International Journal of Communication 6 (2012), 920–935 1932–8036/20120920

Information Capacity
IJoC

The Volume and Value of Information

ANDREW ODLYZKO
University of Minnesota

The measurement of the volume of information is fraught with difficulties. However, trends in availability and usage can be very revealing, as is shown by some examples. This article argues that considerations of the volume of information should not be divorced from those of the value of information. In communications, informative comparisons appear to be possible by classifying technologies in just a few dimensions, associated with the features of cost, speed, availability, and usability. An argument also is made that as a rough approximation, the value of information in terms of its volume is best thought of on a logarithmic scale. This approach provides a rough quantitative guide to the diminishing marginal utility delivered by the rapid progress in computing, storage, and communication technologies. It offers a partial explanation for slow uptake of ultra-high speed broadband and related phenomena.

Introduction

Information has been a concern since ancient times. And so has information overload, as the Bible says that "[o]f making many books there is no end, and much study wearies the body" (Ecclesiastes 12:12). Attempts to quantify information date to at least three centuries before the birth of Christ. Already at the Library of Alexandria, charged by the Ptolemies to collect the world's knowledge, estimates were carried out of what fraction of the serious written works of that period had been accumulated there. (See the introduction to this special section of this journal.)

The real flowering of the quantification movement came in the 19th century, when statistics and statistical data collection underwent rapid development and came to be the basis for public discussion and public policy making. International comparisons proliferated, based on statistics of newspaper circulation, book trade, sizes of libraries, and the volume of letters in the mail system. Later, numbers of telegrams and phone calls began to be collected and scrutinized avidly. The importance of information for the economy and for individual well-being was widely recognized. During the British Parliament hearings that led to the Penny Post reform of 1840, Dionysius Lardner, widely known for popularizing science, was

Andrew Odlyzko: odlyzko@umn.edu
Date submitted: 2012–02–19

asked, "Is a tax on bread or a tax on letters the most oppressive on the people?" His response was classic:[1] "I think a tax on bread is the worst; but postage is a tax on the bread of the mind."

The Penny Post reform of 1840 is far removed from us, but it will be cited extensively, along with another important British reform of that period, the drastic lowering of newspaper taxes in 1836, since it illustrates some of the key points of this paper. Perhaps the main one of these points is that volume is not value, and thus just comparing bits or other simple measures of information is seldom very insightful. (At a deeper level, this calls for distinguishing between the traditional categories of data, information, knowledge, and wisdom, something that this paper will bypass.)

There are advantages to bringing everything to a common denominator, such as bits for information. But this has to be done carefully. As an example, one often finds world energy consumption measured in tons of oil equivalent, or in tons of coal, or in terawatt-hours. That is reasonable when considering the potential for alternative energy sources, say, since there is considerable substitutability among oil, coal, and electricity in the economy. On the other hand, there are few discussions of the entire output of a country's or the world's economy in terms of tons of products. There is far too great a variety of goods produced, and looking at their total weight would not be too insightful, especially since it would be dominated by all the soil and rocks moved in mining and construction projects (or coal and oil, if we restrict to outputs of more direct utility). What we find far more frequently in evaluating economic performance is financial measurements, primarily gross domestic product. This leads to a comparison of disparate products and services by way of money valuations, which reflect values of those products and services. Even there, the technical difficulties are immense (such as compensating for inflation and technological change, and deciding whether to use export prices or purchasing power parity, for example) and occupy many experts. Such comparisons are more subjective and time–dependent than tonnage measurements, for example, since they ultimately are based on human value judgments. However, they are more relevant for comparisons of economic performance and for decision making.

Just as world output, when measured by weight, is dominated by low–value rocks or minerals, measurements of the volume of information are usually dominated by broadcast video. This can be seen in (Hilbert & Lopez, 2012) and other articles in this issue. But such video has low value, as is shown later in this paper. Therefore, it appears necessary to go at least one level deeper than a single quantitative measure of the volume of information and obtain some sense for the values attached to different types of information.

The methodological problems in measuring information have been known for a long time and are emphasized in all the contributions to this section. On the other hand, it has long been argued (e.g., Ito, 1981) that studying growth trends can be very insightful. Even when measurements of different types of information cannot be compared directly, as long as each type is quantified consistently, the comparison of the growth rates for those categories over time can be very revealing, especially when a new technology appears on the scene. Several examples are presented later in this paper.

---

[1] *Select Committee on Postage, First Report, Minutes of Evidence, Appendix*. Parliamentary Papers 1837–1838 (278) XX Pt.I.1, Q5531.

In order to facilitate comparison of the value of different communication technologies, the paper (Odlyzko, 2003b) proposed placing them in a four-dimensional space, with the four axes measuring the following attributes:

- Volume: How much data can it transmit?
- Transaction latency: How long does it take to do something?
- Reach: Where can the service be provided?
- Price: How much does it cost?

This classification—which, it should be emphasized, is designed only for comparing communication methods and is not generally applicable to other information quantification cases—ignores some additional important dimensions, in particular reliability. However, it does appear to capture many of the most important features and provides insights into observations that are sometimes cited as paradoxical.

As an example, Amazon charges more for transmitting data into its cloud via the Internet than via physical storage devices that are shipped to it. But this is just a modern incarnation of the saying that apparently dates back to the 1970s: ". . . never underestimate the bandwidth of a station wagon loaded with magnetic tapes." Given the technology trends in storage and transmission, it has always been, and for the foreseeable future will continue to be, less expensive to transmit large data sets via physical media than over data networks. Thus, if short delays—small transaction latency—are not a critical requirement, postal services can provide inexpensive and ubiquitous first-mile connectivity. Some other examples where this classification serves to explain differential valuations of communication technologies are presented later, for example in the section on displacement of wired by wireless voice, and displacement of wireless voice by texting. (The example of texting substituting for wireless voice suggests that low latency is sometimes undesirable, so that the usual value proposition gets reversed!)

An observation that is reinforced by many of the examples in this paper is that even supposedly revolutionary changes—whether in technology or in business models—usually take far longer to make a perceptible effect on society than their enthusiasts claim. It already was noted by Licklider in 1965 that "[p]eople tend to overestimate what can be done in one year and to underestimate what can be done in five or 10 years," (Licklider, 1965, p. 17.) However, growth trends frequently change much faster and provide indications of how popular and important the new developments will be, and to what extent they will be substitutes versus complements for existing technologies or businesses.

Later sections consider several examples that illustrate the poor state of knowledge about the volumes of communications traffic through the ages. They also demonstrate that even in the absence of precise and reliable data, one often can get interesting insights into technology and economy trends.

Before considering those examples, the next section proposes a rough rule for evaluating the value of information as a function of the volume of this information. It appears to apply to computing and storage as well as to communication. This rule is that value is proportional to the logarithm of volume. It is motivated and justified by analogies to a variety of other areas where similar rules have been shown to hold. Applying this rule to Internet access, we find that the move from dial-access connectivity at 10 Kbps to an initial broadband capacity of 1 Mbps represents (taking decimal logarithms) a step from 4 to 6, an improvement by 50%. On the other hand, a move from the currently common 10 Mbps to 100 Mbps represents a step from 7 to 8, a far smaller gain of only 14%. This might help explain the relatively slow diffusion of high-speed broadband. In general, this rule pictures the progress of information and communication technologies as closer to linear than to the exponential (in the strict mathematical sense of the word) that raw performance numbers provide.

Value as Function of Volume

The essentials of a piece of information can often be distilled to just a few bits. Was it X or Y who won the election? But we usually want more. How many votes did the winner and loser receive? How many did they receive in various regions of the country, or from different sociodemographic groups? What were the victory and concession speeches like? The words that were spoken matter the most, even though transcribed, they use few bits. But the far greater volume of information in the actual spoken speech, and the far greater yet volume of the video, also are desirable. Generally, the historical record has been of people frequently complaining of information overload, but almost universally of asking for more information and often of being willing to pay for it. But there is clearly a diminishing marginal utility to the explosion of information, and the suggestion here is that as a rough approximation it should be valued on a logarithmic scale.

Perhaps the earliest application of a logarithmic scale to measuring measurement was the ancient Greek classification of stars by magnitudes. A single step along their scale corresponded to a doubling of brightness. Since that time, similar measures have been introduced in many fields. For example, the Richter earthquake magnitude scale was modeled explicitly on the stellar magnitude scale. In a similar vein, the decibel comes from an early measure of the smallest differences in audio levels that were perceptible to the human ear. These differences turned out to require roughly the same multiplicative factor increase in power over a large range of levels of loudness. More generally, the Weber–Fechner law states that the perceived intensity of various physical stimuli is proportional to the logarithm of the power of those stimuli. These "laws" have many exceptions and limitations. Still, their wide applicability suggests that they correspond to a fundamental fact about human perception and can be used as first approximations in many applications.

Once we move away from simple measures, such as perception of loudness of sounds or weight of objects, the best quantitative measures become exceedingly complex. Two samples of the voluminous literature on how the quality of sound or video depends on various factors are the papers (Barten, 1990; Jayant, Johnston, & Safranek, &1993). So a logarithmic valuation in volume of a file is a very crude approximation, but might not be inappropriate. It appears to correspond to observed phenomena. High-definition TV sets, as well as 3D-display technologies, have been popular but have not led to a rush to purchase them. Similarly, uptake of higher speeds for Internet access has been slow. These observations all point toward users perceiving some benefits, but not ones commensurate with the gains in technical performance as measured in MIPS, bytes, or Mbps.

Additional support for logarithmic valuations can be obtained from other areas. In economics, geography, sociology, and related areas we find observations such as Zipf's law, which states that if people are ranked by their wealth, or cities by their population, or many other measures, then the *k-th* one will be about *k* times smaller than the top one. As with the Weber–Fechner law, there are some heuristics about distributions of wealth and income, frequencies of words, and so forth, that lead to Zipf's law. But also, as with Weber–Fechner, there are counterexamples. Still, the frequency with which Zipf's law shows up suggests that it is a good first heuristic to use in the absence of more detailed knowledge. (And sometimes even when such detailed knowledge is present, provided it is not at wild variance with Zipf. One of the advantages of simple models, even when they are demonstrably false, is that they provide constructs that are easy to understand and manipulate. Of course, one has to be careful when doing this, as all models are at best approximations.)

Zipf's law and related observations lead to a first–order valuation of *n* times the logarithm of *n* for a communication network with *n* participants, instead of the popular square of *n* valuation of Metcalfe's law (Briscoe, Odlyzko, & Tilly, 2006; Odlyzko & Tilly, 2005). This modification of Metcalfe's law says that the average participant in such a network derives value from membership that is proportional to the logarithm of the size of that network.

Another derivation of this *n* log(*n*) valuation for a network in (Odlyzko & Tilly, 2005) relies on Bradford's law of scattering, which describes where researchers find the most useful information. Bradford's law also leads to a logarithmic valuation of a library collection as a function of its size.

A strong argument that something close to a logarithmic function is appropriate in valuing information comes from general economic considerations. If we consider the economy as a whole, it is growing perhaps around 5% per year for the world gross domestic product. This is very slow compared to the close to 50% per year growth rates observed in recent years for computing, storage, and communications capabilities. Yet those three sectors, while they have had periods of rapid growth, are largely stagnant in terms of revenues. Hence the value those sectors can extract from society is almost constant. Of course, one could conceive that the value they deliver is growing commensurately with their measured capacity (i.e., exponentially in the strict mathematical sense of the word), it's just that almost all of this value goes to purchasers (consumer surplus, in technical terms). However, we just don't find any signs that consumers are deriving giant benefits. Therefore, the value that the ICT (information and communication technologies) sector is delivering must be close to constant. To obtain something close to that when applied to the exponential growth of the computing, storage, and communication capacity, we have to apply something like a logarithmic function.[2]

[2] Even a purely logarithmic function can be manipulated to produce improved fit. The base of the logarithm can be changed. After all, natural, binary, and decimal logarithms are all legitimate logarithms. Next, the units of measurement can be changed. Should storage capacity be measured in bits, bytes, or gigabytes? With the choice of these two parameters one can change a particular logarithmic evaluation w by any linear transformation $aw + b$ for arbitrary constants a and b. This points out again that the suggested logarithmic valuation should be treated just as a rough approximation.

Zipf's law and the associated logarithmic valuation of information suggested above also provide a quantitative re–interpretation of the "Long Tail" theory of Chris Anderson (Anderson, 2006). If we have a billion, 10^9, objects that follow Zipf's law in their valuation, the value of the entire collection is about 9. However, the most valuable one thousand are worth about the logarithm of 10^3, which is 3, and the most valuable million are worth about 6. This suggests that there is validity to the "Long Tail" hypothesis, in that the large majority of items with low valuations are valuable in total, and can collectively outweigh the "big hits" with high valuations. However, the "big hits" are almost never negligible. This approach to the "Long Tail" theory is a very rough one, as in practice careful studies of various marketplaces show that Zipf's law is often not a good approximation to the actual valuations. Still, this approach provides a quick way to think of the extent to which "Long Tail" might apply in various situations, and how much depends on just where the long tail starts.

One can also use logarithmic valuations of technological capabilities to throw doubt on the validity of the singularity theory, which asserts that the exponential progress of technology is reaching a point where humanity will be drastically changed. Ray Kurzweil is the most prominent proponent of it (Kurzweil, 2005). For a variety of views on this theory, see the collection of articles (IEEE, 2008). However, if we apply a logarithmic evaluation to those technological capabilities, we find ourselves with a picture of linear progress, which appears more realistic.

Letters and Newspapers in the 19th Century

An interesting perspective on circulation of information is provided by Figure 1. It shows the number of newspapers published in Great Britain between 1827 and 1848.[3] The big jump in 1837 over 1835 was caused by the success of the extensive campaign against "taxes on knowledge," which led to a lowering of those high levies on newspapers. The main London dailies, such as *The Times*, lowered their prices from seven of the old English pence to five, and, perhaps even more important, increased the number of pages they published.[4]

[3] Numbers derived from *Return of Aggregate Number of Stamps issued for Newspapers in Great Britain and Ireland, 1827-1841*, Parliamentary Papers 1842 (412) XXVI.599, and *Return of Number of Newspaper Stamps issued, 1842-48; Return of Number of London, English Provincial, Scotch and Irish Newspapers; Number of Advertisements and Amount of Advertisement Duty Paid, 1841-1848*, Parliamentary Papers 1849 (160) XXX.349. The numbers in Figure 1 are just for Great Britain, excluding Ireland (which had around 5 million papers published each year). The population of Great Britain at that time was around 20 million, and of Ireland around 8 million.

[4] Even at five pence per copy, newspapers were extremely expensive. As a fraction of GDP per capita, five pence then is comparable to $50 for the U.S. today.

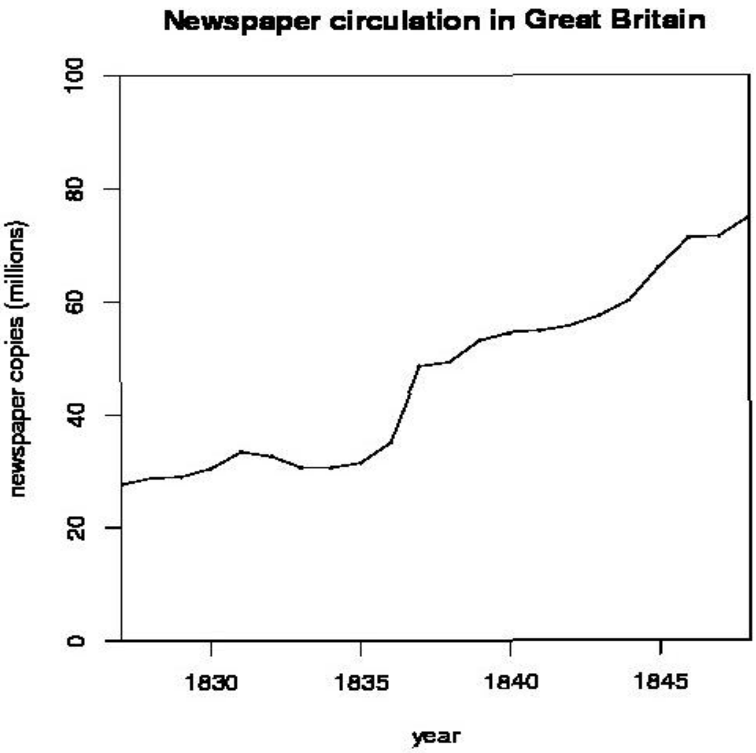

Figure 1. Newspapers printed in Great Britain from 1827 to 1848.

What can we learn about quantifying information from Figure 1? The number of newspapers shown in this figure is very reliable, since the heavy taxation led to careful record keeping. However, there were many publications that did not fall into the legally defined category of newspapers. They were not subject to taxation, and as a result of this and the habitual secrecy of companies about details of their operations, we have only vague estimates for their circulations. Further, Figure 1 does not reflect the general growth in sizes of newspapers. There was a noticeable jump in much of the press in late 1836, when taxes were lowered, but there were also other increases, which have not been measured.[5] Perhaps

[5] As an example, on January 11, 1844, the *Glasgow Argus* switched to a larger page size, following earlier enlargements in September 1836 (when newspaper taxes were lowered) and January 1839. A leader in that issue noted that as a result of this increase, it was delivering to its readers almost exactly twice as much information as in its first format introduced when the paper was established in 1833. It is worth noting that a precise measure of the growth of newspaper information production can in principle be obtained by tedious investigations in libraries.

the greatest uncertainty, though, is about the consumption of newspaper information. Because of high prices, newspapers often were rented by the hour, and read in (usually fee-supported) reading rooms, in pubs, and in coffee houses. Contemporaries often estimated that about 10 people read each paper, but they do not appear to have carried out any systematic studies.

Thus the example of British newspapers from the early 19th century illustrates many of the problems we face today. We have a few reliable statistics, but they do not translate naturally into reliable estimates for either the volume of information generated, or for the volume consumed.

Still, there are interesting trends visible in Figure 1. Perhaps the main one is that the decrease in newspaper taxes did lead to a quantum jump in circulation, but it was not a giant jump. On the other hand, this decrease was associated with a noticeable change, from an essentially static circulation (1.7% compounded annual growth rate between 1827 and 1835, only a bit faster than population growth) toward vigorous increase (5.3% rate between 1838 and 1848).

Similar observations can be made about postal services. Table 1 presents statistics on the letters sent through the British Post Office in the years 1839 through 1855.[6] Surprisingly, we do not have any solid counts on the number of letters before 1839. However, financial statistics show stagnant revenues and profits in the 1820s and 1830s, which strongly suggests that the volume of correspondence sent through the mails did not grow much then. This failure for postal revenues to increase at a time when other measures showed rapid growth in various forms of transportation and other economic activities was one of the main arguments used by Rowland Hill and his supporters in arguing for the Penny Post reform of 1840. They succeeded, and while the immediate jump in the number of letters did not meet their projections,[7] the basic statistics show a new growth trajectory, with growth at a compounded annual rate of 6.9% from 1840 to 1855. Thus we see the same phenomenon with letters as with newspapers: A substantial change in the business model shows its effects more through the long-term growth rates than in immediate jumps in volumes.

Early 19th century postal services show yet another, almost universal phenomenon of volumes dominated by what is called "content," namely material prepared by professionals for wide dissemination. We do not have statistics for the number of newspapers handled by the British Post Office until after 1855. However, around 1837 it was estimated that about 30 million newspaper copies were sent through the mails.[8] The volume of information carried by these 30 million newspapers dwarfed that of the perhaps 80 millions letters. However, it was the letters that were providing the bulk of the funding. A similar situation

---

[6] Statistics obtained from (Mitchell, 1992, p. 731). Unlike in the case of newspapers cited earlier, this time Ireland was included.

[7] Even the observed jump was somewhat misleading, since it almost certainly reflected a decline in illegal smuggling.

[8] Estimate of Rowland Hill, (Hill, 1837, pp. 10, 66). It is likely that this was an overestimate, as Hill's figure for the total number of newspapers printed that year was about 70 million, as opposed to the approximately 50 million shown in Figure 1.

could be observed on the other side of the Atlantic. In the U.S. Postal Service, "[i]n 1832, newspapers generated no more than 15 percent of total postal revenues, while making up as much as 95 percent of the weight" (John, 1995, p. 38). Thus simply counting bits does not give a good grasp on the value in this case, as in most other cases.

Table 1. First Class Letters Delivered by the British Postal Service.

Year	**Letters** (millions)
1839	82
1840	169
1841	196
1842	208
1843	219
1844	243
1845	272
1846	300
1847	321
1848	329
1849	338
1850	346
1851	361
1852	379
1853	411
1854	443
1855	457

The 19th century is a source of many interesting lessons about information. One such lesson that has occurred repeatedly through history is that services that had been thought to be substitutes for each other often grew together. For instance, some opponents of Rowland Hill's reform proposal argued that the steady volume of letters was a result of the spread of railways, which made transportation easier and mail correspondence less important. Instead, both snail mail and railway passenger travel boomed. Later, there were expectations that the electric telegraph would displace postal communication, but both grew over many decades.

Statistics on the volumes of information transmitted by postal and telegraph services are available for many countries in (Mitchell, 1992), although all that is available are figures for the numbers of letters and telegrams delivered. The manuscript (Odlyzko, 2000) also has volume and financial statistics on U.S. postal and telegraph industries. They show that the telegraph never exceeded one-fifth of the postal service in revenues, and was surely far smaller in terms of volumes of information carried. This provides some indication of how important the dimensions listed in the introduction were for individuals and businesses. Low latency, where the telegraph excelled, was not able to overcome disadvantages in volume, reach, and price.

Wireless Volume Declines and the Value of a Few Bits

An interesting tradeoff between different communication technologies is being made as this paper is being written. The press is full of dire predictions of the wireless data tsunami, of spectrum shortages, and so on. But it is an intriguing observation that the most important part of wireless communication—as judged by what people are willing to pay for—is actually shrinking in volume in the U.S.

Table 2. Voice and Texting in the U.S. Wireless Industry

(minutes of voice calls and texts each month by the average user).

Year	Voice Minutes	Texts
2005	641	35
2006	682	60
2007	725	124
2008	699	319
2009	685	471
2010	638	584

Table 2 shows that wireless voice usage is declining, while texting (SMS) is growing rapidly.[9] The number of wireless subscribers grew from 195 million in mid-2005 to 293 million in mid-2010. Therefore, the total volume of wireless voice calls actually grew by more than 50% from 2005 to 2009. While we do not have reliable recent figures for wireline voice usage, all indicators—such as the growth in numbers of households without wireline service—point to a decline, with a substantial fraction of that decline representing substitution by wireless. Thus, users willingly move from wireline voice, with its far superior quality at 64 Kbps, to considerably lower quality wireless, which is estimated at around 8 Kbps. The magnitude of that switch on the volume of information that is transmitted is hard to estimate. However, in wireless, the total volume of voice calls declined from 2009 to 2010—by about 1.5%—while the number of users grew by about 5.8%. So it appears safe to conclude from Table 2 that texting was substituting for voice. Thus a minute of voice, which can be taken to be approximately 120 KB, was being replaced by a few texts of just 120 or so bytes.[10]

[9] Data in Table 2 is derived from (CTIA, 2011). Minutes of use and texts, which are provided only for full calendar years, were divided by the number of subscribers given for the middle of the corresponding year.

[10] Since both sides of a wireless call are allocated exclusively to a customer during a voice call, a minute can be taken as equivalent to two transfers of 60 KB, one in each direction. If we take the 47-minute decline from 2009 to 2010 as being replaced exactly by 113 texts—the increase from 2009 to 2010—we find that a minute of voice is equivalent to about 2.4 texts, which is likely considerably under 120 bytes.

The decline of voice volumes as users move from wireline to wireless is surely the result of the overwhelming advantage of mobility—reach in the dimensions of the introduction. It could easily be reversed in the next few years as wireless capacity grows and better quality wireless voice becomes widely available. A more intriguing case is that of wireless voice giving way to texting. Is this due to the obvious reason, namely the advantage of asynchronous communication—which made e-mail the "killer app" of the Internet? (That is one case where low-transaction latency, which is generally desirable, plays a negative role, in that a phone call is intrusive and does not give the caller a choice of when to respond. An interesting question is whether the communication industry could have developed a popular and lucrative service by providing easy voice messaging, in which people could easily record and transmit to the recipient's voice mail box their voice messages. Human culture is largely oral, after all, and the huge revenues of voice telephony testify to the utility of voice communication. Voice messaging could provide a nice hybrid between the traditional telephony and email/texting, one that even illiterate people could use. Some voice messaging is taking place, but in clumsy forms, and what is needed is something really simple, such as a button on a handset that a person could press to ensure that after the called party number is dialed, a message is recorded and deposited directly in a voice mail system.) Or are there other social communication changes that are playing a role? However, the advantage of mobility that wireless offers is undeniable and may have implications for Internet traffic growth rates.

Internet

The author's studies of Internet traffic volumes and growth rates started in 1997. Initially they were carried out jointly with Kerry Coffman at AT&T Labs Research, and were documented in (Coffman & Odlyzko, 1998, 2002a, 2002b). They were later continued at the University of Minnesota through the Minnesota Internet Traffic Studies project, MINTS, and the latest detailed paper is (Odlyzko, 2003a). The many inadequacies of the available information sources, and the resulting uncertainties in Internet traffic estimates, are documented in those papers. It is worth noting that these and essentially all studies of communication volumes measure the total volumes of customer-useful data that are transmitted, not the smaller volumes that could be obtained by compression, as in (Hilbert & Lopez, 2012).

As of late 2011, the most thorough and most widely cited studies of Internet traffic and of growth projections for the future come from Cisco's Visual Networking Index project (Cisco, 2011). (For international traffic, the data from Telegeography is also widely used, although only brief summaries are made public.) The latest report, from June 1, 2011, estimates the world Internet traffic in 2011 at 20,650 PB/month. This is consistent with the MINTS estimate for year-end 2011 of 15,000 to 25,000 PB/month.

All Internet traffic estimates have to be treated as extremely rough. Even the Cisco studies are based on reports from only 20 service providers, with no information on which ones are involved. The MINTS estimates are based on collections of data from various sites, most prominently Internet exchanges that make their traffic statistics publicly available. These sites are fully identified and can be accessed by anyone. However, that covers only a fraction of the world traffic—and a very small fraction of North American traffic, in particular. Hence, to form an estimate of total traffic, MINTS also has to rely on confidential reports from some service providers. Thus there is no estimate available that is based just on publicly available and verifiable information.

The deficiencies in traffic information are alleviated slightly by some service providers occasionally releasing data on their traffic. There also are two government traffic collection efforts in Australia and Hong Kong that provide valuable information, as well as a collaborative ISP industry project in Japan, that covers a substantial fraction of that country's market.

We do have solid data from the early 1990s. The NSF Internet backbone was well instrumented and published complete statistics. Those did not include various types of traffic that remained on local networks, but in most experts' views, the NSF backbone carried the overwhelming majority of Internet traffic. Growth on this backbone was a very regular doubling each year, from 0.001 PB/month in 1990, to 0.016 PB/month in 1994.

Once the NSF backbone was replaced by private networks, reliable information became scarce as carriers refused to provide any reports, and growth rates jumped. The estimate of (Coffman & Odlyzko, 1998)—which included estimates for other types of communication traffic as well—was that in both 1995 and 1996, growth was close to 1000% per year, the proverbial "doubling of Internet traffic every 100 days." However, already by early 1997, there was a dramatic slowdown in North America, with growth declining to about a doubling each year. (Recent investigations, not documented before, based on discussions with various network engineers who were at large service providers in 1997, showed that, to a large extent, this slowdown was initially caused by capacity constraints.) This doubling each year continued until around 2002 and 2003, when growth slowed to about 50% per year, and more recently dipped a bit lower, to the range of 40%–50% per year. The year-end 2011 estimate from MINTS for North American Internet traffic is in the range of 3,500 to 5,500 PB/month, which is somewhat lower than the Cisco estimate (Cisco, 2011).

Even rough estimates of total traffic volumes and growth rates can be useful. The Internet bubble was inspired and maintained to a substantial degree by the myth of "Internet traffic doubling every 100 days." This myth did have a kernel of truth, as such growth rates appear to have prevailed in the U.S. in 1995–1996. It also was not totally inconceivable that they could have been maintained a couple of years longer. Computing power and storage had been growing rapidly for some decades before the Internet burst on the public consciousness with the Netscape IPO in 1995. On the other hand, long-distance data communications had been a laggard. Hence it was not unreasonable to think that the appearance of the Internet would lead to a catching-up period, when the disconnected computing and storage systems started exchanging information. However, the observations reported in (Coffman & Odlyzko, 1998) showed that even in situations with no communication link constraints, traffic tended to grow at most at single doubling each year. (This was later confirmed in many other situations. For example, in countries such as South Korea and Hong Kong, which deployed broadband rapidly, there usually was a year or two of very rapid growth, and then a decline to rates below 100% per year.) Thus the predictions of a rosy future for many dot–com and networking ventures were bound to be disappointed, as were the expectations that data would be sent around the world on a large scale for processing. Still, the persistence of 100% growth rates in the late 1990s led to the conjecture (Coffman & Odlyzko, 2002a, 2002b) that communication was catching up to computing and storage, even if not as rapidly as popularly thought, and that such growth rates might persist for quite a while. This conjecture has been disproved by the experience of the last decade. Internet traffic growth rates have declined to the 40%–50% range. At that level, they are very close to the growth rates of magnetic storage and of computing power.

Content and Connectivity

An earlier section noted that in the 19th century, "content," in this case newspapers, dominated in measures of volume of information transmitted by postal services, but the value was in the relatively small information carried by first-class letters. This imbalance continues today, part of an ancient tradition of overvaluing content at the expense of connectivity (Odlyzko, 2001). This tradition continues in spite of all the contradictory evidence that has been accumulating. The U.S. wireless statistics, cited earlier, (CTIA, 2011), show that in 2010, this industry brought in revenues of about $160 billion, of which $110 billion was for low-bandwidth voice services. Of the remaining $50 billion, it is estimated that about half came from texting, which, as was already noted, involves extremely low volumes. By contrast, the U.S cable industry had total revenues from residential video of just $53 billion in 2009 (NCTA, 2010), in spite of the high volumes transmitted. This provides yet more evidence that connectivity, and not content, is king. It also provides more support for the suggestion that there might be high value in stimulating the use of videos for social communication, of which YouTube might be just an early example. Such use might be reflected in greater upload volumes from residential subscribers—although not necessarily so, since the communication of the high-volume objects might take place in the cloud, subject to low-volume commands on the uplink from subscribers.

Conclusions

Communications volumes are very hard to measure. There are inherent problems, many common to other areas of information measurement. In addition, service providers refuse to provide statistics, and regulators in almost all cases have decided not to insist on traffic reporting. As a result, we are left with an unsatisfactory state of knowledge.

Still, even with our current inadequate approaches, we should be able to track large changes in the patterns of information flows. There is a continuing evolution of the Internet, with business deals leading to new traffic patterns, as reported in (Labovitz, Iekel-Johnson, McPerson, Oberheide, Jahanian & Karir, 2009), as well as the appearance of new services, like Netflix delivering movies online. Going forward, what developments might take place, and to what extent can they be seen in their early stages?

There have been many warnings of an impending "exaflood" that would swamp networks. They have not taken place, and instead we have seen a deceleration in wireline Internet traffic growth rates. This decline could persist, as Cisco (Cisco, 2011) has been predicting. (Japan, with very high bandwidth residential connections, has been experiencing such low growth rates over the last few years, cf. (MINTS).) The Cisco methodology is based on detailed projections of traffic generated by various devices. But even from a high-level view, one could imagine such a slowdown coming from changes in Moore's law in several fields. Processor clocks are not getting much faster, so the additional computing power is coming from putting more units on a chip, with the result that the communication bandwidth per unit of computing power is declining. In storage, too, even if areal density continues to increase at the current rate—which some skeptical voices are questioning—the bandwidth from the disk may well grow more slowly. Thus we may end up with data becoming "cooler," in database language, not being used as extensively and not traveling as much.

Yet another reason we may see a decline in traffic growth on the Internet is the one discussed earlier on substitution of wireless for wireline voice, and texting for wireless voice. The advantage of mobility may lead to more attention being placed on wireless, where the bandwidth is, and likely will continue to be, far lower than with fiber optics. At the end of 2011, mobile data was only about 3% of the wireline data, but it was far more valuable. Hence, if users decide they would rather have mobility with lower traffic—as they did with voice—developers and service providers may cater to their desires, and neglect wireline connectivity. With few applications being developed for high-bandwidth fiber optic links, that segment of the industry could stagnate.

On the other hand, it is possible to envisage communication advancing faster than computing and storage. The volumes of broadcast video and of PC hard-disk memories are very high, and new applications, or just new usage patterns, could lead to acceleration in traffic growth.

Whichever way the industry develops, it is important to be able to track it reliably, and improved methods for measuring information production and flows are much to be desired.

References

Anderson, C. (2006). *The long tail: Why the future of business is selling less of more*. New York, NY: Hyperion.

Barten, P. G. J. (1990). Evaluation of subjective image quality with the square-root integral method. *J. Opt. Soc. Am. A*, *7*, 2023–2031.

Briscoe, B., Odlyzko, A., & Tilly, B. (2006, July). Metcalfe's Law is wrong. *IEEE Spectrum*, 26–31. Available at http://www.spectrum.ieee.org/jul06/4109

Cisco. (2011). Cisco Visual Networking Index. Available at http://www.cisco.com/en/US/netsol/ns827/networking_solutions_sub_solution.html

Coffman, K. G., & Odlyzko, A. M. (1998, October). The size and growth rate of the Internet. *First Monday, 3*, no. 10. Available at http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/620/541

Coffman, K. G., & Odlyzko, A. M. (2002a). Internet growth: Is there a "Moore's Law" for data traffic? In J. Abello, P. M. Pardalos, & M. G. C. Resende (Eds.), *Handbook of Massive Data Sets*, 47–93. Dordrecht, The Netherlands: Kluwer. Preprint available at http://ssrn.com/abstract=236108

Coffman, K. G., & Odlyzko, A. M. (2002b). Growth of the Internet. In I. P. Kaminow & T. Li (Eds.), *Optical Fiber Telecommunications IV B: Systems and Impairments*, 17–56. New York, NY: Academic Press. Preprint available at http://www.dtc.umn.edu/~odlyzko/doc/oft.internet.growth.pdf

CTIA. (2011). *Year-end 2010 top-line survey results*. Available at http://files.ctia.org/pdf/CTIA_Survey_Year_End_2010_Graphics.pdf

Hilbert, M., & Lopez, P. (2012). How to measure the world's technological capacity to communicate, store and compute information? (This Special Section.)

Hill, R. (1837). *Post office reform: Its importance and practicability*. London, England: W. Clowes.

IEEE. (2008, June). Special report: The singularity. *IEEE spectrum*. Available at http://www.spectrum.ieee.org/jun08/singularityspecialreport

Ito, Y. (1981). The Johoka Shakai approach to the study of communication in Japan. In C. Wilhoit & H. de Bock (Eds.), *Mass Communication Review Yearbook*, vol. 2, 671–698. Beverly Hills, CA: SAGE Publications.

Jayant, N., Johnston, J., & Safranek, R. (1993). Signal compression based on models of human perception. *Proc. IEEE*, *81*, 1385–1422.

John, R. R. (1995). *Spreading the news: The American postal system from Franklin to Morse*. Cambridge, MA: Harvard University Press.

Kurzweil, R. (2005). *The singularity is near: When humans transcend biology*. New York, NY: Viking.

Labovitz, C., Iekel-Johnson, S., McPherson, D., Oberheide, J., Jahanian, F., & Karir, M. (2009). Atlas Internet Observatory 2009 Annual Report. Available at http://www.nanog.org/meetings/nanog47/presentations/Monday/Labovitz_ObserveReport_N47_Mon.pdf

Licklider, J. C. R. (1965). *Libraries of the future*. Cambridge, MA: MIT Press.

Minnesota Internet Traffic Studies (MINTS) project. (2009). Available at http://www.dtc.umn.edu/mints

Mitchell, B. R. (1992). *International historicals statistics: Europe, 1750–1988*, *3rd ed*., London, England: Macmillan.

NCTA. (2010). National Cable & Telecommunications Association, Cable industry revenue 1996–2009. Available at http://www.ncta.com/Stats/CustomerRevenue.aspx

Odlyzko, A. M. (2000). The history of communications and its implications for the Internet. (Unpublished manuscript.) Available at http://ssrn.com/abstract=235284

Odlyzko, A. M. (2001, February). Content is not king. *First Monday*, http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/833/742

Odlyzko, A. M. (2003a). Internet traffic growth: Sources and implications. In B. B. Dingel, W. Weiershausen, A. K. Dutta, & K.-I. Sato (Eds.), *Optical Transmission Systems and Equipment for WDM Networking II*, *Proc. SPIE*, 5247, 1–15. Preprint available at http://www.dtc.umn.edu/~odlyzko/doc/itcom.internet.growth.pdf

Odlyzko, A. M. (2003b, September). The many paradoxes of broadband. *First Monday*, http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/1072/992

Odlyzko, A. M., & Tilly, B. (2005). A refutation of Metcalfe's Law and a better estimate for the value of networks and network interconnections. (Unpublished manuscript.) Available at http://www.dtc.umn.edu/~odlyzko/doc/metcalfe.pdf