

# The growth rate and the nature of Internet traffic

Andrew Odlyzko

School of Mathematics  
University of Minnesota  
Minneapolis, MN 55455, USA  
odlyzko@umn.edu

<http://www.dtc.umn.edu/~odlyzko>  
Revised version, October 3, 2015

**Abstract.** The growth rate of Internet traffic is continuing its long decline. This is occurring in spite of the existence of huge potential sources of additional demand. Possible reasons for this trend, its implications for the Web, and potential future developments are discussed.

## 1 Introduction

How does one measure the Internet? There are a multiplicity of studies. Many, such as the *Internet Trends* series by Mary Meeker, consider various versions of the wide “Internet economy,” ranging from growth in sales of Internet-connected devices to healthcare. As the Internet continues to diffuse more widely and deeply, and becomes an ever more essential infrastructure, it is becoming harder to draw the boundaries. We still talk of e-commerce. However, the growth and proliferation of Amazon.com physical facilities, and the increase in online activities by traditional retailers demonstrate we are already in a system in which online information collection and online transactions are essential and inseparable elements of that industry. That parallels what happened with earlier technologies. Long before the Internet was built, most commercial transactions involved the phone, but there was no discussion of t-commerce (with “t” standing for the telephone). In line with this trend of melding the communications infrastructure with the rest of the economy, the OECD replaced in 2015 its biannual series of *Communications Outlook* reports with a new series, *Digital Economy Outlook*, [12], which considers a considerably wider subject.

The rate of growth of Internet traffic was for a long time a useful indicator of the health of the entire Internet ecosystem, but is becoming less relevant. This growth rate has been decreasing for almost two decades. Even the growth rate in wireless data, which was extremely high in the last few years, shows clear signs of a decline. There is still rapid growth, but it is simply not at the rates observed earlier, or hoped for by many promoters of new technologies and business methods. The latest statistics are surveyed briefly in Section 2.

This decline, while it has some potentially painful implications for many service and systems providers in telecommunications, probably will not have seriously deleterious effects on the economy. Some observers attribute the slowdown in the rate of growth of traffic

to an “approaching maturity in Internet adoption” (cf. [12], p. 113) or similar phenomena. That is likely to be a significant factor, as in 2015, close to half of the world’s population already was accessing the Internet (most often through a mobile phone). Thus the number of users cannot grow fast for long. Furthermore, Moore’s Law, in its various manifestations, such as computing power of processors, number of transistors on a chip, cost of a transistor, and hard disk storage capacity, is experiencing a substantial slowdown, which lessens the growth in the computing and storage capacities. That reduces the pressure for greater communication volumes.

However, there is still plenty of room for growth in volume of traffic. There is far more broadcast video going over the air or over coaxial cable to people’s homes than there is Internet traffic. The volume of data stored on magnetic disks alone is large enough that it would take many months to transmit all of it through the backbones of the Internet at current traffic levels. Furthermore, there are many new applications on the horizon.

So why isn’t Internet traffic skyrocketing, given all these all these obvious sources of potential demand, and the fact that technology is advancing fast enough to allow for fast growth without increases in investment? It is impossible to be certain, since we are observing (very imperfectly, since we do not have good measures of traffic, and even less of the types of traffic) the outcomes of complex interactions of technology, service providers business plans, and user decisions. The factors that are likely to be very important in affecting traffic growth rates are

- the shift to wireless, with complicated implications
- content is not king, as connectivity has the highest value
- the strong negative correlation between value and volume of data traffic

These factors are discussed in Section 3.

The conclusion that is drawn in the final section from all the considerations earlier is that declining traffic growth rates are not necessarily a sign of stagnation. The preoccupation with video, especially with prerecorded entertainment video, which dominates the thinking of industry leaders, researchers, decision makers, the press, and the general public, is leading them all astray. The observed declines in traffic growth rates reflect largely declines in video traffic increases. However, much of the value of current and future applications comes from relatively low-bandwidth applications, and those can prosper in the current environment.

## 2 Internet traffic growth rates

Internet traffic growth was for a long time a very useful measure of the development of the digital economy. Its explosive growth in the mid-1990s demonstrated the attraction of the new communication technology. But its slowdown to what was then regarded as a “mere doubling every year” starting in 1997 provided early warning that the Internet bubble was bound to burst. However, the studies that identified the annual doubling of traffic, and subsequent ones, led to the prediction that this trend would continue for a considerable time, meaning a decade or even more [3,4]. This prediction was based on observed rates

of traffic growth, new application adoption, and the presence of large sources of additional traffic (primarily broadcast video and the contents of magnetic storage systems). Actual traffic trends falsified this conjecture, as the first decade of the 21st century witnessed a substantial slowdown, documented in [8] and [2]. One of the puzzling facts from the beginning was that Japan, which has been among the top countries in providing fiber connectivity to homes, and has very high bandwidth Internet access, has had low traffic growth rates, and currently has far lower traffic (on a per capita or per connection) basis than countries such as the United States [7]. The frequent predictions about “exafloods” overwhelming the networks that were frequent a decade ago have simply not come to pass. At the 20 to 30% per year growth rates that are observed today in industrialized countries, technology is advancing faster than demand, so there is no need for increasing the volume of investments, nor for the fine-grained traffic control schemes that are beloved by industry managers as well as researchers.

The arrival of smart phones ushered in, around 2010, a period of explosive growth in wireless data traffic (which is taken here to mean technologies such as 3G, 4G, and LTE, and excludes WiFi). There were again many cries about unsustainable trends, and demands for more spectrum (even though the most ambitious conceivable re-allocation of spectrum would have at most doubled the cellular bands, which would have accommodated only a year of the projected 100+% annual growth). However, in the last few years we have observed a marked deceleration of wireless data traffic growth. On a world-wide basis, the June 2015 Ericsson report [6] estimates there was about a 65% growth from 2013 to 2014, and projects 45% compound annual growth rates for the 2014 to 2020 period. Perhaps even more remarkably, U.S. wireless data traffic, which more than doubled from 2012 to 2013, increased just 26% from 2013 to 2014 [5]. This was a surprise to many observers, especially since there is still more than 10 times as much wireline Internet traffic as wireless traffic (cf. [6]).

Many factors may have contributed to the slowdown in wireless traffic growth. Service providers imposing tight data caps was surely one. But there were probably other significant ones. For example, mobile devices have to cope not just with limited transmission capacity, but also with small screens, battery limits, and the like. This may have led to changes of behavior not just of users, but also of app developers. They likely have been working on services that can function well with modest bandwidth. Since mobility is a great attraction by itself (see [11] for some statistics and references), this likely led to less attention being devoted to high-bandwidth wireline services, and contributed to the slowing of traffic growth in that area.

### 3 Content and connectivity

In telecommunications, content has never been king. Historically, whenever we can obtain reliable statistics, we find that people were spending more on connectivity than on content [10]. This is also true today. As just one example, global revenues of telecommunications service providers are about \$2 trillion per year, while advertising is only about one quarter of that. (For more precise figures and references, see [11].) In the U.S., which has a very high degree of cable TV penetration, the revenues of this industry from providing

Internet access and voice telephony are likely soon to surpass revenues from video. Furthermore, profits of Internet and voice services are already far higher than of video, since there is no need to pay for the expensive “content.”

One of the simplest ways to demonstrate the falsity of the “content is king” myth is by comparing the prices people are willing to pay for the extremely low-bandwidth texting versus high-bandwidth video. (See the table in [11].) Another is by comparing the revenues of the entire video industry with those of cellular operators just before the arrival of smart phones, when the wireless industry was being paid almost exclusively for carrying the low-bandwidth voice and texting services.

Yet throughout history, policy makers and telecom managers have almost universally been fixated on content. The dominant role of basic connectivity is rediscovered every once in a while (cf. [1]), but it seems never to garner the attention it deserves. This continues to lead the industry (including researchers) astray. The myth of content as king is especially damaging when it is combined with the dominant myth of streaming video. (Most of the video on the Internet is transmitted, and will surely continue to be transmitted, in the form of faster-than-real-time partial downloads, the most sensible form for handling all static material.) The dream of the industry, and the likely inspiration for many of the publicity campaigns about exaffloods of traffic overwhelming the networks, is precisely that myth of streaming real-time video. That would justify gold-plated networks with tight control over traffic, and could lead to higher revenues and profits.

The problem for the industry is that many of the valuable services that are either already growing or can be reasonably foreseen can prosper with only medium speeds, and with substantial transmission impairments. That is even true of much of telemedicine. Many of the IoT (“Internet of Things”) services can adjust their traffic demands, often trading off local processing and storage for transmission. And, since for the most part they rely on radio communication, they simply have to be resilient enough to function even in the presence of substantial communication breakdowns. Therefore any realistic service quality guarantees will have to be statistical, rather than absolute.

Much of the current preoccupation of telecom service providers with content can be explained away as following historical precedents, succumbing to the glamor of “content,” and so on. But there is likely another pressing reason that applies today. With connection speeds growing, and the ability to charge according to the value of traffic being constrained either directly by laws and regulations, or the fear of such, the industry is in a desperate search for ways not to be a “dumb pipe.” This search is almost certainly doomed to fail, since “dumb pipes” is precisely what society needs, and since the non-trivial skills of the telecom industry that are required to provide ubiquitous physical connectivity are far from those involved in generating novel services. But the industry is determined to fight against its natural role of being “dumb pipe” providers, with its implication of being thrown into a commodity market. Managers in this field see video distribution as a way to develop differentiating offerings that will keep users from switching. An analogy might be with restaurants that all offered the same main courses, and competed just on the basis of their desserts. Desserts are not as important as the basic meal, or even just bread and water,

but when dealing with an affluent and sated society, they can sway decisions of where to go for dinner.

The fixation with video means the telecom industry is concentrating too much on limiting user traffic. In many ways, the danger for the industry, especially in the wireline arena, is from too little traffic, not too much. The many debates as to whether users really need 100 Mbps connections, much less 1 Gbps ones, reveal lack of appreciation that burst capability is the main function of modern telecom, serving human impatience. Although pre-recorded video dominates in the volume of traffic, the future of the Net is likely to be bursts of traffic coming from cascades of interactions between computers reacting to human demands, cf. [9].

## 4 Conclusions

The declining growth rate of Internet traffic is puzzling. There are several sources of additional traffic that seems bound to move to the Net, and will dwarf what is visible right now. However, this decline reflects primarily the slow move of broadcast video to video-on-demand services and the growing dominance of wireless access. There is already plenty of capacity for the growing range of novel services that will determine the future of the Web.

## 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. Cisco, Visual Networking Index, regular series of reports, available at <http://www.cisco.com/c/en/us/solutions/service-provider/visual-networking-index-vni/index.html>.
3. K. G. Coffman and A. M. Odlyzko, “The size and growth rate of the Internet,” *First Monday*, vol. 3, no. 10, October 1998. Available at <http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/620/541>.
4. K. G. Coffman and A. M. Odlyzko, “Internet growth: Is there a “Moore’s Law” for data traffic?,” in *Handbook of Massive Data Sets*, J. Abello, P. M. Pardalos, and M. G. C. Resende, eds., Kluwer, 2002, pp. 47–93. Available online at <http://ssrn.com/abstract=236108>.
5. CTIA, The Wireless Association, annual wireless industry survey, available at <http://www.ctia.org/your-wireless-life/how-wireless-works/annual-wireless-industry-survey>.
6. Ericsson Mobility Report, regular series of reports, available at <http://www.ericsson.com/mobility-report>.
7. ISP Backbone Traffic in Japan, regularly updated statistics at <http://www.hongo.wide.ad.jp/InternetTraffic/data/>.
8. MINTS, Minnesota Internet Traffic Studies, <http://www.dtc.umn.edu/mints/>.

9. A. M. Odlyzko, “The current state and likely evolution of the Internet,” in *Proc. Globecom’99*, IEEE, 1999, pp. 1869–1875. Available online at <http://www.dtc.umn.edu/~odlyzko/doc/globecom99.pdf>.
10. A. M. Odlyzko, “Content is not king,” *First Monday*, vol. 6, no. 2, Feb. 2001, <http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/833/742>.
11. A. M. Odlyzko, “Will smart pricing finally take off?,” in *Smart Data Pricing*, S. Sen, C. Joe-Wong, S. Ha, and M. Chiang, eds., Wiley, 2014, pp. 3–33. Available online at <http://www.dtc.umn.edu/~odlyzko/doc/smart.pricing.pdf>.
12. OECD, *OECD Digital Economy Outlook 2015*, <http://www.oecd.org/internet/oecd-digital-economy-outlook-2015-9789264232440-en.htm>.