



Research Article

The history of the internet: the missing narratives

Martin Campbell-Kelly¹, Daniel D Garcia-Swartz²

¹Department of Computer Science, University of Warwick, Coventry, UK

²Compass Lexecon, 332 S. Michigan Avenue, Chicago, IL 60604, USA

Correspondence:

M Campbell-Kelly, Department of Computer Science, University of Warwick, Room CS326, Coventry CV4 7AL, UK.

Tel: 02476523193

Fax: 02476573024

Abstract

The origins of the Internet are only partially understood. It is often believed that the Internet grew as a tree from a tiny acorn, the ARPANET network set up in 1969. In this study, we argue that this interpretation is incomplete at best and seriously flawed at worst. Our article makes three contributions. First, on the basis of a wide variety of primary and secondary sources we reconstruct the history of computer networks between the late 1950s and the early 1990s. We show that the ARPANET network was one among a myriad of (commercial and non-commercial) networks that developed over that period of time – the integration of these networks into an internet was likely to happen, whether ARPANET existed or not. Second, we make a systematic effort to quantify the significance of these various networks. This allows us to visualize more clearly the extent to which the ARPANET network was one among many, and not a particularly large one at that. Third, we provide a nuanced interpretation of the rise of various technologies, including the Transmission Control Protocol/Internet Protocol and the World Wide Web, as ‘dominant designs.’ Their rise should be interpreted within the economic framework of industries with network effects, in which historical accidents bring about tipping points that lead to universal acceptance. We thus show that history matters for understanding why information systems function in the way they do.

Journal of Information Technology (2013) **28**, 18–33. doi:10.1057/jit.2013.4

Keywords: history of the Internet; information systems; computer networks

Introduction

Most of the current crop of histories of the Internet can be characterized as ‘teleologies’ or ‘Whig history.’ This is a style of writing that takes the present to be the end point toward which history has been inevitably unfolding (Lamoreaux *et al.*, 2004: 377). Teleological histories seek uncomplicated explanations, often based on a single cause for an historical epoch. In the case of Internet history, the epoch-making event is usually said to be the demonstration of the 4-node ARPANET network in 1969. From that single incident the global Internet developed. Examples of (otherwise carefully written) accounts that follow this approach are Salus (1995) and Hafner and Lyon (1996).

In this simplistic explanation, almost all of the networking activities that were simultaneously happening around the globe, commercial and non-commercial, have been written out of the story. The sense conveyed is that the Internet has grown like a tree. From a tiny acorn planted in

1969, we now have the giant oak of the global Internet. But a tree is the wrong metaphor. When the Internet took off in the early 1990s the world was covered by thousands of isolated networks and the integration of these networks into a global entity was likely to happen, whether ARPANET existed or not. A better metaphor is that the networked world was like a super-saturated salt solution. It just needed a single crystal of salt to make the whole change its state. As it happens, that crystal was the ARPANET’s Transmission Control Protocol (TCP)/Internet Protocol (IP). But there were other protocols and technologies that could have established an internet.

This article attempts to fill out some of the missing narratives of Internet history, as far as this is possible in an article of modest length. We begin the story with the rise and fall of ‘computer utilities’ in the 1960s. These developments promised facilities that foreshadowed the Internet

by 30 years, but which lay dormant until technology finally caught up in the 1990s. Although the 1960s dreams of consumer networks did not materialize, the 1970s saw the development of commercial networks on a global scale. We describe this largely hidden infrastructure, which enabled the electronic commerce that became commonplace by the 1980s – from automated teller machines to supermarket checkouts.

By the late 1970s, the idea of an internet (not then graced with a capital I) was in the air. (Throughout this study, we maintain the distinction between an internet, which is any interconnected set of networks, and the Internet, the network of networks operating under the TCP/IP protocol.) We describe the competition between the telecommunications monopolies, computer manufacturers, international standards organizations – and the ARPANET community – to define the Internet's architecture. This contest was still unresolved, but given new force, by the arrival of desktop personal computers in the 1980s. We describe the emergence of numerous forms of network provision for this new community of users – consumer networks, bulletin board systems (BBSs), email services, and videotex. Early adopters of these services were 'online' long before the commercial Internet took off, and the facilities they enjoyed – such as electronic forums and online shopping – later shaped the Internet experience.

The diffusion of desktop PCs finally created the environment in which an internet could flourish. We describe the process by which the Internet – with a capital I – converged around the TCP/IP protocols. In 1990, however, the Internet was a gray and dreary place devoid of content – like a TV station without programs. We conclude with the development of the World Wide Web, which promised a user friendly, visually compelling, content rich experience to which ordinary users could relate. We argue, however, that the World Wide Web was not the unique, killer application claimed by many writers. There were other possibilities, but the accidents of history gave us the Web.

Our study makes three contributions. First, it reconstructs the history of computer networks between the late 1950s and the early 1990s, and thus puts the development of the ARPANET network in perspective. Second, it quantifies, to the extent possible, the significance of these various networks (in terms of investments made to set them up, number of locations reached, number of hosts, and/or number of users). This allows us to visualize more clearly the extent to which the ARPANET network was one among many, and not a particularly large one at that. Third, it interprets the rise of various technologies, including the TCP/IP protocol and the World Wide Web, as 'dominant designs' in the framework of industries with network effects, in which historical accidents may generate tipping points that lead to universal acceptance (Farrell and Klemperer, 2007).

The economic literature on industries with network effects (or network industries) points out that a network effect is present if the value that an individual or company derives from consuming a good or service depends, to some extent at least, on how many other individuals or companies consume it. The higher the number of computers that are connected to a computer network, for example, the more valuable the network is to each participant. Network industries have a number of peculiar

characteristics. First, they tend to be unstable and to tip to one supplier only. Second, for that very same reason, companies or organizations in network industries often compete 'for the market' rather than 'in the market.' Third, history matters in several ways. Initial conditions and small historical accidents, for example, may have a decisive influence in shaping the evolution of a network industry. Fourth, history matters also in the sense that current adopters of a network good/service will likely care about past adoption (that is, about the installed base). Fifth, expectations matter as well, as current adopters may care not only about past adoption but also about the likelihood of future adoption. Finally, it is difficult for industry players to get traction at the starting point – the 'chicken and egg' problem – and it is even more challenging, although not impossible, for an entrant to displace a dominant incumbent. Our analytical narrative of the evolution of the Internet – the archetypical network of networks – provides abundant examples of how these features play a role in shaping historical outcomes.

Our study shows that history matters for information systems because history allows us to develop a richer understanding of how these systems come to be. They are rarely the outcome of a single vision of a group of actors who 'know best' – they are rather the result of competing visions that are selected out by market forces, social interactions, and happenstance.

In order to make it easier for the reader to follow the developments that we analyze in this study, Table 1 presents a timeline. Precisely to make the point that the ARPA Internet was just one of the paths that could have led to the Internet, we track both ARPA-related and more general network developments over time.

The rise and fall of the computer utility

There were two historic developments that fed into the World Wide Web, one conceptual and another technological. The conceptual development was the idea of a global store of information, accessible to any person at his or her desktop. This vision was articulated as the 'World Brain' by HG Wells in the 1930s, as the 'memex' by Vannevar Bush in the mid-1940s, and as a 'thinking center' by JCR Licklider in 1960. The technical development was computer networking, which eventually made the vision possible.

World brain and memex

In the 1930s, the writer and futurist HG Wells (1866–1946) conceived the idea of a World Brain, 'a complete planetary memory for all mankind' (Wells, 1938: 60). Wells (1938: 54) wrote, in a prescient vision of things to come: 'The time is close at hand when any student, in any part of the world, will be able to sit with his projector in his own study at his or her own convenience to examine *any* book, *any* document, in an exact replica.' Wells' vision faded from sight during WWII, but Vannevar Bush (1890–1974) articulated a similar concept in the United States after the war. Bush was probably the most important scientific administrator of the twentieth century (Zachary, 1997). An MIT graduate and professor of electrical engineering, Bush rose to become head of the wartime Office of Scientific Research and

Table 1 An internet timeline

Year	General network developments	ARPA-related network developments
<i>(A) 1938–1969</i>		
1938	Wells, <i>World Brain</i>	
1945	Bush, ‘As We May Think’	
1958		ARPA established
1960		Licklider, <i>Man Computer Symbiosis</i>
1961	MIT Compatible Time-Sharing	
1962	SAGE system ERMA banking network	Licklider Program Director at ARPA IPTO formed at ARPA Baran’s packet-switching paper
1965	SABRE reservations system MIT Project MAC Commercial time-sharing Remote processing systems Online information systems	ARPA network study Davies’ packet-switched network
1967		ARPA adopts packet-switching
1969		First ARPANET node at UCLA
<i>(B) 1971–1995</i>		
1971	Tymnet packet-switched network	
1972		
1973	First EDI network (railroads)	
1974	IBM announces SNA	TCP specified
1975	Telenet packet-switched network X.25 defined by the CCITT	
1976	UPC in supermarkets	
1977	OSI project	
1978	First Bulletin Board System MicroNET (CompuServe)	TCP/IP specified
1979	First videotex systems	
1983	AOL online service MCI email service	Milnet separated from ARPANET
1984	Prodigy online service	
1986		NSFNET established by the NSF X.400 protocol specified
1990	WWW functional at CERN	ARPA internet subsumed into NSFNET
1991	Gopher live WWW released by CERN WAIS invented	
1993	MOSAIC browser	
1994	First Internet service provider	
1995	The Internet a private entity	The Internet a private entity

Source: See text.

Development, the organization responsible for coordinating WWII scientific research.

As the war was drawing to its close, Bush had begun to reflect on what he saw as the most important post-war scientific problem – the information explosion. He published an article ‘As We May Think’ in *Atlantic Monthly* in July 1945 and in *Life* magazine a few months later, in which he described an information storage and retrieval machine called the ‘memex’ (Bush, 1945a, b). The article caused a tremendous amount of interest, and although Bush never credited Wells or his contemporaries in Europe, ‘the rhetoric seemed straight from Wells himself’ (Zachary, 1997: 265).

Government and private networks

In August 1949, Russia exploded its first atomic bomb. This event triggered the Cold War, and massive defense spending on computer networks. The United States urgently needed a new air defense system. A swift review by the Air Force Scientific Advisory Board at the end of 1949 set in train a process that would – with an investment of US\$8 billion – lead to the SAGE air defense system, first deployed in 1962 (Edwards, 1996: 74–111).

SAGE was utterly groundbreaking. At a time when just a handful of prototype computers had been built for mathematical calculations, SAGE would use computers for the entirely new purpose of real-time control and the



integration of multiple data sources. The technologies pulled through by SAGE would give the United States a dominance in computer science, technology, and industry that has persisted to the present day (Flamm, 1987: 176–180, 1988: 89–90). SAGE was just the beginning of a huge investment in similar computerized command-and-control systems, encompassing at least 25 other military systems built in the late 1950s and early 1960s (Edwards, 1996: 107).

There was extensive spin-off from SAGE in the civilian and business sectors. However, even before SAGE there had been forays into computerized real-time operations, such as the teashop distribution system devised for the Lyons Company in the United Kingdom in 1954 (Caminer *et al.*, 1998: 45–47, 375–383). Such applications represented a latent demand waiting for technology to catch up. The SAGE project was this catching up. The IBM–American Airlines SABRE airlines reservations system was the first project to capitalize on the SAGE technology (McKenney *et al.*, 1995: 97–140). When fully deployed in 1965, the system enabled 1100 travel agents, scattered throughout the United States, to access a mainframe computer at the American Airlines reservations center in Briarcliff, New York. American Airlines' competitors were forced to follow suit – by 1962 it was estimated that there were a dozen airline reservation projects completed or underway (Datamation, 1962: 53–55). Where the airlines led, American enterprise followed. The banks, for example, quickly saw the significance of computer networks: General Electric and Bank of America pioneered with the ERMA system, a network of 32 computers that was operational in 1962 (McKenney *et al.*, 1995: 64). During the second half of the 1960s, businesses, utilities, and governments in many parts of the industrialized world established dozens of computer networks.

Computer time-sharing

In the late 1950s, largely independently of government and corporate real-time networks, universities and research organizations with much smaller budgets had begun to experiment with 'time-sharing' computers. In a time-sharing system, a large mainframe computer was shared among many users equipped with typewriter-like terminals. This allowed a user to type a program directly into the computer and get results back in a few seconds. MIT was the first university to develop a rudimentary system, the Compatible Time-Sharing System, in November 1961 (Fano and Corbato, 1966). Organizations pursuing the same goal included Carnegie Mellon University, Dartmouth College, RAND, and several others.

MIT, however, was to become the primary locus of time-sharing development, for reasons that reflected its connections with the SAGE project as much as its technological capability. A leading investigator in the SAGE project was JCR Licklider (1915–1990), an MIT alumnus now often dubbed the father of the Internet (Waldrop, 2001). A psychologist and an engineer, Licklider established the psychology program for engineers at MIT in 1950. He subsequently directed man–machine communication studies (as people called them at the time) for the SAGE project. After that, he became a vice-president at Bolt, Beranek and

Newman (BBN), then a small consulting firm that was on the cusp of broadening its mission to include computers. In 1960, Licklider (1960) wrote one of the classics of computer science, *Man Computer Symbiosis*, in which he advocated the use of the computer as an everyday tool for knowledge work – an early vision of what we now call personal computing. In 1962, he took leave from BBN to serve as Program Director of the Advanced Research Projects Agency (ARPA), where he would exercise considerable influence on the future of computing (Norberg and O'Neill, 2000; Abbate, 1999).

ARPA had been established in 1958 as a response to Sputnik, with the 'stated mission of keeping the United States ahead of its military rivals by pursuing research projects that promised significant advances in defense-related fields' (Abbate, 1999: 36; Lukasik, 2011). The Information Processing Techniques Office (IPTO) was formed in 1962, with an initial annual budget of \$7 million. Licklider was granted a great deal of autonomy to set the research direction, and he chose to work on the advancement of human–computer interaction technologies. Most contracts were given to a charmed circle of elite institutions – MIT, and the universities of Stanford, Carnegie Mellon, and Utah.

Licklider awarded \$3 million to MIT to develop a state-of-the-art time-sharing system, Project MAC (Fano and Corbato, 1966: 77). The system became operational in 1965, with 160 terminals located around the MIT campus. When Licklider's term of office was over, his replacement Ivan Sutherland – another MIT alumnus – provided \$7 million for MULTICS, a giant time-sharing system intended to support a thousand or more simultaneous users.

The computer utility

By the mid-1960s, the concept of a 'computer utility' was rapidly gathering momentum. The idea of a computer utility was to provide computing from a giant centralized mainframe, rather like an electric power plant supplied electricity. Probably Martin Greenberger (1964), an Associate Professor at MIT's School of Industrial Management, set the ball rolling with an article in *Atlantic Monthly* in May 1964 'The Computers of Tomorrow.' His somewhat temperate vision predicted applications such as information retrieval, bill payments, and stock trading. After that, all-comers jumped on the computer-utility bandwagon, and there was a flurry of conferences and books about computers in the home (Parkhill, 1966; Barnett *et al.*, 1967). However, market reality soon set in. The cost of domestic computing – \$10 an hour and upward – was completely prohibitive, and so the dream of domestic online computing would lie dormant for approximately 20 years.

Meanwhile, there was a real market for engineering calculations and financial analysis using time-shared computers. In 1965, General Electric launched its GE-265 time-sharing service, which was soon offering online computing in 20 US cities (GE Information Services, 1985). Other new entrants included Tymshare, University Computing, Compu-Serv, and a dozen others. By 1970, the market for time sharing was estimated at \$240 million, about 4% of total expenditures on computers and computer services (Campbell-Kelly and Garcia-Swartz, 2008).

Building the electronic infrastructure in the 1970s

During the 1970s, the world became overlaid by hundreds of large-scale networks. Because so many industries, organizations, and computer manufacturers were simultaneously developing the new technology of networks, many competing solutions and standards inevitably emerged. Organizations anxious not to be left behind simply pressed ahead with *ad hoc* solutions based on the emerging standards that most nearly fit their requirements.

Private and public information services

Private and public information services led to two different models of network development. Private networks typically established their own permanent data communications infrastructure, and used it exclusively. Public information services, which were available to any organization or member of the public, could not physically establish a link to each potential customer and therefore had to rely on data communications supplied by the telephone and telegraph monopolies.

The IBM–American Airlines SABRE system is a well-documented example of a private network: it supplied a reservation service to its own sales offices and to selected travel agents, which were permanently connected by lines leased from AT&T (Plugge and Perry, 1961). Public online services presented a very different problem, in that they potentially served thousands of users on an intermittent basis with no permanent connection. Long distance and international calls were prohibitively expensive, and would have deterred most users when added to the already high online costs.

The time-sharing vendors were the most aggressive in building private infrastructures that interfaced with the ordinary telephone networks. The history of General Electric Time Sharing, one of the market leaders, is typical (GE Information Services, 1985). It started in 1965 with

a strategy of deploying computers in major American cities (and Europe) – in effect, the computer was taken to the customer. In 1969, GE consolidated its operations into a single ‘supercenter’ in Cleveland, Ohio, with one very powerful centralized computer that would serve the entire network. Local-call access was provided in major US cities, and a satellite link was leased to London for European customers. Two further supercenters were subsequently added in Washington DC and Amsterdam. Renamed General Electric Information Systems, by the late 1970s the service was available with local call access from over a hundred locations, from Puerto Rico to Helsinki. It had 5000 employees, sales and service offices in every major location, 6000 corporate customers, and claimed to be the world’s largest supplier of computing power. General Electric had several major competitors that also grew prodigiously during the 1970s – Tymshare, Comshare, and University Computing among them. These ‘interactive services’ had aggregate sales of over \$1 billion (in nominal terms) by the mid-1970s (Campbell-Kelly and Garcia-Swartz, 2008). Table 2 tracks the evolution of the time-sharing industry between the mid-1960s and the late 1970s (in the context of the evolution of the data processing services industry more generally).

Another set of suppliers that built private networks comprised those providing ‘remote processing services,’ where firms off-loaded payroll and other routine processing to a remote service provider. Market leaders included ADP, CSC, EDS, Keydata, and University Computing, with aggregate revenues of more than \$300 million in the mid-1970s. By the late 1970s, online services – interactive plus remote processing – generated sales approaching \$3 billion (again, in nominal terms), which represented about 16% of the total expenditure on computers and computer services (Campbell-Kelly and Garcia-Swartz, 2008).

Another emerging market in the 1970s was online information services. Industry leaders included Dialog,

Table 2 Data processing services industry, 1963–1978 (\$ billion at 1963 prices)

Year	Remote processing	Interactive	Total online services	Total data-processing services	Total computer shipments	Total online percentage
1963		0.005	0.005	0.270	1.300	0.32
1964		0.010	0.010	0.312	1.644	0.50
1965		0.014	0.014	0.393	1.976	0.61
1966		0.019	0.019	0.500	3.085	0.52
1967		0.045	0.045	0.663	3.636	1.05
1968	0.009	0.095	0.103	0.893	4.154	2.04
1969	0.041	0.130	0.171	1.189	4.007	3.29
1970	0.069	0.185	0.255	1.466	3.360	5.28
1971	0.084	0.244	0.328	1.521	3.120	7.08
1972	0.101	0.307	0.408	1.730	3.919	7.22
1973	0.134	0.412	0.546	2.027	3.888	9.23
1974	0.168	0.495	0.663	2.310	4.218	10.16
1975	0.191	0.560	0.751	2.524	3.640	12.19
1976	0.290	0.596	0.886	2.802	3.636	13.76
1977	0.404	0.679	1.083	3.125	4.292	14.61
1978	0.483	0.779	1.262	3.446	4.624	15.64

Source: Phister (1979: 277, 610); total data-processing services include traditional mail-batch services, software, and facilities management; total shipments include shipments of mainframes and minicomputers.



Lexis, SDC Orbit, and there were a score of lesser entities (Bourne and Hahn, 2003). Most of these services had developed experimental systems in the 1960s, and during 1971–1973 they attempted to commercialize them. For example, Dialog, a subsidiary of Lockheed Corp., supplied online bibliographic information. Starting with two databases in 1972, by 1976 it was offering 50 databases containing a total of 12 million records, and had thousands of customers in North America and Europe (Bourne and Hahn, 2003: 141–184). Lexis was launched in 1973 and initially offered a service to lawyers in New York, Washington DC, and Ohio (Bourne and Hahn, 2003: 229–258). Lexis was unique in that it offered full text case law and employed offshore keyboard operators to enter the texts. The system was expensive, but popular with lawyers who could afford it, and 3000 customers were signed up in the first year. The service soon went nationwide.

The electronic data interchange (EDI) movement

The first commercial networks, such as those operated by airlines and banks, were relatively simple in term of business process. These industries were relatively homogeneous, and most transactions were, in any case, conducted within the enterprise.

The need for dissimilar organizations to exchange messages electronically led to the EDI movement in the 1970s. EDI was simultaneously a concept, a set of standards, and a set of networks. The EDI concept was to enable electronic transactions between businesses within a single industrial sector. Once an EDI standard had been negotiated within an industry, a network was usually built and operated by a third party, paid for by the trade associations and industry members (Kimberley, 1990).

The first EDI standard was established for the US railroads in 1973 and served as the prototype for many industrial sectors embarking on EDI (Cortada, 2004: 237). The second half of the 1970s saw EDI networks established for many industry sectors: IVANS for the insurance industry, ORDERNET for the pharmaceutical industry, networks for the automotive industry, the chemical industry, and so on.

The prospects for EDI were transformed by the introduction of the Universal Product Code (UPC; the barcode) in US supermarkets in 1976 (Cortada, 2004: 296–302; Brown, 1997). The histories of the UPC, EDI, and the grocery industry are intimately intertwined. The UPC enabled data on goods sold to be captured, allowing for automatic stock control within a retail operation. However, EDI allowed the process to go a stage further: sales data could be sent back to suppliers so that they could replenish stocks ‘just in time.’ In the early 1980s, EDI and barcodes spread throughout the world.

The new paradigm: packet-switched networks

In the 1970s, the telecommunications monopolies saw no opportunity for dedicated data communication networks, and blocked the entry of potential providers that did. After much lobbying, in 1974 the FCC enabled computer firms to offer data communications services. The most successful of these new entrants were Telenet and Tymnet, both of which used the new technology of packet switching.

Donald Davies of the National Physical Laboratory (the British equivalent of the National Bureau of Standards) introduced the packet-switching concept (and the term ‘packet’) for data communications in the United Kingdom (Yates, 1997). (Paul Baran developed a similar approach in the United States; see, for example, Lukasik, 2011: 9–10.) Packet switching reduced data transmission costs by an order of magnitude. In the existing ‘circuit switched’ networks, a direct connection was maintained between a user and a remote computer for the duration of the interaction. This was very expensive when making long distance connections, because most of a user’s time was actually spent thinking, not typing. Packet switching replaced direct person-to-computer connections with a diffuse packet-switched network. Messages that went into the system would be broken up into packets of about 100 characters. The packets would be injected into the network, like a batch of tiny telegrams, each bearing the address of the destination and the sender. High-speed computers would switch the packets toward their destination, taking just a few milliseconds at each node in the network. When packets arrived at their destination, the original message would be reconstituted from the packets. Davies’ long-term aim was to have the UK telecommunications monopoly establish a national network. Although this did not happen for more than a decade, he did manage to convince others.

Most importantly, the ARPA networking project decided to adopt packet switching in 1967 and subcontracted the development to BBN. The Tymshare time-sharing system also developed a network, Tymnet, that was initially only for Tymshare’s subscribers but was soon opened to non-Tymshare computers. By 1979, Tymnet offered local-call access in 180 American cities, had 250 host computers connected to it, and offered 23 foreign access points (Tymshare, 1979: 15). Telenet, a subsidiary of BBN, was perhaps the most important US public packet-switched network. Larry Roberts, the project leader of ARPANET, became Chairman of Telenet, and the service was launched in mid-1975. By 1978, the system was available in 156 US cities and 14 other countries (Roberts, 1978).

The Telenet and Tymnet services were very important in liberating information services such as Dialog and Nexis from reliance on long distance calls. They reduced the data communications costs from \$30 an hour to \$8–10, and even Europe could be connected at \$22 an hour, a massive saving on the previous \$160-plus charges (Bourne and Hahn, 2003: 361).

Beyond the United States, national telephone providers had begun to offer packet-switched services by the end of the 1970s – in the United Kingdom and Canada (1977), France (1978), and Japan (1979) (Abbate, 1999: 154; Quarterman, 1990). Even computer manufacturers developed proprietary protocols for packet switching. IBM’s Systems Network Architecture (SNA) and Digital Equipment’s DECNET were the best known (Jones, 1984; Martin, 1987).

By the end of the 1970s, packet switching was the dominant design in data communications. The next issue to confront the computer community was the need for *inter-networking*, so that it would be possible for networks to communicate with each other. There was intense competition to define internet standards, and thereby liberate data

communications or gain a commercial advantage in the vast computer and communications market.

The social construction of the Internet Protocols

There were many potential models for an internet protocol, although there were only four groups that had the resources and international plausibility to set a global standard. These four were the International Telecommunications Union and the national post, telegraph, and telecommunications authorities (the PTTs), the computer manufacturers (principally IBM), the International Standards Organization, and the Department of Defense acting through the ARPANET community. All of these participants had sensible and viable visions for an internet, and several outcomes were possible.

The CCITT and X.25

During the early 1970s, the PTTs at last began to see data communications as a significant market and set about establishing international standards. The International Telecommunications Union operated through the CCITT standards-setting authority. The CCITT's vision was inevitably constrained by its history of providing old-fashioned telephone services. This legacy, CCITT's ponderous and labyrinthine operating procedures, and its mountains of documentation made it a risible target for the slightly anarchic ARPANET community (Malamud, 1993: 124–133). However, looked at objectively through the long-lens of history, the international telephone network was a staggering technical success. It was possible with the simplest of interfaces – a telephone keypad or dial – to make a direct connection to almost anywhere on the planet. This was largely possible because all the 'intelligence' was inside the network. In reality, the phone network was an ugly mishmash of electro-mechanical and solid-state switches, repeaters and multiplexers, copper wire and optical pipes, and submarine and satellite links. Yet all of this was completely invisible to users: they just dialed a number, period.

The CCITT wanted something analogous for national data communications networks – a 'virtual circuit' that connected a user with an information service, or connected two or more networks to each other. Inside the network there would be an ever-evolving mix of technologies, just as in the telephone networks, but the *interface* for users of the network would be simple and extremely stable. Although packet switching would be used inside the network, this would be a matter of plumbing that was invisible to the user. A message sent into the network would emerge exactly as it went in. The service would be responsible for converting an incoming message into packets, and if any of the packets got lost or damaged in transit, this was the responsibility of the network, and part of the consumer's guarantee of service.

In 1975, a coalition within CCITT consisting of the PTTs of Canada, France, and Britain, plus the US-based Telenet, defined an end-to-end packet-switching protocol known as X.25 (Deasington, 1985). During 1977–1979, X.25 was incorporated in the national packet-switched services of Canada, Britain, France, and Japan, as well as Telenet in the United States. Furthermore, computer equipment vendors incorporated the protocol as a common interconnection

standard. Somewhat later, in 1984, the X.400 email protocol was introduced and it too achieved significant take-up but was later largely replaced by the SMTP email protocols of the Internet (Jakobs, 2013). The fact was that the computer world would never cede its entire network architecture to the PPTs.

IBM and SNA

By the mid-1970s, most computer manufacturers had embarked on developing proprietary architectures – IBM with SNA, the Digital Equipment Corporation with Digital Network Architecture, Xerox Data Systems with XNS, Burroughs with BNA, and so on.

For all the manufacturers, proprietary protocols were partly a matter of expedience (because there were no published standards to work to), but they were also a way to differentiate their products and lock in users. The smaller players were more willing to consider public standards because they had the most to gain from an open market of interconnecting data communications products. For this reason they quite liked X.25 (Deasington, 1985: 15).

In 1974, IBM announced its SNA. Over the next 15 years, SNA evolved from a set of protocols for centralized, mainframe-based networks (like SABRE), to an internet-like network of networks – with the important qualification that only SNA-compliant networks could join. IBM supplied several hundred communications products, and did not readily allow third-party competitors into that market; when obliged to do so it charged hefty licensing fees for use of its protocols.

IBM had an unparalleled amount of networking experience, and knew what industrial users wanted in terms of performance and security. For example, IBM's most important (and profitable) networking software product was CICS. In 1977, it had 6100 licensed users and by 1981 it had 16,550. In the same years, ARPANET had just 111 and 213 hosts, or approximately 389 and 746 users, respectively (see Campbell-Kelly, 2003: 150 for the number of CICS users, and Zakon, 1997: 16 for the number of ARPANET hosts; Salus, 1995: 218–221 suggests that, on average, one IP address per host and 3.5 users per IP address are reasonable assumptions). By 1989, IBM had installed over 25,000 SNA networks (Computer Economics, 1989), at a time when less than 1000 networks were connected to the Internet (Zakon, 1997: 16). With SNA, IBM had everything it needed to establish an internet. But IBM's monopolistic reputation, its lack of diplomacy, and its unwillingness to negotiate in technical committees, doomed its chances of defining the Internet.

On the whole, IBM has been written out of popular Internet history. But if it had played its cards differently, it is entirely possible that IBM and not the ARPANET community would have defined the architecture of the Internet.

The open systems interconnection (OSI) standards

Computer users and the smaller computer manufacturers were caught between the devil of the PTTs and the deep-blue sea of the dominant mainframe makers, IBM and Digital Equipment. The CCITT X.25 standard effectively vested ownership of the network infrastructure in the PTTs, which would leave users with little control over tariffs



and computer researchers with little freedom to evolve distinctive network architectures. But control vested in IBM was more worrisome still. IBM's use of standards to sustain its market dominance was a universal, if sometimes anecdotal, belief (Flamm, 1988: 214; DeLamar, 1986: 283). There was also technical opposition to the idea of a homogenous network, whether it was run by the PTTs or IBM. The research communities, especially the ARPANET community, were opposed to this homogenous view – they saw the future as a world where many different network architectures should be able to flourish and communicate with one another – a pluralistic view of which much has been made in the popular histories.

In 1977, a group of computer manufacturers, major users, and academics organized the OSI project under the auspices of the International Standards Organization. European computer makers were particularly well represented, because the OSI project held the prospect of reducing the dominance of IBM. Eventually, 12 European computer manufacturers, and 5 American, enlisted (Gannon, 1997: 236–238).

The OSI view was that it was premature to define a set of binding standards, like SNA or TCP/IP, while the technology was in a state of flux. Instead, OSI proposed an architectural framework – the 7-layer model – that could accommodate evolving and existing standards. The OSI model was widely accepted, both in Europe and the United States, and redefined the way people thought about network architecture. It also shifted power away from the PTTs and IBM and toward users and smaller manufacturers.

Throughout the 1980s, there was every reason to suppose that networking would eventually converge to OSI standards. However, with so many conflicting commercial and national interests, the OSI project became mired in bureaucratic and technological wrangling. The result was that the ARPANET's TCP/IP protocols ran away with the prize just as it seemed to be within the grasp of the OSI movement. It was largely a matter of timing; networks were developing at a speed faster than the OSI's negotiating processes. This slow progress was perhaps an inevitable consequence of the multinational, democratic participation in OSI. For speed, the OSI could not compete with the small, tight-knit ARPANET community.

ARPA and TCP/IP

In 1977 when the OSI project started, the TCP/IP protocol did not formally exist – it would not emerge as a robust standard until about 1983. In the mid-1980s, and even as late as 1990, few networking experts would have expected TCP/IP to emerge as the global standard.

ARPA's IPTO had taken an early interest in computer networks by funding time-sharing systems, such as MIT's Project MAC and the SDC time-sharing system used at UCLA, Berkeley, and Stanford (Norberg and O'Neil, 2000: 94). In 1964, the ARPANET project was created to link together the computers in the several laboratories funded by ARPA. ARPANET would also act as a laboratory for the development of computer networks.

In 1972, Robert Kahn, another MIT and BBN alumnus, became Program Manager for ARPA's IPTO. In 1976, he recruited Vint Cerf, a Stanford University computer

scientist, to lead the networking activity. Kahn and Cerf have often been described, along with Licklider, as the 'founding fathers of the Internet.' By the time of Cerf's arrival, ARPA had created two additional experimental networks, PRNET and SATNET, which used incompatible protocols designed for radio- and satellite-based packet switching, respectively. It was a logical next step to merge these networks as a 'seamless whole,' for which a protocol was specified by Cerf (Russell, 2006). The result was a tiny 'internet' of three networks. In 1978, the protocol was divided into two parts – the TCP, which operated inside the network, and the IP, which operated between networks.

At this time, the end of the 1970s, ARPANET was not notably different to several other embryonic inter-networking research projects – such as the EPSS network in the United Kingdom, the Cyclades network in France, and the European Informatics Network (Quarterman and Hoskins, 1986). The European academic-research networks were actively participating in, as well as being paced by, the OSI standardization effort.

The ARPANET community operated very differently, and to a different beat of time. The most important innovation was organizational, not technical. TCP/IP was an *ad hoc* solution to inter-networking, not unlike those being devised elsewhere; the fact that it would eventually emerge as a global standard was largely because of the way in which it was designed and accepted by organizational consensus.

In 1983, the military nodes of the ARPANET were hived off into a separate network known as Milnet. Freed from security and defense concerns, this allowed what was now called the 'ARPA Internet,' or simply the Internet, to become a network experiment on a global scale. By 1985, it was estimated that there were nearly 2000 hosts on the ARPA Internet (Lottor, 1992; Zakon, 1997). However, this amounted to only a small fraction of the world's networked computers. In fact, it has been pointed out that 'throughout the 1970s and 1980s, many observers thought of the Internet as a transitional network, an interesting experiment that would fade away once OSI protocols were standardized and implemented by users and manufacturers' (Russell, 2006).

The impact of the PC in the 1980s

Until the advent of the personal computer, network access was confined to the relatively few users inside organizations who had the use of a computer or time-sharing terminal. The diffusion of millions of consumer PCs in the 1980s created a mass-market for online services. This market was largely ignored by the computer manufacturers and academic networks, but was fostered by computer hobbyists, commercial online services, and the PTTs.

BBSs and consumer networks

For PC users who had no institutional network connection, there were two kinds of online access: small, free, informal, and usually local BBSs or national commercial networks such as CompuServe. Both took off in the early 1980s.

BBSs were accessible to technically capable and dedicated computer users, mainly hobbyists. The first system, which went online in February 1978, is attributed to two

Chicago-based computer hobbyists (Aboba, 1993: 59). Because only one user at a time could access the system, the operating software emulated a cork-and-thumbtack bulletin board so that users could leave messages for one another. The BBS idea spread like wildfire. The more capable systems allowed several simultaneous users and provided online conferences for them to interact, thereby extending the bulletin board metaphor. From 1984 on, it was possible for bulletin boards to connect to one another by low-speed telephone networks (Aboba, 1993: 535–538, 545–548). There were also commercial BBSs. Vendors of hardware and software operated some of the BBSs to distribute product information and software upgrades; others operated as online stores; others served up multi-user online game playing. The best documented, and probably the most professionalized, BBS was *The Well*, which was run by the publishers of the *Whole Earth Review* and developed a cult reputation (Hafner, 2001; Reingold, 1993).

The universe of BBSs was extremely fragmented and specialized, and no one system provided a complete user experience; users would typically sign up with several systems. In 1993, *Boardwatch* (a magazine for the BBS community) estimated that there were 57,000 BBSs in North America, and perhaps 100,000 worldwide with 10 million users (Aboba, 1993: 59). In effect, each BBS operated as would a website today, providing a specialized service or a sliver of knowledge.

The most successful commercially run consumer network of the 1980s was CompuServe. Originally established in 1969 as Compu-Serv in Columbus, Ohio, in the 1970s it operated in a niche market as a time-sharing service for insurance companies. The motivation for creating a consumer network was to generate additional revenues by more fully utilizing its computer plant, which was little used outside office hours. In 1978, CompuServe collaborated with the Midwest Association of Computer Clubs to establish a bulletin-board-style service called MicroNET. The service was popular and went nationwide in 1980. By summer 1984, CompuServe claimed to have 130,000 subscribers, 600 employees, and 26 mainframe computers in its Columbus headquarters (Levering *et al.*, 1984: 414–420). By September 1985, the number of subscribers had grown to 225,000.

Whereas a BBS provided a narrow specialized service, CompuServe aimed to provide all of a user's online needs. As well as developing all the services of a BBS, CompuServe integrated and resold the services of other information providers. Its basic services included: email to other users, conferences and forums, online chat rooms, a National Bulletin Board for posting classified ads and notices, and computer games (Bowen and Peyton, 1984). Other services supplied by third parties and repackaged by CompuServe included weather forecasts, AP wire services, newspapers and magazines, stock quotes, banking, and online shopping (the 'electronic mall').

There were competitors of CompuServe, though they were few in number because the entry costs for building a national network were very high. They were principally *The Source* (a joint venture between *The Readers Digest Association* and *Control Data Corp.*), *Genie* (a spin-off from *GE Information Services*), *Delphi* (a subsidiary of

Table 3 Consumer online networks, 1979–1989, number of subscribers (households)

Year	CompuServe	The Source	Delphi	Genie	Total
1979	N/A	1500			1500
1980	1200	7000			8200
1981	17,467	8000			25,467
1982	33,733	20,000			53,733
1983	50,000	35,000			85,000
1984	120,000	49,000			169,000
1985	200,000	63,000	4500		267,500
1986	275,000	65,000	6333	3000	349,333
1987	350,000	60,000	8167	56,000	474,167
1988	430,000	75,000	10,000	115,000	630,000
1989	500,000	53,000	120,000	150,000	823,000

Sources: Glossbrenner (1983, 1985, 1990); H&R Block (1983–1992); and a wide variety of newspaper and trade-press articles available from the authors upon request.

Table 4 Key consumer online networks, 1990–1995, number of subscribers (households)

Year	CompuServe	Prodigy	AOL	Total
1990	555,000	267,000	30,000	852,000
1991	800,000	588,000	75,000	1,463,000
1992	985,000	882,352	181,000	2,048,352
1993	924,000	1,230,000	302,000	2,456,000
1994	1,378,000	1,230,000	900,000	3,508,000
1995	2,221,000	1,230,000	3,500,000	6,951,000

Sources: H&R Block (1983–1992); CompuServe (1997); Glossbrenner (1995); AOL (1990–1996); and a wide variety of newspaper and trade-press articles available from the authors upon request.

General Videotex Corp), and *Prodigy* (a joint venture between IBM and Sears, Roebuck). None of these services was remotely as successful as CompuServe, which benefited partly from its superior content, but also from positive feedback effects. More users funded better services, and better services attracted more users in a virtuous cycle (Shapiro and Varian, 1998: 173–225; Campbell-Kelly *et al.*, 2008). Table 3 tracks the evolution of the consumer online networks in the 1980s.

America Online (AOL), a relative latecomer to consumer networks, provided much more effective competition for CompuServe (Swisher, 1998; Glossbrenner, 1995: 105–110). AOL developed an exceptionally user-friendly and accessible interface. It was not unlike the accessibility leap that the web browser was later to make for the Internet. AOL's network and content may not have been much differentiated from its competitors, but the service excelled in getting consumers online with the minimum of effort and disappointment. Table 4 tracks the rise of AOL *vis-à-vis* CompuServe and Prodigy.

Demonstrating that ease-of-use and accessibility could overcome network effects, AOL grew extremely rapidly, overtaking CompuServe's 2 million plus subscribers in 1995. In the second half of the 1990s, AOL would introduce millions to the Internet.

Electronic mail

Proprietary email began to lift off in the mid-1980s in a separate market niche to consumer networks. In the 1970s, email had been available on time-sharing systems and occasionally on corporate mainframes. The impact of the PC was to create a new constituency of potential email users consisting of SMEs and individual consumers, as well as a much wider range of users in organizations that had never previously had network access.

In September 1983, MCI introduced the first commercial email service – described as a ‘Digital Post Office’ – aimed at PC users (Cantelon, 1993: 368–383). The service was costly and risky to set up. MCI invested \$32 million building the service, and reportedly spent a further \$30 million in the first year of operation on advertising and promotion (Cantelon, 1993: 251). Competitors quickly piled in – Western Union with Easylink and FedEx with ZapMail, with reported investments of \$250 million and \$100 million, respectively (*Wall Street Journal*, 1984 and 1985).

By the late 1980s, there were some 20 US vendors of email services, including major players such as Sprint Telemail, MCI Mail, Dialcom, and AT&T Mail. Collectively, they had 1.3 million subscribers, each exchanging an average of 20 email messages per month (Barteski, 1988: 48). However, this represented barely 3% of the installed base of 45 million personal computers, and subscribers used email only intermittently. Table 5 describes the main features of the key email providers in 1988.

None of the email services was as successful as initially anticipated because few of them attracted more than 100,000 customers, and those users could not send email messages to the subscribers of other systems. The services were separate ‘islands of communication’ (Guidi, 1987: 18). Establishing gateways between the networks was very expensive because of the lack of a standard protocol. At the turn of the decade, suppliers adopted the CCITT X.400 email protocol and then interconnection speeded up. By the mid-1990s, anyone could send an email to anyone.

Videotex: the consumer internet that almost was

In the 1970s, an interactive communications technology known as Videotex evolved from the one-way text information

services developed for broadcast TV in the previous decade. Early videotex services were accessed using a standard TV set augmented with a keyboard or a dedicated terminal, but these soon gave way to the increasingly ubiquitous personal computer, which required only a software application and a modem to provide access.

Between 1979 and 1984, national videotex systems were established in some 15 countries, including Britain, France, Germany, Canada, Australia, Japan, Norway, Ireland, and Brazil (Sigel, 1980). Conspicuously, the United States did not develop a national videotex system. National videotex systems were developed in complex public–private partnerships, with centralized network infrastructure funded and controlled by national PTTs and supported by private sector information and equipment suppliers. By 1987, the CCITT had designed a videotex inter-networking protocol (the VI protocol), and the stage was set for the networks to consolidate as a worldwide consumer-oriented internet, a decade before *the* Internet took-off (Shimell, 1987: 122).

Videotex systems had a similar trajectory almost everywhere they were developed: initial enthusiasm and large-scale public investment, followed by a failure to live up to initial projections, then a lack-luster service that staggered on for a decade, and finally conversion into a web-based service. The one exception to this general pattern was France, which was not overtaken by the privatization rush of the 1980s and supported a national videotex system, *Télétext*, as a *grand projet* (Fletcher, 2002: 103–107). The government provided basic monochrome ‘Minitel’ terminals free of charge, and by 1988 there were 4.2 million terminals and more than 9500 service providers. Besides an online phone book, the system provided similar services to other consumer networks: classified ads, adult content, and chat rooms.

Why, elsewhere, did the videotex experiments fail? The reasons are complex, but probably the most important one was that the United States did not embrace the videotex idea. If America had bought into videotex, it would have boosted confidence everywhere. However, a publicly funded, national videotex service was politically unthinkable in the United States. There can be no doubt, however, that had videotex gained traction, a consumer-style internet would have been possible a decade before the web.

Table 5 Leading email providers (1988)

Service	Subscription charges	Connect time charge	Minimum message charges	Mailboxes	Messages per month
AT&T Mail	\$30 per year	Free 800 service	40 cents	20,000	400,000
CompuServe’s EasyPlex	\$39.95	\$6–\$12 per hour	None	330,000	6,600,000
Dialcom	\$25 per year	\$15 per hour	5 cents	120,000	2,400,000
McDonnell Douglas’ OnTymePlus	\$300 per month	\$3–\$4.50 per hour	25 cents	60,000	1,200,000
MCI Mail	\$25 per year	N/A	45 cents	90,000	1,800,000
U.S. Sprint’s Telemail	\$15 per month low usage; \$140 corporate	\$7–\$14 per hour	5 cents	100,000	2,000,000
Western Union’s EasyLink	\$25 per month min charge	\$7.50 per hour	45 cents	155,000	3,100,000
Total				1,327,000	21,890,000

Sources: Barteski (1988) and Network World (1988).

From ARPANET to internet

The diffusion of desktop personal computers in the 1980s created a mass market for organizational and domestic online services, which was being supplied by closed corporate networks, consumer networks, email services, bulletin boards, and national videotex services (outside the United States). It was perhaps inevitable that these disparate networks would converge sooner or later, just as the world's telegraph networks converged following the formation of the International Telegraph Union in 1865, and later the world's telephone networks became connected. Not least, network economics helped drive this integration.

There could be only one winner in the four-horse race to define the architecture of the Internet: the CCITT, IBM, OSI, or the ARPA community. In the late 1980s, the OSI 7-layer model was the front-runner. Videotex, which was also a technically viable option, never gained sufficient momentum in the United States to be a realistic challenger. One can identify three reasons for the ascendancy of ARPA's TCP/IP protocols. First, the ARPA Internet community was politically adroit, and as its networking model pervaded universities and research organizations it brought on board a powerful group of early adopters and expert application developers. Second, from about 1991, there was a strong political push – especially from the Clinton-Gore administration – that forced the pace of development (Gore, 1991). While the Clinton-Gore rhetoric was always couched in terms of an 'information superhighway,' and expressed no preference as to protocols, the OSI protocols were not yet ready for deployment so TCP/IP filled the vacuum. Third, the decentralized management of the ARPA-based Internet eliminated most of the bureaucratic friction to joining the network community. If a network wanted to attach itself to the Internet, this was largely a local, cooperative decision – no permission was needed from an over-arching network owner.

Early development of the internet

In the mid-1980s, the ARPA Internet was a substantial network with about 2000 host computers (Quarterman and Hoskins, 1986). Although the ARPA Internet was the 'best-known example' of an academic-research network, it was far from unique and not especially large (Quarterman and Hoskins, 1986: 932). It was a somewhat exclusive community of organizations that were ARPA grant holders.

During the 1980s, several US university consortia created networks to provide their faculty and students with similar facilities to those enjoyed by ARPA grant holders, and several of them grew to comparable size to the ARPA network. In 1986, for example, the UUCP network had more than 7000 hosts and the USENET network more than 2500 (Quarterman and Hoskins, 1986: 936). NSFNET, established by the National Science Foundation (NSF) and operational from 1986 on, was the most important of these coexisting, but isolated, networks. NSFNET provided network facilities to all universities in the United States. This provision was to consume some \$200 million of NSF's budget for computers and infrastructure over a period of a decade (Abbate, 1999: 191). NSF established a number of regional networks and, in a cost-sharing arrangement, made use of the ARPA Internet 'backbone' to carry inter-network traffic.

Table 6 The ARPA Internet: hosts, users, and networks (1969–1997)

<i>Date</i>	<i>Hosts</i>	<i>Users</i>	<i>Networks</i>
December 1969	4	14	N/A
April 1971	23	81	N/A
June 1974	62	217	N/A
March 1977	111	389	N/A
August 1981	213	746	N/A
May 1982	235	823	N/A
August 1983	562	1967	N/A
October 1984	1024	3584	N/A
October 1985	1961	6864	N/A
November 1986	5080	17,780	N/A
December 1987	28,174	98,609	N/A
October 1988	56,000	196,000	N/A
October 1989	159,000	556,500	837
October 1990	313,000	1,095,500	2,063
October 1991	617,000	2,159,500	3,556
October 1992	1,136,000	3,976,000	7,505
October 1993	2,056,000	7,196,000	16,533
October 1994	3,864,000	13,524,000	37,022
July 1995	6,642,000	23,247,000	61,538
July 1996	12,881,000	45,083,500	134,365
July 1997	19,540,000	68,390,000	N/A

Sources: Zakon (1997: 16) and Salus (1995: 217–223); users estimated under the assumption of 3.5 users per IP address and one IP address per host.

In 1987, as network traffic increased, the NSF commissioned IBM and MCI to build a new backbone, using TCP/IP protocols. Because the original ARPA backbone was nearly 20 years old, it was decommissioned and the original ARPA Internet was subsumed into NSFNET. From that point on, it was no longer necessary to prefix Internet with ARPA, and NSFNET effectively became *the* Internet – a government owned network for the US academic and research community. Network effects now began to work their magic, and the number of hosts on the Internet increased exponentially as academic–research networks attached themselves. Table 6 tracks the number of Internet hosts, users, and networks over time.

It is important to recognize that this 'growth' of the Internet represented more a transfer of allegiance than new networking activity. Certainly, because networks were growing rapidly, this was not a zero-sum game, but during the period 1985–1990 most Internet growth represented the attachment to the Internet of existing networks by the installation of new software and protocols, not new networking plant.

Throughout the 1980s, the Internet was a US Government-owned entity and the NSF's conditions of acceptable use forbade commercial activities. The next stage in its development was the establishment of management structures and privatization.

Internet management structures

The organizational history of Internet management is complex, although the principles are simple and persistent (Abbate, 1999: 205–208; Thomas and Wyatt, 1999). In 1979, an Internet Configuration Control Board (ICCB) was

established by ARPA to externalize some of the technical decision making to the wider ARPA Internet community. As the ARPA Internet grew in size, the ICCB was replaced by the Internet Activities Board (IAB), which was more inclusive and brought in representatives from the technical community in universities and research organizations. Membership was by invitation and was representative rather than democratic. The IAB acted as a 'council of elders' overseeing task forces that controlled different technical and administrative activities. Within the task forces, working groups took on individual technical assignments. The most important of the task forces was the Internet Engineering Task Force, which was responsible for the short-term engineering development of the Internet.

Two key innovations of Internet management have been encapsulated by the phrase 'rough consensus and running code,' which has become a motif or motto of the community (Russell, 2006). In order to eliminate the bureaucratic hold-up that can arise from conflicts in conventional standards making bodies, Internet standards were accepted by a consensus (said to be 80–90% of the participation) (Russell, 2006).

The notion of 'running code' was the most significant distinction between the Internet and OSI standardization processes. In the OSI process, technical standards were negotiated, voted upon, and then mandated top-down for implementation in the field. The process was ponderous, subject to occasional hold up, and the standards, although heavily informed by technical expertise, were considered to be somewhat theoretical and detached from engineering reality. By contrast, Internet standards were developed from the bottom up, only after at least two successful implementations had been demonstrated. This process kept standards and their implementations closely coupled.

In 1990, the IAB was replaced by the Internet Society, and the hierarchy of subordinate task forces and working groups continued to evolve, addressing new issues as they surfaced – security, intellectual property, internationalization, and so on.

Privatization

Because the NSF's acceptable use policy forbade commercial activities, privatization was a *sine qua non* of expanding and globalizing the Internet. Commercial providers supplemented the NSFNET backbone with a separate (though connected) national network capacity and privatization was thus achieved without loss of continuity. By summer 1995, there were 'at least 14 national and super-regional high-speed TCP/IP networks in the United States' (MacKie-Mason and Varian, 1997: 31). Simultaneously with this provision, the NSF quietly stepped out of the scene by selling off its assets, a process that was completed by 30 April 1995 at which point the Internet was unequivocally a private entity.

At first, commercial information-technology and communications suppliers had been reluctant to embrace the academic–research Internet because it offered little commercial potential. The NSF, however, by involving commercial entities such as IBM and MCI in infrastructure development, had produced something of a Trojan Horse effect (Abbate, 2001: 172–173). For example, IBM, which

had exclusively pushed its proprietary SNA architecture, perforce developed an expertise in TCP/IP and subsequently became a major supplier of Internet hardware, software, and consulting services. In 1994, MCI offered one of the first Internet consumer services, and by 1997 it claimed to be carrying 40% of all US Internet traffic (MacKie-Mason and Varian, 1997: 31). The NSF also spun off the first Internet Service Providers (ISPs) (Greenstein, 2008: 51–54).

In July 1991, a consortium of three ISPs established the Commercial Internet Exchange (CIX) by which they carried one another's traffic on a mutual basis, enabling them to avoid using NSFNET's infrastructure. Soon dozens of ISPs, both national and international, had joined CIX.

The United States remained the global hub of the Internet, and hence, for example, ISPs in the United Kingdom had to use the United States as the backbone for their national traffic. In November 1994, the London Internet Exchange was formed by five UK-based ISPs. The process of creating national Internet exchanges was mirrored in nations all over the world (OECD, 1998). These Internet exchanges planned and commissioned the growing Internet infrastructure, and agreed accounting processes to allocate costs between the member companies.

The rise of the World Wide Web

By 1990, the Internet based on the TCP/IP protocol was growing fast. The Internet pioneers, always more focused on technology than users, had created a scalable architecture and the network was expanding rapidly without centralized control. The Internet contained millions of potentially useful files and documents, but the network was unfathomable. Without directories and a means of navigation, Internet content was inaccessible to anyone other than the cognoscenti; it was 'like the ancient world before the Library in Alexandria' (Frana, 2004: 24).

Between 1991 and 1995, the means for navigating the Internet was heavily contested by both public and private actors. In 1993, the most popular navigation aid was a system called 'gopher,' and its history is instructive because it shows that the World Wide Web was not the only potential means of organizing the Internet.

Rise and fall of gopher

Gopher was designed at the University of Minnesota as a means of finding documents on the Internet – that is, files that contained human readable text. The system was initially developed for a campus-wide information system at the University of Minnesota.

In a gopher server, files were each provided with a one-line descriptor – such as the title of the report the file contained. Gopher was very much a university activity, and after the system went live in April 1991, other universities took up the software for their campus information systems. After that – somewhat to the surprise of the original developers – the number of gopher servers grew exponentially, not just in universities but also in organizations of every kind.

The experience of using gopher was like browsing a menu-based library catalog. Users would begin at the 'mother gopher' at the University of Minnesota, say, and

from there locate another gopher site likely to lead them toward their goal – the National Weather Service gopher, for example, if they were seeking weather reports. There users would then drill down to locate the document wanted.

The gopher system was superbly adapted for the particular moment that it arrived on the Internet, and this explains its initial popularity. Its menu-driven interface was lean and easy to navigate, it consumed very little bandwidth, and needed minimal infrastructure – the entire University of Minnesota system ran on two desktop Macintosh computers, for example (Frana, 2004: 22). By mid-1992, it was far and away the most popular Internet platform, and it attracted two important complementary services, Veronica and Jughead, both aimed at overcoming the gopher's principal shortcoming – an inability to conduct searches (Glistner, 1994: 77–98). These services enabled a user to enter a search string ('Minneapolis weather,' say) and the system would return a list of candidate documents and their locations. At its peak in 1995, there were more than 9000 gopher servers on the Internet and hundreds of Veronica and Jughead databases.

WAIS (wide area information systems)

WAIS (pronounced 'ways') was probably the first private enterprise Internet investment that addressed content rather than infrastructure (Krol, 1992: 211–226; Aboba, 1993: 255–263; Glistner, 1994: 99–132; Mulner, 1995: 247–250). Like gopher, it ultimately lost out to the web, but also like gopher it played an important role in serving as a model of knowledge organization for the Internet from which the market would eventually select.

WAIS was the concept of Brewster Kahle, then a senior Software Engineer with the Thinking Machines Corporation (a super-computer maker) and later Founder of the Internet Archive. In the WAIS system, for the first time, documents could be selected by their *content* rather than a filename or document description. A WAIS server stored a set of documents, much as a gopher server, but in addition maintained an index of every word in every document. Now a search string such as 'Minneapolis weather' would cause the titles of documents containing those words to be listed. WAIS made use of an existing document retrieval protocol known as Z39.50, and therefore represented one of the first attempts to integrate the Internet with existing information services.

In 1992, WAIS Inc. was established as a joint venture between Thinking Machines, Apple Computer, KPMG, and Dow Jones. By 1992, there were more than 250 free WAIS-compatible databases available on the Internet, as well as several proprietary systems such as the Dow Jones Index and the *Wall Street Journal*. However, with the rise of the World Wide Web in 1993, WAIS lost momentum. It was acquired by AOL in May 1995 for \$15 million, and like several of AOL's acquisitions it is unclear what, if anything, happened to it thereafter.

The World Wide Web

In 1990, the main impediment to effective use of the Internet was the lack of a global directory. During 1992 and 1993, when gopher became popular, global directories were established with the Veronica and Jughead systems. Thus,

gopher-space was already evolving into a satisfactory way of organizing cyberspace and, had the World Wide Web not supervened, it was fully capable of adapting to provide an experience comparable to the World Wide Web.

The web sprang from a scenario similar to the one that gave rise to gopher – the development of a networked organizational information system, in this case for the CERN European Particle Physics Laboratory in Geneva (Berners-Lee, 1999; Gillies and Cailliau, 2000). The inventor of the World Wide Web, Tim Berners-Lee, was born of cultured parents, both of them pioneers of UK computing in the 1950s, and had graduated in physics from Oxford University.

The unique feature of the web, as Berners-Lee conceived it, was that it did not require a directory of any kind, local or global. Instead, documents contained links to other documents. Thus, once one had found one document, it was possible to navigate the information space without a directory. The idea of embedding links in a document in this way is very old – certainly links were explicitly used in Vannevar Bush's memex proposal in 1945. The web is such a close embodiment of Bush's concept, of which Berners-Lee was unaware, that historians inevitably seek a causal influence. The connection appears to be the development of 'hypertext' in the 1960s by the computer scientist (and inventor of the mouse) Douglas Englebart and the computer celebrity and guru Ted Nelson (Bardini, 2000). During the 1970s and 1980s, hypertext was a well-researched area, which had been commercialized in products such as CD-ROM encyclopedias and educational software.

In 1984, Berners-Lee took up a fellowship at CERN where he developed a research interest in hypertext systems as a sideline to his official responsibilities for developing the local computer network. He later described the web as a 'marriage of hypertext and the Internet' (Berners-Lee, 1999: 28). An experimental system, with the potential to organize the Laboratory's online information, was working by late 1990. Berners-Lee and his colleague Robert Cailliau submitted a paper describing the system to the Hypertext '91 conference. The referees turned down the paper, mainly because the system was not particularly novel except that it happened to run on a TCP/IP network (Gillies and Cailliau, 2000: 219). The software was made available for download, however, and the software jackdaws of the Internet began to use it.

By the summer of 1993, there were 130, mostly experimental, web servers worldwide, which was no threat to gopher's 2000 plus. The turning point for the web came with the development of the MOSAIC web browser. MOSAIC was developed at the University of Illinois by a team of programmers led by Marc Andreessen, a gifted computer science undergraduate and later a co-founder of Netscape. The MOSAIC browser was a textbook example of user-friendly, point-and-click software. It was placed on a public server in November 1993, and about a million copies were downloaded (Gillies and Cailliau, 2000: 241).

Why the web won

Several reasons have been advanced for the decline of gopher and the rise of the web, though none is wholly

Table 7 Number of gopher and web servers installed (1992–1995)

Date	Number of gopher servers installed	Number of web servers installed
November 1992	258	N/A
May 1993	1100	50
June 1993	1559	130
July 1993	2018	200
November 1993	4337	270
December 1993	6657	623
April 1994	6958	1681
June 1994	5723	2738
September 1994	4488	6380
December 1994	4773	10,022
January 1995	5057	16,761
June 1995	7052	23,500
July 1995	9046	N/A

Source: Frana (2004: 29); some numbers have been linearly interpolated.

convincing (Frana, 2004: 27–29). It seems that there were multiple causes, but that once a tipping point had been reached, the appeal of the web interface and network effects drove out gopher and all other forms of competition.

One reason gopher's momentum stalled was that the University of Minnesota decided to assert its IP rights, whereas Berners-Lee persuaded CERN to waive its IP rights in the World Wide Web. Another factor in the web's favor was that the browsing experience was much more user-friendly and accessible than a gopher session. MOSAIC and competing web browsers were generally point-and-click, and web pages could contain a variety of font sizes and images – tame by later standards, but even then much more attractive than gopher's plain text documents. In fact, graphical interfaces for gopher were under development and it would have just been a matter of time before more attractive formats (such as PDF documents) came forward (Frana, 2004: 32). As late as mid-1993, informed opinion favored the prosaic gopher over the somewhat gimmicky web.

However, the ease and informality with which documents could be written or repurposed and 'published' on the web was perhaps the single most important factor that led to its dominance. There were almost no barriers to web publishing. Because there was no need to register web documents in a directory, there was no bureaucratic obstacle or corporate permission needed to publish documents. Table 7 tracks the rise of the web and the decline of gopher – in mid-1994, the web overtook gopher as the most popular Internet format.

A central irony of the web is that it took off because the absence of a centralized directory removed all the bureaucratic obstacles that might have prevented individuals from creating web pages. But in practice, directories were needed – without them it would have been even harder to navigate the web than gopher-space, because the information space was completely unstructured and non-hierarchical. In 1994–1995, a number of academic research efforts for cataloging web pages – such as Yahoo!, Lycos, and Excite – made the web navigable. These would be among the first major industries that the web created.

Conclusions

The Internet of today is the computer utility dream of the 1960s made concrete. This article has attempted to show that the evolution of the Internet is a much richer story than portrayed in the standard histories. The Internet did not come out of a vacuum. Its existence owes much to early public investments in military defense systems such as SAGE. Furthermore, the Internet drew on market forays by private sector operators over a period of more than 40 years.

Our history further illustrates the push–pull relationship between infrastructure and applications. We see in these narratives examples of government and commercial applications – such as ATMs and Videotex – that demanded new infrastructure; this infrastructure enabled new applications, adding to further demand for infrastructure, in a virtuous circle. This phenomenon is not confined to the present study, but is to be seen in IT innovations past, present, and future. The spillover effects of infrastructure and application development are as unpredictable as they are economically beneficial. This is why the study of history matters in IT.

We have shown that, contrary to the popular view that interprets today's Internet as growing linearly (and 'inevitably') from the ARPA Internet, the Internet as we know it has a considerably more complex history. At various points in time, 'dominant designs' emerged (in crucial historical junctures) from the confluence of social and market forces, and these designs in turn contributed to shape today's Internet.

During the 1970s, for example, packet switching replaced circuit switching as the dominant design in data communications. The packet-switching technology, which reduced data transmission costs by an order of magnitude, served as the foundation for the emergence of public networks such as Tymnet and Telenet. These networks, in turn, boosted the expansion of online information services, such as Dialog and Nexis, and consumer online networks, such as CompuServe.

In the early 1990s, the TCP/IP protocol emerged as the dominant design in the inter-networking standards battle. This protocol was an *ad hoc* solution like many others, and its rise to dominance was completely unexpected at the time. The ARPA network itself was one among many academic and research networks that could have given rise to the Internet. The ARPA networking model benefited from having a critical mass of early adopters in universities and research centers and from the fact that, when the 'information superhighway' rhetoric came to the forefront, the OSI model was not ready for prime time. The lack of bureaucratic obstacles to joining the ARPA Internet played a key role as well in boosting the network's expansion. Once the TCP/IP-based ARPA network (subsumed into NSFNET in 1990) started emerging as 'the' leader among the many academic and research networks in existence, network effects took over and further solidified its dominance.

Finally, in 1993 the World Wide Web emerged as the dominant design in terms of Internet organization. Just like the lack of bureaucratic obstacles to joining boosted the rise of the ARPA Internet, the ease and informality with which documents could be published on the web explain why the web won. Its triumph, however, was anything but

a sure thing, and as late as mid-1993 informed observers tended to favor gopher over the web.

We have interpreted the rise of these dominant designs within the economic framework of network industries. We have highlighted junctures in the history of the Internet in which initial conditions and small historical accidents appear to have played a key role in shaping historical outcomes.

The core message of this study is *not* that a counterfactual Internet based on, say, IBM's SNA and gopher would have been better than the Internet as we know it but that the features of today's Internet are anything but inevitable. History helps us understand how key information systems of today – could there be any more crucial than the Internet? – have become what they are. And thus a careful study of history may be required in order to understand what needs to be done in order to improve these systems in the future.

References

- Abbate, J. (1999). *Inventing the Internet*, Cambridge: The MIT Press.
- Abbate, J. (2001). Government, Business, and the Making of the Internet, *Business History Review* 75(1): 147–176.
- Aboba, B. (1993). *The Online User's Encyclopedia: Bulletin boards and beyond*, Reading: Addison-Wesley.
- AOL (1990–1996). *10-Q Reports*.
- Bardini, T. (2000). *Bootstrapping: Douglas englebart, coevolution, and the origins of personal computing*, Stanford: Stanford University Press.
- Barnett, C., Anderson, B. and Bancroft, W. (1967). *The Future of the Computer Utility*, New York: American Management Association.
- Barteski, E. (1988). A Do-It-Yourself Guide to E-Mail, *Business Credit* (January): 47–50.
- Berners-Lee, T. (1999). *Weaving the Web: The past, present and future of the World Wide Web by its inventor*, London: Harper Paperbacks.
- Bourne, C. and Hahn, T.B. (2003). *A History of Online Information Services, 1963–1976*, Cambridge: The MIT Press.
- Bowen, C. and Peyton, D. (1984). *How to Get the Most Out of CompuServe*, New York: Bantam Books.
- Brown, S. (1997). *Revolution at the Checkout Counter: The explosion of the bar code*, Cambridge: Harvard University Wertheim Committee.
- Bush, V. (1945a). As We May Think, *Atlantic Monthly* (July): 101–108.
- Bush, V. (1945b). As We May Think, *Life* (Sept. 10): 112–114, 116, 118, 121 and 123–124.
- Caminer, D., Aris, J., Hermon, P. and Land, F. (1998). *LEO: The incredible story of the world's first business computer*, New York: McGraw-Hill.
- Campbell-Kelly, M. (2003). *From Airline Reservations to Sonic the Hedgehog*, Cambridge: The MIT Press.
- Campbell-Kelly, M. and Garcia-Swartz, D. (2008). Economic Perspectives on the History of the Computer Time-Sharing Industry, *IEEE Annals of the History of Computing* 30(1): 16–36.
- Campbell-Kelly, M., Garcia-Swartz, D. and Layne-Farrar, A. (2008). The Evolution of Network Industries: Lessons from the conquest of the online frontier, 1975–1995, *Industry and Innovation* 15(4): 435–455.
- Cantelon, P. (1993). *The History of MCI: The early years, 1968–1988*, Dallas: MCI Communications Corporation.
- CompuServe (1997). *Annual Report*.
- Computer Economics (1989). IBM's System Network Architecture in the 1990s, *Computer Economics* 11(5): 1–4.
- Cortada, J. (2004). *The Digital Hand: How computers changed the work of American manufacturing, transportation and retail industries*, Oxford: Oxford University Press.
- Datamation (1962). A Survey of Airline Reservation Systems, *Datamation* 8(August): 53–55.
- Deasington, R. (1985). *X.25 Explained: Protocols for packet switching networks*, Chichester: Ellis Horwood.
- DeLamarter, R. (1986). *Big Blue: IBM's use and abuse of power*, New York: Dodd, Mead & Company.
- Edwards, P. (1996). *The Closed World: Computers and the politics of discourse in Cold War America*, Cambridge: The MIT Press.
- Fano, R. and Corbato, P. (1966). Time-Sharing on Computers, in *Scientific American, Information*, San Francisco: W.H. Freeman.
- Farrell, J. and Klemperer, P. (2007). Coordination and Lock-In: Competition with switching costs and network effects, in M. Armstrong and R. Porter (eds.) *Handbook of Industrial Organization*, Volume 3 Amsterdam: North-Holland.
- Flamm, K. (1987). *Targeting the Computer: Government support and international competition*, Washington DC: Brookings Institution Press.
- Flamm, K. (1988). *Creating the Computer: Government, industry, and high technology*, Washington DC: Brookings Institution Press.
- Fletcher, A. (2002). France Enters the Information Age: A political history of Minitel, *History and Technology* 18(2): 103–107.
- Frana, P. (2004). Before the Web There was Gopher, *IEEE Annals of the History of Computing* 26(1): 20–41.
- Gannon, P. (1997). *Trojan Horses and National Champions: The crisis in Europe's computing and telecommunications industry*, London: Apt-Amatic.
- GE Information Services (1985). *Spectrum: A special edition commemorating the twentieth anniversary of general electric information services company*, December.
- Gillies, J. and Cailliau, R. (2000). *How the Web was Born*, Oxford: Oxford University Press.
- Glister, P. (1994). *Finding it on the Internet*, New York: Wiley.
- Glossbrenner, A. (1983). *The Complete Handbook of Personal Computer Communications*, New York: St. Martin's Press.
- Glossbrenner, A. (1985). *The Complete Handbook of Personal Computer Communications*, 2nd edn, New York: St. Martin's Press.
- Glossbrenner, A. (1990). *The Complete Handbook of Personal Computer Communications*, 3rd edn, New York: St. Martin's Press.
- Glossbrenner, A. (1995). *The Little Online Book*, Berkeley: Peachpit Press.
- Gore, A. (1991). Infrastructure for the Global Village, *Scientific American* (September): 108–111.
- Greenberger, M. (1964). The Computers of Tomorrow, *Atlantic Monthly* (July): 63–67.
- Greenstein, S. (2008). Innovation and Evolution of Market Structure for Internet Access in the United States, in W. Aspray and P. Ceruzzi (eds.) *The Internet and American Business*, Cambridge: The MIT Press.
- Guidi, P. (1987). Electronic Email Begins a New Era, *Telecommunication Products and Technology* (December): 16–20.
- Hafner, K. (2001). *The Well: A story of love, death and real life in the seminal online community*, New York: Carroll & Graf Publishers.
- Hafner, K. and Lyon, M. (1996). *Where Wizards Stay Up Late*, New York: Simon & Schuster.
- H&R Block (1983–1992). *Annual Reports*.
- Jakobs, K. (2013). Why Then Did X.400 Fail? *Journal of Information Technology* 28(1): 63–73.
- Jones, R. (1984). SNA Review, *Data Processing* 26(4): 5–7.
- Kimberley, P. (1990). *Electronic Data Interchange*, New York: McGraw-Hill.
- Krol, E. (ed.) (1992). *The Whole Internet: User's guide and catalog*, Sebastopol: O'Reilly and Associates.
- Lamoreaux, N., Raff, D. and Temin, P. (2004). Against Whig History, *Enterprise and Society* 5(3): 376–387.
- Levering, R., Katz, M. and Moskowitz, M. (1984). *The Computer Entrepreneurs*, New York: Dutton Adult.
- Licklider, J. (1960). Man Computer Symbiosis, *IRE Transactions on Human Factors in Electronics* (March): 4–11.
- Lottor, M. (1992). Internet Growth (1981–1991), Network Working Group RFC 1296, <http://tools.ietf.org/rfc/rfc1296.txt>, accessed 30 January 2013.
- Lukasik, S. (2011). Why the Arpanet was Built, *IEEE Annals of the History of Computing* 33(3): 4–21.
- MacKie-Mason, J. and Varian, H. (1997). Economic FAQs about the Internet, in L. McKnight and J. Bailey (eds.) *Internet Economics*, Cambridge: The MIT Press.
- Malamud, C. (1993). *Exploring the Internet: A technical travelogue*, Englewood Cliffs: Prentice-Hall.
- Martin, J. (1987). *SNA: IBM's networking solution*, Englewood Cliffs: Prentice Hall.



- McKenney, J., Copeland, D. and Mason, R. (1995). *Waves of Change: Business evolution through information technology*, Boston: Harvard Business Press.
- Mulner, H. (1995). WAIS and Information Retrieval on the Internet, *Journal of Universal Computer Science* 1(4): 247–250.
- Network World (1988). Easy Email: Modern electronic mail products and services, *Network World* (December 5): 35.
- Norberg, A. and O'Neill, J. (2000). *Transforming Computer Technology: Information processing for the pentagon, 1962–1986*, Baltimore: The Johns Hopkins University Press.
- OECD (1998). *Internet Traffic Exchange*, Paris: OECD.
- Parkhill, D.F. (1966). *The Challenge of the Computer Utility*, Reading: Addison-Wesley.
- Phister, M. (1979). *Data Processing: Technology and economics*, Bedford: Digital Press.
- Plugge, W. and Perry, M. (1961). American Airlines' 'SABRE' Electronic Reservations System, Proceedings of the AFIPS, Western Joint Computer Conference. Washington DC: 592–601.
- Quarterman, J. (1990). *The Matrix: Computer networks and conferencing systems worldwide*, Boston: Digital Press.
- Quarterman, J. and Hoskins, J. (1986). Notable Computer Networks, *Communications of the ACM* 29(10): 932–971.
- Reingold, H. (1993). *The Virtual Community*, Reading: Addison-Wesley.
- Roberts, L. (1978). The Evolution of Packet Switching, *Proceedings of the IEEE* 66(11): 1307–1313.
- Russell, A. (2006). Rough Consensus and Running Code' and the Internet-OSI Standards War, *IEEE Annals of the History of Computing* 28(3): 48–61.
- Salus, P. (1995). *Castig the Net*, Reading: Addison-Wesley.
- Shapiro, C. and Varian, H. (1998). *Information Rules: A strategic guide to the network economy*, Boston: Harvard Business Press.
- Shimell, P. (1987). A Key to Open the Videotex Marketplace, *Telephone Engineer and Management* (15 April): 122.
- Sigel, E. (ed.) (1980). *Videotext: The coming revolution in home/office information retrieval*, New York: Knowledge Industry Publications.
- Swisher, K. (1998). *Aol.com*, New York: Three Rivers Press.
- Thomas, G. and Wyatt, S. (1999). Shaping Cyberspace – Interpreting and transforming the internet, *Research Policy* 28(7): 681–698.
- Tymshare (1979). *A Tymshare, Inc. Presentation for the New York society of security analysts*, New York: Tymshare.
- Waldrop, M. (2001). *The Dream Machine: J.C.R. Licklider and the revolution that made computing personal*, New York: Penguin.
- Wall Street Journal (1984). Western Union Expects Electronic-Mail to Break Even in 1985, *Wall Street Journal* (October 31): 1.
- Wall Street Journal (1985). Federal Express Specifies the Loss ZapMail Incurred, *Wall Street Journal* (January 18): 1.
- Wells, H.G. (1938). *World Brain*, London: Ayer.
- Yates, D. (1997). *Turing's Legacy: A history of computing at the national physical laboratory 1945–1995*, London: Science Museum.
- Zachary, G.P. (1997). *Endless Frontier: Vannevar Bush, engineer of the American century*, New York: Free Press.
- Zakon, R. (1997). Hobbes' Internet Timeline, Network Working Group RFC 2235, <http://www.ietf.org/rfc/rfc2235.txt>, accessed 30 January 2013.

About the authors

Martin Campbell-Kelly is an Emeritus Professor in the Department of Computer Science at the University of Warwick, where he specializes in the history of computing and the IT industries. His books include *Computer: A History of the Information Machine* (3rd edn, 2013) co-authored with William Aspray and others, *From Airline Reservations to Sonic the Hedgehog: A History of the Software Industry* (2003) and *ICL: A Business and Technical History* (1989). He is a Member of the editorial board of the *IEEE Annals of the History of Computing*.

Daniel Garcia-Swartz is an Economic Consultant with Compass Lexecon in Chicago. He holds a joint Ph.D. in Economics and History from the University of Chicago. His publications focus on the economics of high-technology industries, including computer hardware and software. He is currently working with Martin Campbell-Kelly on a book on the economic and business history of the international computer industry.