Web History and Economics

Andrew Odlyzko

School of Mathematics
University of Minnesota
Minneapolis, MN 55455, USA
odlyzko@umn.edu
http://www.dtc.umn.edu/∼odlyzko
Revised version, October 9, 2012

Abstract. In retrospect, the Web appears a very natural development, a byproduct of the growth of the information and communication technologies (ICT) sector. Still, the success of the Web was actually attained through a series of fundamental misunderstandings about the direction that various technologies would take, and how society would use them. It was a tool that happened to be "good enough" to meet some urgent needs, but not ideal. A brief survey of the history and pre-history of the Web is presented, with an emphasis on its economics and the many misleading notions that played key roles in its development. This survey leads to some speculative thoughts about the future of the Web.

1 Introduction

The Web is a great success, although there are many views on just what the Web is. There are also many ongoing attempts, such as Web 2.0, Web 3.0, and the Semantic Web, to provide extensions to the Web that would remedy its shortcomings and provide new capabilities.

The only safe prediction about the future of the Web is that most predictions will continue to be wrong, and that success will usually be attained by accident. That has been the historical pattern, as developments of technologies and especially of how people use them have been notoriously hard to foresee.

A nice example of "stumbling to success" is provided by a paper about Web search from an academic conference in 1998. It declared that

> [t]he goals of the advertising business model do not always correspond to providing quality search to users. ... we expect that advertising funded search engines will be inherently biased towards the advertisers and away from the needs of the consumers. ... we believe the issue of advertising causes enough mixed incentives that it is crucial to have a competitive search engine that is transparent and in the academic realm.

Yet Google, the great high-tech success story of the first decade of the 21st century, and central to effective functioning of the Web today, did just the opposite of what this paper recommends. Google is extremely opaque in its operations, it is not in the academic realm, and it is supported almost exclusively by advertising revenues. Yet it has earned wide trust and is by far the most popular search engine on the Web. As a result, it is extraordinarily lucrative. Hence one might think that the authors of the passage cited above shorted Google stock at its IPO, and are now bankrupt. But such a supposition would be wrong, for the authors were Sergey Brin and Larry Page, the founders of Google [1].

This little vignette illustrates the main point of this paper, namely that success often comes by accident. Not only were the Brin and Page predictions about the Web search business wrong, so were those of the overwhelming majority of contemporary observers. At the time of the dot-com bubble crash in 2000–2001, Google was the leading search engine, but there was wide skepticism as to whether search could be a viable business. Even online advertising as a whole was thought to have very limited prospects.

This was not an isolated incident. History is littered with examples of "stumbling to success." One of the greatest technology surprises has been the success of the Internet. Its roots go back all the way to the Arpanet of the 1960s, and it continued to be developed over the decades. However, as late as the early 1990s the general consensus, in the ICT (information and communication technologies) industry as well as in academia, was that it was just an experimental tool, and that other networks would dominate. For example, Andrew Tanenbaum, a distinguished computer networking pioneer, wrote as late as 1989 ([22], p. xiii) that the future belonged to the OSI Reference Model, and that "[i]n the near future, almost all other network architectures will disappear, and computers from one vendor will be able to communicate effortlessly with computers from another vendor, thus stimulating network usage even more."

The range of mistakes that have been made in technology prediction is very wide. Still, there are some persistent patterns one can discern among those mistakes, and in particular among those dealing with the Web and its predecessors. They concern the rate of change, the role of content, the degree of control by service providers, the importance of locality, and the role of voice in communications.

This brief note surveys some of these patterns of mistakes and how they affected the development of the predecessors of the Web and of the Web itself. It concludes with some speculations about the ongoing evolution of the Web.

2 Breakthrough innovation versus historical inevitability

The history of science and technology is full of instances of independent and almost simultaneous inventions. Calculus, the theory of evolution, the telephone, and the lightbulb are just four examples of this phenomena. As a field develops, and theory and available tools improve, new problems become amenable to solution, and if there is enough entrepreneurial drive, somebody will find the solution. Very often the margin of victory has been and likely will continue to be very small in the future.

Not only do we have simultaneous or nearly simultaneous inventions and innovations, we can find examples where being far ahead of the field fails to have an impact. That can be seen in much of Leonardo da Vinci's futuristic visions. In the ICT area, a compelling example is Babbage's pioneering work on digital computers. Babbage was a brilliant inventor. At the same time, he was a poor engineer and manager, so none of his machines were built. On the other hand, two Swedes, Per Georg Scheutz and his teenage son, Edvard, built a calculating machine based on modifications of Babbage's design that worked [7,13]. However, they sold only two copies, and their business folded. Later in the 19th century, Martin Wiberg, another Swede, came up with another, more compact and practical version of this computer, and failed to sell any copies. The problem was that these machines were designed for producing mathematical tables, and the demand for new tables was very limited. Hence Babbage did influence intellectual development of computing. However, the far simpler punch card and the associated tabulating machine dominated in terms of practical computing in the 19th century, and led directly to the computing revolution of the 20th century.

The Web also had early predecessors that fizzled out. A particularly noteworthy example is Vannevar Bush's Memex [15]. It attracted considerable attention when it was published at the end of World War II, and it did influence thinking about information through its proposal for a proto-hypertext system. However, an electromechanical device similar to what Bush had in mind could not be built with the technology of the time.

A more direct predecessor of the Web was Ted Nelson's Xanadu system. It was conceived for electronic digital computers, and the reasons for its failure are still disputed. There does seem to be wide agreement that computers were not sufficiently powerful to allow Nelson's ambitious goals to be realized when he proposed the project and started working on it. In addition, it appears likely that his stringent requirements for permanence, rights management, and other features were simply not what the world needed, whereas the Web was "good enough." Still, Xanadu was a major inspiration for many computing pioneers, and so a forerunner of the Web.

Thus it seems safe to conclude that something like the Web was inevitable. The development of ICT areas, and the growing information volumes of modern society required some convenient tool. So if the Web had not been invented in its original form, another technology would have been found, just as the card catalog (and many other tools, including the paper clip) had been invented to cope with the earlier stages of the information flood. And indeed, in the formative period when the Web was created, there was extensive development and use of tools such as Gopher, WAIS, and others that served to provide access to the growing volume of online information. Further, even after the Web and the Netscape browser had already attained wide usage, there were many observers (including the author of this piece, in [16]) who thought that even better solutions would show up soon. And in a certain way this prediction was fulfilled, except that it was fulfilled by modifications of the basic Web and the basic browser.

So why did the Web as we know it take over? Most likely because it passed a certain magic level of usability that made it attractive to users, and because there was a large enough mass of users on the Internet to adopt it. Even so, it did take several years. Perhaps this stretch would have been shortened had a browser been available for platforms more widespread than the NeXT computer used by Berners-Lee and his colleagues at CERN. But even that is not certain.

Cases where the first-mover advantage is not decisive are common. Google was not the first search engine (Alta Vista was the first one to get wide usage), eBay was not the first online auction site, Facebook was not the first social network, and the iPhone was not the first smart phone. Yet all are still dominant, but in evolved forms. What appears to matter most is getting a technological edge in a market large enough so that a sufficiently large user base can be captured to enable network effects to operate.

Similarly, the Internet was developing slowly in relative obscurity, while the telecom industry was working on ATM. But then the world discovered that this research tool was good enough to handle the growing volumes of data traffic, and adopted it, leaving ATM for niche applications, such as serving as a carrier for Internet traffic.

Hence we can conclude that if the Web had not been invented, we likely would have ended up with something with a different name, but likely similar functionality. Likewise, had the Internet not been around, we might have adopted ATM, and then modified ATM to better serve the real traffic needs, and to get rid of its costly encumbrances.

3 The speed and direction of change

Some types of technological change are nicely predictable. In particular, Moore's Law for semiconductors has fit predictions very well for almost half a century. (It has to be said, though, that the rate of progress has slowed down. Also, in this case Moore's Law is not a natural law, but rather the result of a giant collaborative effort of many companies and research groups in several technologies to achieve the goals they agree on.) Similar "Moore's laws" in some other areas have held with less regularity, but there has been steady progress. On the other hand, in many fields, especially those dealing with society's adoption of new technologies, prediction has been much harder. We do not fly in supersonic planes, nor commute in helicopters, and the "War on Cancer" that it widely regarded as having been declared by U.S. President Nixon in 1971 is still nowhere close to won, even though substantial progress has been made. In general, promoters tend to be overoptimistic about how attractive their technology will be. In some areas, though, there has been a persistent tendency to underestimate the popularity of new devices and technologies, and this applies in particular to computing and communications [21].

J. C. R. Licklider (footnote on p. 17 of [11]) had a very appropriate warning:

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

Generally, promoters of new technologies tend to forget this. Many of the rosy predictions of the Internet bubble years, for example about online shopping, are finally coming true.

There are other pitfalls in prediction. Licklider's book [11] is at one level an unusually successful piece of futurology. He predicted that digital libraries would become feasible around the year 2000, and he was almost exactly right. However, he also expected great progress in AI, resulting in real language understanding by machines. That is still not on the horizon. Progress has come far more from growth in computing and only moderately intelligent pattern matching than from real machine understanding. It appears likely that this trend will continue, and that many predictions of "The Singularity" [9] will not come true. There are arguments that the value of computing, storage, and communication devices and services grows logarithmically in their raw capacities [19]. Thus upgrading from a 1 Mbps link to a 10 Mbps one might correspond to moving from a value of 6 (the decimal logarithm of a million) to 7, and not a 10-fold improvement. That appears to match most people's personal experiences.

In general, the hardest predictions about technologies are about what society will do with them. Sometimes new products and services are embraced with enthusiasm, and sometimes they are ignored. Yet even there one can discern some patterns. Gleick's popular science book *The Information* has the grandiose claim that "history is the story of information becoming aware of itself." Yet what the many engagingly told stories of that book show is that in practice, people use communication technologies primarily for very mundane tasks, most of all for simple connectivity.

4 Content and connectivity

The Web was initially designed at CERN for the distribution of data among physicists. They had special needs, and Berners-Lee developed (after a while with the collaboration of several colleagues) a solution that turned out to be "good enough" for the world at large. But physics is a small area on the global scene. In general, the professionally produced information sources are not a huge industry. As an example, the entire worldwide scholarly publishing market, as measured in terms of publisher revenues, is around $20 billion per year [23]. The global financial data market is around $16 billion [3]. This is large, but small compared to the telecommunications industry, where, in 2011, mobile services alone brought in revenues of over $1 trillion. In the United States alone, wireless industry revenues were about $170 billion in 2011 [4].

Most of public and corporate policy discussions in the telecom area concentrate on "content," material prepared by professionals for wide distribution. Content, especially video content, is perceived as driving the evolution of telecom networks. Yet content is not king [18]. Traditionally content has attracted disproportionate attention from policy makers, but it has never spurred as much willingness to spend as simple connectivity. The $170 billion in 2011 wireless revenues in the U.S. comes overwhelmingly from voice and texting. By contrast, cable industry in the U.S. obtained revenues from residential video services of just $57 billion in 2011 [14].

Table 1. Price per MB.

SMS	$1,000.00
cellular voice	1.00
wireline voice	0.10
residential Internet	0.01
backbone Internet	0.0001

Thus content is large, but not as large as connectivity. The Web has opened up the whole spectrum of communications, and created players such as Google, which is neither content nor connectivity in the traditional sense. "Google envy" drives much of the maneuvering of telecom companies. Yet Google's worldwide revenues in 2011 came to only $38 billion. Admittedly the profit margins and the growth rate of Google are high, but the disparity between what users are willing to pay for simple voice or texting compared to what they pay indirectly through advertisers to use Google is striking.

The volume of information on telecom networks is increasingly video. But volume is not value, and similar disparities have been seen in the past. For example, in the U.S., "[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" ([8], p. 38). Congress, as a matter of public policy, taxed the low-bandwidth but high-value first class letters to subsidize newspaper delivery. A similar policy was pursued by British policy makers [6].

If we take the amount people pay for various communication services as an approximation of how much they value them, Table 1 shows huge disparities. (This table is based on very rough approximations and extrapolations of what people pay in the U.S. and Western Europe, with the figures rounded to powers of 10.)

The underestimation of the value of connectivity is a consistent pattern in history. For example, the popularity of email, acknowledged as the first "killer app" of the Internet, were not expected by the builders of the Internet's predecessor, the Arpanet. (See [17] for details, in particular how email was specifically ruled out from the design requirements of the Arpanet.) What is most surprising about this is that Licklider, who initiated the ARPA research that led to the Arpanet, was the first to point out that the key role of the computer was as a communication device [10,12]. Yet perhaps that should not be too surprising, as it follows the usual historical pattern, see [17]. Typically new technologies are used first for military and government communication. Then, as they become less expensive and become deployed more widely, business connectivity dominates. Further down the road, comes social connectivity. Distribution of content often dominates planning, but (with a few exceptions) is not the key revenue producer.

Table 2. Voice and texting in U.S. wireless industry.

year	voice minutes billions	texts billions
2005	1,495	81
2006	1,798	159
2007	2,119	363
2008	2,203	1,005
2009	2,275	1,563
2010	2,241	2,052
2011	2,296	2,304

While greater bandwidth is desirable, we can even find cases where people willingly move to lower capacities. Table 2, based on data from [4], shows clear signs of texting substituting for voice in the cellular industry in the U.S. Since the number of subscribers has been growing all along during the years represented in the table, the average length of time spent talking on mobile phones has started declining, although slowly, while texting has continued to grow vigorously. This likely represents the perceived advantages of asynchronous communication that is represented by texting (as well as by email).

In the future evolution of the Web, content is almost universally expected to be the crucial component. And indeed it is likely to be very important, since growing network capacity will have to be filled with something, and video is the most likely candidate, and is clearly already beginning to play that role. Further, because high quality video is hard to produce, content (in the sense used here, of material prepared by professionals for wide distribution) is likely to dominate volumes of traffic. But it would go against historical precedents if it dominated in value.

5 Locality, orality, and social communication

The Web is evolving. Further, whether by design or by bitter experience, it is evolving to meet many of the basic communication needs of society. The rise of social networks, for example, is a very natural phenomenon. The flat nature of the Internet, where any node can communicate with any other node, is its main advantage. At the same time it is a key defect, as it facilitates spam and distributed denial of service attacks. While there is value in being potentially able to reach anyone, this value appears to grow not at the n^2 rate that Metcalfe's Law predicts, but at the more sedate pace of $n \log(n)$ [2]. Even with ancient technologies, there has always been a strong locality effect in communication, with higher traffic densities among those located nearby [17].

The Web opened up the whole spectrum of communications, from the extreme of one-to-one basic connectivity of snail mail and voice telephony, to the one-to-many content distribution of newspapers and TV. Social networks and collaboration tools provide a means of more controlled distribution.

Whether Facebook remains the dominant social network is an open question. But even if it does, it will surely have to evolve. How people adopt new communication tools is hard to predict, and just as "The Singularity" [9] is unlikely to arrive, so is a definitive system for social interaction. Today's social networks have their roots in email discussion groups or online bulletin boards (and going even further, in ordinary physically-based social networks, of course). The first widely used commercial collaborative online platform was Lotus Notes, introduced back in 1989. It still exists, now as IBM Lotus Notes, but its long and continuing evolution, and the leap-frogging it has experienced at the hands of competitors, also suggest that much change awaits us.

One area where it appears safe to say that current technologies and ongoing developments are deficient is in the neglect of voice communication. With the ascendancy of the Internet, and now with some instances where wireless customers are replacing voice communications with texting, the industry has been paying little attention to voice. Yet, as can be seen in the works of McLuhan, Ong [20], and others, human culture is at its base an oral one. The huge telecom industry revenues are still coming overwhelmingly from charges for low-bandwidth voice services. The problem this industry is facing is that of moving to a new environment, in which available bandwidth is high, and the cash cow service, namely voice, is just a minor contributor to the traffic. That is simply the inevitable future. However, it appears likely that much more could be done with voice to enhance all those fancy new broadband services, with deployment of high quality voice, voice messaging, and integration of these into other systems.

Another development that a broad historical perspective indicates would be desirable is provision of better tools for users to create videos. Visual communication is extremely powerful, and YouTube shows some of its promise. However, good video is hard to create, so making available to the public better tools for handling video, as well as extensive collections of professionally prepared content they could reuse (as is already done in a modest way in mashups) appear natural developments.

6 Conclusions

The Web is evolving, and is likely to do so for a long time. Human needs and reactions are hard to predict, and so much experimentation will be needed. Historical precedents suggest that this area will often "stumble to success," even in cases where the mistakes are very obvious ones (such as in overemphasis on content and neglect of basic connectivity).

References

1. S. Brin and L. Page, "The anatomy of a large-scale hypertextual Web search engine," presented at the 7th International World-Wide Web Conference, Brisbane, Australia, April 1998. Full version available at http://www7.scu.edu.au/1921/com1921.htm.
2. B. Briscoe, A. Odlyzko, and B. Tilly, "Metcalfe's Law is wrong," *IEEE Spectrum,* July 2006, pp. 26–31. Available at http://www.spectrum.ieee.org/jul06/4109.
3. S. Clifford and J. Creswell, "At Bloomberg, modest strategy to rule the world," *New York Times,*, July 17, 2010.
4. CTIA - The Wireless Association, Year-End 2011 Top-Line Survey Results, http://files.ctia.org/pdf/CTIA_Survey_Year_End_2011_Graphics.pdf.
5. J. Gleick, *The Information: A History, A Theory, A Flood,* Pantheon Books, 2011.
6. R. Hill, *Post Office Reform: Its Importance and Practicability,* 1837.
7. A. Hyman, *Charles Babbage: Pioneer of the Computer,* Princeton Univ. Press, 1982.
8. R. R. John, *Spreading the News: The American Postal System from Franklin to Morse,* Harvard Univ. Press, 1995.
9. R. Kurzweil, *The Singularity is Near: When Humans Transcend Biology,* Viking, 2005.
10. J. C. R. Licklider, Man–computer symbiosis, *IRE Transactions on Human Factors* **1** (March 1960) 4–11. Available at http://memex.org/licklider.pdf.
11. J. C. R. Licklider, *Libraries of the Future,* MIT Press, 1965.
12. J. C. R. Licklider and R. W. Taylor, The computer as a communication device, *Science and Technology* **76** (April 1968) 21–31. Available at http://memex.org/licklider.pdf.
13. M. Lindgren, *Glory and Failure: The Difference Engines of Johann Müller, Charles Babbage and Georg and Edvard Scheutz,* MIT Press, 1990.
14. National Cable & Telecommunications Association, cable industry revenue 1996 - 2011, http://www.ncta.com/Stats/CustomerRevenue.aspx. Downloaded Sept. 28, 2012.

15. J. M. Nyce and P. Kahn, eds., *From Memex to Hypertext: Vannevar Bush and the Mind's Machine,* Academic Press, 1991.
16. A. Odlyzko, "Tragic loss or good riddance? The impending demise of traditional scholarly journals," *Intern. J. Human-Computer Studies,* 42 (1995), pp. 71–122. Available at http://www.dtc.umn.edu/~odlyzko/doc/tragic.loss.long.pdf.
17. A. M. Odlyzko, "The history of communications and its implications for the Internet," 2000 unpublished manuscript, available at http://ssrn.com/abstract=235284.
18. A. M. Odlyzko, "Content is not king," *First Monday,* vol. 6, no. 2, February 2001, http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/833/742.
19. A. M. Odlyzko, "The volume and value of information," *International Journal of Communication*, vol. 6, 2012, http://ijoc.org/ojs/index.php/ijoc/article/view/1570/740.
20. W. J. Ong, *Orality and Literacy: The Technologizing of the Word,* Methuen, 1982.
21. S. P. Schnaars, *Megamistakes,* The Free Press, 1989.
22. A. S. Tanenbaum, *Computer Networks,* 2nd ed., Prentice Hall, 1989.
23. M. Ware and M. Mabe, *The STM Report: An Overview of Scientific and Scholarly Journals Publishing,* International Association of Scientific, Technical and Medical Publishers, 2009. Available at http://www.stm-assoc.org/2009_10_13_MWC_STM_Report.pdf.