, in the form of first class mail for business and social purposes. Content (newspapers in that case) is what the federal government decided should be subsidized for public policy regions with the profits from first class mail.

For all the hoopla about Hollywood, all the movie theater ticket sales and all the DVD sales in the U.S. for a full year do not come amount to even one month of the revenues of the telecom industry. And those telecom revenues are still over 70% based on voice, definitely a connectivity service. In wireless, there is very rapid growth in data service revenues, but most of those revenues are from texting, another connectivity service (and one that the industry did not design, but stumbled into).

Yet the "content is king" dogma appears to be impossible to shake. It deludes academics as well as government and industry leaders. For example, almost all the scholarly papers on net neutrality (see [15] for some references) model the Internet as a content delivery mechanism. And the new head of Time Warner, is planning to spin off its cable operations in "an effort to focus more sharply on 'content creation' (or what nonsuits still like to call movies and television shows)" [2]. Yet in the current Time Warner conglomerate, "cable networks have much higher margins" than the 'content creation' pieces [2]. So the move represents, in Samuel Johnson's words, "a triumph of hope over experience." Now in the context that Johnson made his quip, hope does triumph over experience with a reasonably high frequency. In the content versus connectivity area, though, the chances of success are far slimmer. And let us note that cable industry margins, even though higher than those of movie making, are not all that high. As one recent Wall Street report put it, "video is inherently a *much* lower margin product than is voice or data to begin with" (p. 24 of [7]). (Among other things, cable operators spend about 40% of their revenues acquiring the content they sell, p. 25 of [7], an obvious point that somehow is missing from most of the discussions of the wonders of a content distribution business. In the voice telephony and Internet access business, no content is needed, users fill the pipes themselves.) Occasionally the collision with reality is painful enough that people wake up. For example, in a presentation by Takeshi Natsuno, "one of the principal architects behind DoCoMos wildly successful 1999 launch of i-mode" (which is commonly regarded as a pioneering and successful content service, although texting was key to its success), "one message became abundantly clear: content is not king" [18]. But that message is awfully slow to spread, and we can be confident, based on all the historical precedents, that content will continue to get disproportionate attention. And we can also be confident that content will not be a gold mine, and will not bring in enough money to pay for super-expensive new networks.

Now there is a serious argument that new high capacity networks are not all that expensive. See [15], for example, or note that the cable industry did manage to build their networks on the basis of movie distribution, and that arguments have been made that the costs of upgrading those networks to higher speeds are not all that high. But the industry argues otherwise, that the costs are astronomical, and since this paper examines only the

plausibility of their claims about video, it accepts the (almost certainly false) premise that costs are very high.

3 Two video transmission approaches and their advocates

The next section will explain why faster-than-real-time progressive downloads of music or video are far preferable to real-time streaming. But first let consider the strange situation in which this issue is not discussed publicly, and the advocates of each of the two types of video transmission mostly seem unaware there is a real alternative to their preferred solution, and that there is a serious decision that has to be made for each video service.

That faster-than-real-time downloads have compelling advantages is not a new observation. It has been made many time before, and apparently independently by many people. (Two decade-old papers on this are [9, 10]. Note that slower-than-real-time downloads have been discussed much more widely, typically as a solution for inadequate bandwidth, with videos downloaded overnight, say, for future viewing. The first printed reference to transmitting movies as files that I have seen was by George Gilder in the early 1990s, but it is not clear whether he meant transmissions that are faster-than-real-time.) But the issue does not seem to have hit public attention. A few years ago, I heard a distinguished computer scientist say that he saw no point in transmitting video faster than real-time. This prompted me to start a series of informal polls at my networking presentations. I have been asking listeners to raise their hands if they saw any point at all in faster-than-real-time transmission of multimedia. I always explain very carefully that I mean this in a very broad sense, not whether this leads to viable business models, or anything specific, but just whether the audience sees any point, from the standpoint of someone, whether it be a residential user, or service provider, or content delivery agent, in using this technique. The highest positive response rate I have observed was at a networking seminar at the Royal Institute of Technology in Stockholm, in September 2007. It was about 30%. Twice, at networking seminars at CMU and Stanford, the rate was about 20%. Usually it is far lower, almost always under 10%. And sometimes it is close to zero. I had two similar audiences, on two separate continents, of about 100 people in each case, consisting of (mostly non-technical) mid-level telecom managers as well as government research agency staff and others connected with communications, where among the approximately 200 attendees in all, just one hand went up, and that one very tentatively.

In discussions with individuals, advocates of streaming seem generally to be unaware that there is any alternative. On the other hand, advocates of faster-than-real-time file transfers are aware of streaming, but generally regard it as a bizarre aberration. How this mutual misunderstanding could have persisted for years without a public debate is a real mystery. It is especially strange because of the very wide use of faster-than-real-time transmission.

Devotees of streaming are also almost uniformly astounded and disbelieving when told that most of the multimedia traffic on the Internet consists of faster-than-real-time file transfers. But that has been the case at least since Napster appeared. In those days, music MP3 files were typically 128 Kbps and perhaps occasionally 192 Kbps, but were usually moved around at megabit speeds. And today, when video on the Internet is still often

under 0.5 Mbps, and seldom more than 2 Mbps, transmission speeds are usually higher than that. Moreover, many services, such as YouTube, which appear to do streaming, are in fact using progressive downloads with faster-than-real-time transfers. There does appear to be growth in traffic that is truly real-time streaming, but it still forms a small fraction of the total. So faster-than-real-time transmission is used widely, but people who use it are mostly not aware of what is happening.

4 Dreaming of streaming

The press is full of claims that video will require a complete transformation of the Internet. As just one example, a 2005 news report [3] said that “Mr. Chambers [CEO of Cisco] predicts the demands of video will transform the Internet over the next decade,” and quoted Chambers directly as saying that “[m]aking [video] work is really, really, really difficult.”

But making video work on the Internet is not difficult at all, as many services (such as YouTube, for example) have demonstrated. One just has to do it properly. The story [3] does not make it clear why Chambers thought that video is difficult, but it is basically certain that this was due to the assumption that video over the Internet has to be delivered the way it is over the air or over cable TV, namely in real-time streaming mode. And indeed, if one is to use that approach, “making video work is really, really, really difficult.” One has to assure that packet loss rates, latency, and jitter are all low, and that is hard in a “best-effort” network based on statistical multiplexing. Consider the 2006 story [16] about the AT&T U-verse IPTV service, in which Microsoft was supplying most of the software:

Word has it the U-verse network loses roughly two packets of data per minute. ...
For the viewing public that can mean little annoyances like screen pixelation and jitter – or, at worst, full screen freezes.

...

One source close to the situation says Microsoft has already built in a 15 to 30 second delay to live video streams to allow some time for dealing with packet loss. AT&T, the source says, is uneasy about the scalability of the setup. Microsoft TV Edition product manager Jim Baldwin says his company’s middleware platform adds roughly a quarter of a second delay for packet error correction and another second of delay for instant channel changing, but that’s it.

So here is a system that was developed and deployed at tremendous cost in order to provide live streaming, and yet it has to introduce large delays, delays that eliminate the “live” from “live streaming.” And yet acceptance of far smaller delays would make far simpler solutions possible.

It is not just corporate CEO’s interested in selling fancy expensive new gear that assume video over the Internet has to be delivered in real-time streaming mode. Networking researchers also widely share this view. Let us just consider Simon Lam’s acceptance speech for the 2004 annual ACM SIGCOMM Award for lifetime technical achievement in data communications [6]. He called for a redesign of the Internet, but unlike Chambers, was explicit in his reasoning and recommendations. In particular (slides 9 and 14 of [6]) his

concern was that voice and video traffic would dominate on the Internet. Since (in his vision) such traffic would use UDP, which “does not perform congestion control,” and is “preferred by voice and video applications,” the Internet would be subject to congestion collapse. So Lam recommended use of a flow-oriented service (slides 16 and 28). The basic assumption underlying Lam’s argument, though, was that in the absence of a redesign of the Internet, video would use UDP, instead of TCP, which acts cooperatively in the presence of congestion. And indeed UDP is the preferred method for delivering real-time streaming. But the question that Lam did not ask is whether it makes sense to deliver video in that mode.

There are certainly services, such as voice calls and video conferencing, where human interaction is involved, and there real-time streaming, or a close approximation to it, is required. People are very sensitive to degradation in quality, and there are careful studies, going back decades, of how much latency, for example, they are willing to tolerate. However, voice calls, although they still provide the bulk of the revenue for telecom service providers, are now only a small and rapidly shrinking, though still noticeable, fraction of the traffic. On the other hand, video conferencing is growing rapidly, but is small, and there is no evidence it will ever be very large. (Video telephony falls in the same category as video conferencing, and there we have several decades of market experience, as well as more careful human usability studies, which show that the attractiveness of this service is limited, so we should not expect it to generate a huge amount of traffic.)

The vast bulk of video that is consumed by people today, and is likely to be consumed in the future, does not require real-time transmission. Most of it is movies and video clips that are pre-recorded, and thus can easily tolerate a certain amount of buffering. Even many apparently real-time transmissions are actually delayed. (For example, in the U.S., after the Janet Jackson episode, networks have been delaying broadcasts of events by several seconds, in order to be able to intervene and block objectionable images from being seen.)

Every time even a small delay can be tolerated, progressive faster-than-real-time transfers are the preferred solution. Consider a situation in which a standard-resolution movie, of about 2 Mbps (standard for today’s TV) is to be transmitted, but that the viewer has a 10 Mbps link to the source (very low speed in places like Korea or Japan, although still high for the U.S. as of mid-2008). If a 1-second delay can be tolerated, during that second, 5 seconds’ worth of the signal can be sent to a local buffer. Even if there is some congestion on the link, one can usually count on being able to transmit at least 3 seconds’ worth during that one second. And if one has 3 seconds’ worth of signal in a buffer, one can play the movie for 3 seconds with perfect fidelity, even if there is a complete network outage.

But in practice, there is no need for even a 1-second delay. As is done in YouTube and other services, one can start playing the movie right away, from the buffer, as that buffer gets filled. After one second, if 3 seconds’ worth of signal has been received, 1 second of it will have been displayed, but there will be enough to play the next 2 seconds of the movie. And by the end of those 2 seconds, almost surely at least an additional 6 seconds’ worth of signal will have been received. And so on. In cases where the signal is live, but delayed, say by 1 second, the computation is slightly different, but the buffering and faster-than-real-time transmission allow for easy compensation for any packet losses or jitter. One can

build theoretical models and do simulations, as is done in [17], for example, to see what kind of performance one obtains even with current TCP versions, without any changes to the Internet. The bottom line is that if one has a transmission link with capacity higher than the live signal rate, and one can tolerate some buffering, then using the standard TCP that handles congestion well works to provide high quality. There is no need for any fancy new technologies. And almost universally (unlike the old broadcast and phone networks, where the signal speed was exactly matched to the channel bandwidth, and which have clouded even expert visions, as is discussed in the next section), we do have higher network speeds than the signal, and today we have plentiful local storage.

Why did the Swedish audience (which, like those at CMU and Stanford, consisted largely of graduate students and faculty), show the relatively high recognition of the advantages of faster-than-real-time file transfers? Many of them had been working on projects in wireless communication, in sensor and ad-hoc networks. Hence they were forced to face the problem of intermittent connectivity, as nodes move out of range, or face interference. The obvious solution in such situations is to transmit at maximal feasible rates while there is a connection. And once they adopted this natural solution for their situation, it was obvious to them that it was also the best solution for wireline communications even when there is constant connectivity. And that seems to be the common pattern, that when technical people are faced with the task of delivering video economically, they usually reinvent faster-than-real-time progressive downloads.

Lots of reasons have been tossed around for real-time streaming. But none of them are persuasive. For example:

- Interrupted transmissions: Evidence shows that most videos are not watched in their entirety. So why download a complete video if only a quarter of it will be enjoyed by the customer? But of course there is no need to download the entire video, one can set limits on how much material will be stored at any time in the buffer. Faster-than-real-time progressive download often already do precisely that.
- Security: Streaming does not leave the video on the customer's equipment. This apparently makes content providers feel safer, as leaving movies in buffers appears to invite attackers to crack their protection schemes. But the degree of protection depends only on the security of the cryptographic algorithms and protocols. Attackers sophisticated enough to break those would have no problem intercepting a signal that is being streamed. And certainly the contents of the buffer can be encrypted.
- ...
- And last, but not least, a reason that is usually not explicitly mentioned, but likely provides a large part of the motivation for real-time streaming: This technique requires complicated and expensive gear from system providers, and justifies high prices and a high degree of control over traffic by service providers. That may be the most persuasive reason for streaming, but of course it only makes sense for those interested in high costs, not for network users.

On the other side, faster-than-real-time file transfers offer advantages beyond simpler and less expensive networks, advantages for both users and service providers. They lead to new services and stimulate demand for higher speed links. Suppose that you have just 5

minutes before you have to rush out of the house to catch a train or a taxi to the airport, and you plan to watch a movie on your laptop or other portable device during the trip. If it is a standard resolution 2 Mbps movie, and you have a 5 Mbps connection, there is no way you can download it during those 5 minutes. And of course there is no way to download it in 5 minutes if your service provider only lets you do real-time streaming of movies. But if you have a 50 Mbps connection, and the content provider allows it, you can get the movie to your portable device in those 5 minutes. And if you are really impatient, and want to download that movie in under a minute, you may be induced to pay for a 500 Mbps connection. This will be discussed in more detail in the next section. For the time being, though, the basic conclusion to be drawn from the discussion is that faster-than-real-time file transfers are a far more sensible way to move movies and video clips than real-time streaming.

Of course there is, and will continue to be, some truly real-time traffic, such as voice telephony. But that type of traffic should not be expected to occupy too much of the capacity of future networks (see [9, 10], for example), and there are various ways to accommodate it inexpensively without building special networks. The two key elements are the relatively slow rate of growth in resolutions of display devices, and the far faster rate of growth in transmission capacity.

5 Data networks and human impatience

Why do we have the widespread dogma of real-time streaming video? It appears to be inherited from the two old networks that have dominated imagination and discussion among the public as well as experts. One was the broadcast network, the other was the voice phone network. The traditional voice network does have a real-time requirement, as people do not tolerate substantial latency. For broadcast (radio or video) this requirement did not exist (except for things such as call-in shows), but lack of storage meant that local buffering was not an option. And so both these networks grew to provide real-time streaming using technology that basically was engineered for constant bandwidth streams. And that mental image appears to have influenced packet data network designers.

But that is not how packet data networks developed. That should have been obvious from early on, from an examination of utilization rates. But amazingly enough, those rates were not tracked, and for a long time there reigned the myth that data networks were chronically congested. As recently as 1998, the Chairman of the IETF (Internet Engineering Task Force) who was also a top Cisco router expert expressed the belief that data networks were heavily utilized [13]. Yet a modest effort sufficed to show that in fact corporate wide area networks were run at utilization rates approaching those of local area networks, and even Internet backbones were run far below the utilization rates of the voice network [8, 13]. And further evidence has been accumulating ever since, so that it is now recognized that data networks are lightly loaded. (Residential broadband connections are generally run at under 2% utilization in the U.S., and a tenth of that in Japan. Backbones appear to be up to about 25% utilization, as slower growth has lessened the impact of the factors [13] that led to the lower rates observed a decade ago.) But the implications of this observation are still not absorbed. These “low utilization rates show that what matters to users is the

peak bandwidth, the ability to carry out transactions quickly, and not the ability to send many bits,” [9]. To put it bluntly,

The purpose of data networks is to satisfy human impatience.

That should not be surprising. The computer industry understands this. PCs are bought for their peak performance, which is used relatively rarely. And Google understands this, as it designs its systems to deliver search results in a fraction of a second. Human time is a very limited resource. In a nice phrase of George Gilder’s, “You waste that which is plentiful.” And today, computing, storage, and transmission (other than where it is controlled by telecom service providers) are plentiful. Human time is not.

Once one accepts that data networks exist to satisfy human impatience, many phenomena are easy to explain. For example, the communications industry has been laboring to construct scenarios for why residential users might ever want 100 Mbps connections. By adding up several HDTV channels and some other services, they came up with semi-plausible scenarios as to how such a demand might arise a decade in the future. Yet today, with practically no high-definition videos around, 100 Mbps links are sold in large numbers in places like Japan and Korea. And why not? With 100 Mbps, one can transmit a movie 10 times faster than with a 10 Mbps link. Now the utility of doing so is not unlimited, and it is probably best to think of the value of a link as proportional to the logarithm of the speed of the link, so a 10 bps link (such as the electric telegraph) might be worth 1, a 1 Mbps link might be worth 6, and a 100 Mbps link would come in at 7. But there is a value to higher speed, and there is no limit to what might not be demanded at some point in the future. This opens up new vistas for service providers. They do not have to worry about the next speed upgrade being the last one. And they can now segment the market by the speed of the connection (something they have been moving into, but slowly and reluctantly).

As another, even more concrete example, note that in home wireless networks, the industry there (not the usual telecommunications supplier industry, it should be noted) has successfully persuaded people to move from first-generation 11 Mbps 802.11b WiFi systems to 150+ Mbps ones (even though admittedly those are maximal speeds only, the only guarantee being that they will not be exceeded). There is no streaming traffic inside the home that requires anywhere near that speed. What the new systems enable is low transaction latency, to satisfy human impatience. What we are seeing evolve on home networks is what we have seen in corporate ones before, namely a variety of transmissions, often machine-to-machine, but all ultimately driven by human impatience. Some of those transmissions contain content, but hardly any of that content is, or needs to be, streamed. But the speeds are growing, even though they are far above what is needed for streaming today’s movies.

The natural scenario outlined above, of link speeds growing with advances in technology (assisted by proper marketing) leads to a continuation and even an extension of what we have seen for a long time, namely light utilization. In this environment, the faster-than-real-time progressive downloads are the natural solution for video delivery. Real-time streaming is a damaging dead-end.

Table 1. Value of bits: Price per megabyte of various services.

service	revenue per MB
wireless texting	\$1000.00
wireless voice	1.00
wireline voice	0.10
residential Internet	0.01
backbone Internet	0.0001

6 Conclusions

Service providers may very well believe their story about the need to avoid net neutrality in order to build networks that can stream movies. The two myths, that movies are a gold mine, and that they should be delivered in streaming mode, are very widely held. But at the same time, it seems clear that service providers are aware this is not even the most promising avenue to explore in search for new revenues and profits. They have been devoting a lot of attention to the potential of DPI (deep packet inspection). Now DPI is not needed if you believe that you cannot have a successful video service without special channels for streaming delivery. If you do believe that, then you just build a network in which you control access to those special features that enable quality streaming. On the other hand, you do need DPI in either of two situations:

- You want to prevent faster-than-real-time progressive downloads that provide low-cost alternative to your expensive service.
- You want to control low-bandwidth lucrative services that do not need the special video streaming features.

Communications service providers do have a problem. But it is not that of a flood of video. Instead, it is that of the erosion of their main revenue and profit source, namely voice. Voice is migrating to wireless. Second lines, and to an increasing extent, even primary landlines, are being abandoned. And voice is (with today’s technologies) a low-bandwidth service, that takes just a tiny fraction of the capacity that modern broadband links provide.

Table 1 (taken from [15]) shows some rough approximations to the revenues the industry derives from various services. A full understanding of the industry also requires looking at revenues and profits, but this one table already shows it is the low bandwidth services that are most lucrative. And although it is not in the table (since it is not offered by current telecom service providers), the basic Google search is also a very low bandwidth service.

In trying to face a future in which the very profitable voice of today is just an inexpensive service riding on top of a broadband link, it is very tempting to try to control current and future low bandwidth services. And to control those, you do need “walled gardens” and DPI. And to succeed in this strategy, you need to stop net neutrality. So far the actions of service providers are consistent with such a course of action. Should they succeed, they could gain new sources of revenues and profits, not just those that Google commands today, but additional ones that come from more intensive exploitation of customer data (see [15], for example).

Whether service providers should be allowed to pursue this strategy is another question. The aim of this paper was just to examine their claim that they need to defeat net neutrality to be able to build special networks for streaming video. And that claim is simply not credible, whether those service providers believe it themselves or not.

References

1. J. Alleman and P. Rappoport, "The future of communications in next generation networks," white paper for the 2007 ITU Workshop on The Future of Voice, available at <http://www.itu.int/osg/spu/ni/voice/papers/FoV-Alleman-Rappoport-Final.pdf>.
2. T. Arango, "Holy Cash Cow, Batman! Content is back," *New York Times*, Aug. 10, 2008.
3. M. Boslet, "Cisco girds to handle surge in Web video," *Wall Street Journal*, Dec. 8, 2005.
4. ITU Workshop on The Future of Voice, 2007. Presentations and background papers available at <http://www.itu.int/spu/voice>.
5. R. R. John, *Spreading the News: The American Postal System from Franklin to Morse*, Harvard Univ. Press, 1995.
6. S. S. Lam, "Back to the future part 4: The Internet," 2004 ACM SIGCOMM keynote, presentation deck available at <http://www.sigcomm.org/about/awards/sigcomm-awards/lam-sigcomm04.pdf>.
7. C. Moffett, M. W. Parker, and J. Rifkin, "Verizon (VZ): Project FiOS... Great for consumers, but what about investors?," Bernstein Research financial investment report, Jan. 14, 2008.
8. 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/doc/recent.html>.
9. 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/doc/globecom99.pdf>.
10. 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 TPRC. Available at <http://www.dtc.umn.edu/~odlyzko/doc/internet.rates.pdf>.
11. A. M. Odlyzko, "The history of communications and its implications for the Internet," 2000 unpublished manuscript, available at <http://www.dtc.umn.edu/~odlyzko/doc/history.communications0.pdf>.
12. A. M. Odlyzko, "Content is not king," *First Monday*, **6**, no. 2, February 2001, http://firstmonday.org/issues/issue6_2/odlyzko/.
13. A. M. Odlyzko, "Data networks are lightly utilized, and will stay that way," *Review of Network Economics*, vol. 2, no. 3, 2003. http://www.rnejournal.com/articles/andrew_final_sept03.pdf. Original 1998 preprint available at <http://www.dtc.umn.edu/~odlyzko/doc/network.utilization.pdf>.
14. A. M. Odlyzko, "Telecom dogmas and spectrum allocations," written for the *Wireless Unleashed* blog, 2004, <http://www.wirelessunleashed.com/>. Also available at <http://www.dtc.umn.edu/~odlyzko/doc/telecom.dogmas.spectrum.pdf>.

15. A. M. Odlyzko, “Network neutrality, search neutrality, and the never-ending conflict between efficiency and fairness in markets,” *Review of Network Economics*, to appear. Available at <http://www.dtc.umn.edu/~odlyzko/doc/net.neutrality.pdf>.
16. M. Sullivan, “AT&T still has IPTV ‘jitters’,” *Light Reading*, Aug. 11, 2006, http://www.lightreading.com/document.asp?doc_id=101056.
17. B. Wang, J. Kurose, P. Shenoy, and D. Towsley, “Multimedia streaming via TCP: An analytic performance study,” *Proc. ACM Multimedia Conference*, 2004. Available at ftp://gaia.cs.umass.edu/pub/bwang/Wang04_tcp_streaming.pdf.
18. B. Warner, “Are the mobile networks backing the wrong horse?,” *Times Online*, Oct. 17, 2007, http://technology.timesonline.co.uk/tol/news/tech_and_web/personal_tech/article2680101.ece.
19. C. Wilson, “TELECOMNEXT: Net neutrality a bogus debate,” *Telephony Online*, March 22, 2006. Available at http://telephonyonline.com/home/news/net_neutrality_debate_032206/.