

Nurturing the Accumulation of Innovations: Lessons from the Internet

by Shane Greenstein¹

As the Internet diffused throughout the 1990s, it touched a wide breadth of economic activities. The diffusion transformed the use of information technology throughout the economy. It led to improvements in products, lower prices, the development of new capabilities, and the development of many innovations that enabled productivity improvements among business users. It diffused to the majority of homes and businesses, altering the way people shop, research, play, and relate socially.

The Internet began as a government sponsored operation in the 1970s and 1980s and grew into a commercial industry in the 1990s. At first, the Internet lacked market-oriented focusing devices and/or economic inducement mechanisms typically associated with directing efforts toward the most valuable innovative outcomes.² There were contracts for carrier services between government buyers and commercial suppliers, for example, but no general market orientation towards the pricing of the exchange of traffic between carriers. There also were a few providers of Internet equipment for government users, but no waves of inventive entrepreneurial entry. There were managers who understood the specific needs of their niche user communities, but no possibility for tailoring new products and services to every potential new set of users.

How could an institutional setting that lacked market-orientation yield a set of innovations that supported the creation of massive market value only a few years later? This chapter helps explain the progression. The chapter divides the Internet's development into a pre-commercial and commercial era. The pre-commercial period encompasses the 1970s and 1980s and some of the early 1990s, when the government controlled the research and development of the Internet and its components. The commercial period, which arose after the government opened control of the network to commercial interests, encompasses the mid-1990s and onward to present day.

These two eras illustrate two distinct models for accumulating innovations over the long haul. The pre-commercial era illustrates the operation of several useful non-market institutional arrangements. It also illustrates a potential drawback to government sponsorship – in this instance, truncation of exploratory activity. The commercial era illustrates a rather different set of lessons. It highlights the extraordinary power of market-oriented and widely distributed investment and adoption, which illustrates the power of market experimentation to foster innovative activity. It also illustrates a few of the conditions necessary to unleash such growth, such as standards development and competition, and nurturing legal and regulatory policies.

The Pre-commercial Internet under DARPA³

It may be tempting to compare the Internet to historically archetypical big inventions sponsored by government, such as the Manhattan and the Apollo projects. However, these archetypes for developing technical breakthroughs simply are not good models for understanding what happened during the creation of the Internet. The Internet was not a single urgent project in a single lab devoted to engineering a single object. In fact, the early development of the Internet is rather less exciting than commonly assumed.

The Internet began as a government sponsored project with a restricted set of users and uses. It occurred at a time when theory pointed in a few promising directions, but nobody, not even the experts, knew where implementation would lead. The project involved a vastly dispersed set of technically adept

¹ Elinor and Wendell Hobbs Professor, Kellogg School of Management, Northwestern University. I thank Guy Arie, Tim Bresnahan, David Clark, David Crocker, Rebecca Henderson, Franco Malerba, Richard Newell, Bonnie Nevel, John Quarterman, Craig Partridge, Richard Schmalensee, Alicia Shems, Scott Stern, and Stephen Wolff for extraordinarily useful conversations and comments. This essay is part of a larger project funded by the Kaufman Foundation, the Searle Center at Northwestern University, and the Dean's Office at the Kellogg School of Management at Northwestern. I am grateful for the funding. I am responsible for all remaining errors.

² For more on focusing devices in general, see e.g., Rosenberg (1977).

³ Much of the material in this first section summarizes Greenstein (2009a).

participants with a shared interest in the project, but otherwise heterogeneous needs and outlooks. The Internet developed slowly, and through a rather mundane process, accumulating capabilities over time from a vast number of contributors.

The Internet's early development fit into an archetype called "collective invention." Collective invention is "a process in which improvements or experimental findings about a production process or tool are regularly shared."⁴ There was no single user who demanded a technology such as the Internet, nor any single inventor for it. Rather, five partially overlapping groups played a role in shaping attributes that each valued, with an accumulation of innovative contributions over time. The first two groups were the primary decision makers at funding agencies: the Department of Defense and the National Science Foundation. The other three were programmers/inventors, administrators, and application users. Many were funded by the government agencies and given considerable discretion. Others became participants over time and added their own contributions within their own budgetary limitations.

The U.S. military budget served as the first source of funds for the pre-commercial Internet, while the National Science Foundation largely served as the source of funds from 1986 to the end of government involvement in 1995. At a very basic level, the U.S. government paid for much of the research and development of the Internet during this period, and the government was the organization behind the early "demand" for the breakthrough technical achievements that became recognized as the Internet.

This effort did not begin like a military procurement project, as if the Internet were a military rocket procured from several suppliers. In such a procurement process, a group of expert U.S. military personnel consider and deliberate at great length with great foresight about the needs of the government, and then they issue a set of specifications for a product or service with a pre-determined set of attributes. However, in the case of the Internet, the Department of Defense's Department of Advanced Research Projects Agency (DARPA),⁵ funded the core research and development work that led to the Internet – indeed, this work was but one of many DARPA projects on the frontiers of computer science.⁶ The project first intended to build prototypes for a packet-switching data-communications network of networks, which pushed the boundaries of computing at the time. Both "packet-switching" and "a network of networks" were budding theoretic concepts. DARPA's administrators wanted innovations in the form of ideas, new designs, and new software. The administrators also desired all of these innovations be portable to military operations in the long run, as required under the Mansfield Amendment (stipulating that Department of Defense funding be relevant to military's mission). While the demand for these innovative solutions was quite general, the specifics were undetermined. The U.S. military faced issues with its own computing facilities and operations that justified the research and development expense, even though its own managers could not concretely describe the object that would result.⁷

The military sought a robust design for a communications network, and its potential value was self-evident. Keeping communications functioning in spite of a blown/cut line, for example, has military value in hostile battlefield conditions. Packet switching held the promise to achieve this attribute by allowing data to flow along multiple paths, unlike a circuit-switched telephone network in which calls follow a pre-set path programmed into central office telephone switches. In principle, an inexpensive packet switching network could also cover vast geographic distances, which could support the sharing of expensive computing resources over such distances. That too had self-evident military value. For example, military users in many locations – even potentially dangerous locations – could access databases housed in another (potentially safer) location.

⁴ See Allen (1983) for the introduction and illustration of the concept. The quote is from Meyer (2003), who provides further extensions to modern examples.

⁵ This organization was originally founded as the Advanced Research Projects Agency (ARPA), and became DARPA in 1972. For the sake of simplicity, I use only the name DARPA throughout this chapter.

⁶ Norberg et al. (1996), Roland and Shiman (2002).

⁷ Norberg et al. (1996) repeatedly stress that DARPA's funding of packet switching research in the 1960s and 1970s met concerns about whether the funding was relevant to military mission, as required by the Mansfield Amendment, which was proposed several times, and eventually passed in 1973. The research anticipated enhancing the "command and control" capabilities of commanders increasingly reliant on their computing resources.

An additional technical and pragmatic aspiration also played a role. An ideal network of networks could facilitate the movement of data between computer systems – mostly mainframes in this era – that otherwise could not interoperate seamlessly. A system that could enable the exchange data and communication between computing systems without frequent human intervention would save the military time and personnel expenses, and help realize new strategic capabilities. Coordinating the exchange, combination, and filtering of data between computer systems generated numerous logistical and organizational problems for military operations that increasingly depended on computing.

Although these innovations later would be portable to non-defense uses, that was not among the relevant criteria for the program at the outset. Initially there was no explicit requirement that the innovation work with all (or even most) computer systems in non-military uses, although that was a likely byproduct, as non-military uses of computing overlapped in some applications and functionality with military uses. However, the program was informed by a general understanding that shaped all activities in information technology within DARPA: having a healthy U.S. information technology sector was a valuable military advantage in the long run.

Eventually several prototypes for a packet switching network were engineered. With additional work, these innovative designs turned into a prototype of an operating network, operated by managers from Bolt Beranek and Newman (BBN), a research contractor that sub-contracted through DARPA. A number of researchers and their students became familiar with its design and operational principles. The network grew from there, covering more locations and more participants throughout the 1970s. The network extended into many research laboratories and universities with funding from DARPA.

A skunk works and wild ducks

From the outset when DARPA's "packet switching" and "network of networks" project began, the desired attributes for the Internet represented a radical technological departure from existing practice. To understand the development of radical technologies within the military, it is best to begin with an understanding of two terms: a skunk works and wild ducks.

A *skunk works* is an organizational home for frontier development projects.⁸ It is housed away from the main operations of an organization, sometimes in secret or with organizational barriers, and often with top management support for these barriers. Typically the development projects involve something of value to the future of the organization but are not directly connected to its present operational or service missions. In this case, DARPA itself is the military's skunk works. The mission of the agency is research oriented, not operational, although in the case of the Internet operational issues eventually became salient as well. Broadly construed, DARPA's mission is to develop radical new concepts and operations to transform military operations through development of new technologies. The agency had been established after Sputnik, and it was deliberately not beholden to the short run operational needs of any of the armed services, although its innovations were required to eventually enhance some military function.

Wild ducks are a particular group of technically adept and innovative contributors, often considered social outsiders by those controlling funding. Wild ducks can encompass a range of behaviors and social differences that are regarded as potentially disruptive and costly to the regular operations of an organization. The reverse also often holds – that is, wild ducks often regard the practices of those involved in regular operations as interfering with their inventive activity. During DARPA's research into what became the Internet, the inventive individuals behind it were largely wild ducks.

Wild ducks had been a colloquial term of art in computing for decades, coined by IBM CEO

⁸ The phrase originated from a project for the Air Force at a division of Lockheed Martin. The division had called itself the "Skonk Works" after a phrase from Al Capp's *Lil' Abner* cartoon – the skonk works was a "secret laboratory" that operated a backwoods still. The label became well known throughout the industry, in part because it was considered humorous and saucy. *Lil' Abner*'s publisher eventually asked Lockheed Martin to change it, and "skunk works" emerged from there (Rich and Janus 1994).

Thomas Watson Sr. to describe the innovative practices of his technical team.⁹ At IBM, wild ducks were avant-garde thinkers in their field, often chasing visions they saw as aesthetic. The value of their ideas could defy evaluation *ex-ante*, and in many cases, it was hard to evaluate even after a prototype was developed. However, to fully realize the innovative potential of wild ducks and allow them to coexist within a mainstream organization, IBM separated the wild ducks from others, which kept valuable inventions temporarily hidden, unthreatening to others in the business, and, therefore, flowing until needed. IBM used wild ducks to develop innovative products and used the mainstream sales force to systematically and uniformly sell these products. IBM thereby kept control of the computing platform and ensured its commercial success by making sure it did not remain static.

The wild ducks arrangement worked well for the Department of Defense. If the wild ducks failed to realize their grandest innovative vision, then almost nobody had to be bothered. If, on the other hand, the wild ducks invented something that others within the mainstream organization could appreciate and use, then the best scenario would be achieved. However, there was one troubling scenario: What would the mainstream organization do if the most valuable inventions could not be integrated into existing operations? In this case, the wild duck arrangement allowed the military to at least defer questions about a costly integration until the time when or if the innovation proved fruitful. If the need arose, then such knotty questions would have to be addressed, but not before then.

As it turned out, a particularly inventive group of wild ducks in the pre-commercial Internet accumulated a range of inventions, eventually bringing about a large economic gain for participants. More to the point, these innovations, which comprised the basic building blocks for the Internet, turned out to have enormous value when transferred to non-participants. Some of these inventors were established university researchers (such as Paul Baran, Joseph Licklider, and Leonard Kleinrock), and their reputations would be further enhanced by their involvement in designing packet switching networks. Other researchers (including Steve and David Crocker, David Clark, John Postel, and Vint Cerf, and many others) would become affiliated with the project right at the outset of their careers and remain involved throughout their careers. The Internet Protocol Suite known as TCP/IP (Transmission Control Protocol and Internet Protocol) emerged from their efforts.¹⁰ Their achievements would be recognized by contemporaries, and they gained reputations over time from those achievements.

DARPA's program for fostering innovations in computing departed from the archetype of a skunk works practiced among military contractors.¹¹ The continuity in DARPA's managerial procedures and policies borrowed considerably from practices for R&D and military procurement, melding them into a goal-oriented research and development project administered by technically capable program officers executing a general vision. One key departure involved the amount of discretion given program officers. Though they were reviewed eventually, in the short run many had freedom to make the decisions they thought would work best. Another principal departure involved geography. DARPA's skunk works was not physically housed in a single organization in Washington, D.C. Instead, it was *administered* from D.C., but the work was geographically dispersed to many locations in research organizations and universities across the country. DARPA sent money for projects organized by key researchers, who maintained their laboratories at their own universities. Money was also sent to contracting research organizations, such as BBN (in Cambridge, MA), the Rand Corporation (in Santa Monica, CA), and Stanford Research Institute (in Menlo Park, CA). While DARPA provided funds to support labs, buy equipment, and pay graduate students at these locations, the government agency was able to take

⁹ Watson encouraged social conformity in his firm because he believed it made his sales force more effective (for example, all salesmen had to wear blue suits). But Watson came to a different understanding with his technical talent. His wild ducks had permission to be diverse, so long as they invented. For many stories related to wild ducks, see Maney (2003).

¹⁰ TCP determines a set of procedures for moving data across a network and what to do when problems arise. If there are errors or specific congestion issues, TCP contains procedures for retransmitting the data. While serving the same function of a postal envelope and address, IP also shapes the format of the message inside. It specifies the address for the packet, its origin and destination, a few details about how the message format works, and, in conjunction with routers, the likely path for the packet towards its destination. An extensive explanation of TCP/IP can be found in many publications. See, e.g., Leiner et al. (2003) or Abbate (1999).

¹¹ Norberg et al. (1996).

advantage of building on what was already there, both in terms of institutions and brainpower.

Dispersed geography mattered in several other ways. Innovative improvements arose and accumulated in different places, yielding a variety of lessons and insights at a time when theory pointed in many directions and implementations were scarce. Collectively this program began accumulating improvements and suggestions from a diversity of sources, which were loosely coupled to one another through their common funding source and, non-trivially, shared scientific and engineering goals.

Program officers encouraged this sharing through arranged face-to-face meetings and communications.¹² Despite the geographic dispersion, participants shared a sense of identity about the whole project, and they were encouraged to share innovations with one another. Inventors also were encouraged to pay close attention to how their meta-design facilitated inventive specialization across the entire program. In addition, participants developed norms for documentation to facilitate knowledge retention and improvements built on earlier advances.

Unlike typical project management, program officers in this case did not initially rely on some of institutions typically affiliated with academic science, such as peer-review, formal proposals with multiple stages of review, and panels of reviewers. They did solicit research proposals on occasion, but not necessarily proposals promising specific incremental advances within short time horizons. Instead, the program officers often asked for short broad proposals, picked stars, made general agreements with them about the long term goals, funded their labs with uncommonly large amounts of money (for the discipline at the time), and gave them large amounts of discretion to pursue those goals in the manner they saw fit. In exchange for this funding, the researchers were required to attempt technically ambitious projects, participate in certain conferences, document and share their results with each other, and contribute to the training of a new generation of researchers, among other things.

Large sums of money invested by DOD sustained continuity in its operation and continued improvement. The level of funding is notable because no program officer ever asked for concrete invention on a specific time frame, for example, and most of the inventors would have considered meeting such requirements to be pointless and absurd bureaucratic milestones. At the same time, many program officers were technically sophisticated enough to follow specific advanced developments. Some of them even contributed inventions to the efforts. In fact, DOD program officers often did the evaluation themselves or with a small set of consultations, and not necessarily using evaluation by peers.

Nurturing useful prototypes

Precisely because a skunk works seeks to break with established patterns to facilitate experimentation and protect it from the objections of other organizations or their parent entity, a skunk works faces numerous challenges meeting existing user needs. Its challenges are even greater when the participants in the skunk works create inventions for needs that most potential users have not yet even recognized. In the case of the DARPA Internet project, however, innovation and operations began to overlap. As a result, instead of meeting bureaucratic requirements, inventors were held to a different test: they had to eat what they grew. That is, their innovations were put into use comparatively quickly. The overlap between operations and invention played a key role in fostering *useful* innovations.

The first and second generation of Internet researchers¹³ quickly became familiar with a second unusual feature of their skunk works: new ideas grew out of their own experiences and often stemmed from their own needs. Because inventors were also users, they were motivated to develop working prototypes into operational pieces that they and others could employ. Working prototypes were crude models of innovation in need of refined improvement. Often oriented towards demonstrating the proof of

¹² Building coherent scientific communities around nascent technologies was an explicit part of the mission of every program officer in this era. See, e.g., Norberg et al. (1996), and Roland and Shiman (2002).

¹³ There is no clean line between generations. The “the first generation” of Internet researchers grappled with engineering, creating the first packet switching applications and prototypes, and demonstrating the viability of the concepts. The second generation contributed to the existing infrastructure, and, along with the first generation, built applications and scale (Crocker 2008b).

a new concept, these were deliberate interim manifestations of proposals, aimed to explore and, if possible, solve a piece of a problem and to help the inventors learn. Through their own use, many of these inventors became interested in issues that moved beyond simply illustrating a concept with their prototypes. They were introduced to issues associated with refining and maintaining workable versions of their inventions in a functioning and operational network—and not just any network, but a network they all used.

In the short run, mixing inventive activities with operational activities had a very direct effect on orienting innovation. Although using a common network, each group of researchers began working in its own direction, with its own working prototypes, for its own use as well as use by others. Due to their common affiliation with DARPA and common use of the network (which became known as the DARPAnet), the researchers began to make their prototypes interoperate with each other.

Many analysts of computing markets today stress the importance of a “killer application” – an application so compelling it justifies complementary investments. Early Internet innovators quickly developed several killer applications -- file transfer, something close to what we today recognize as instant messaging, and electronic communication that became electronic mail.¹⁴ Arguably, electronic mail was not even the most central innovation of the skunk works, but it was one that every participant used. Its pragmatic value was widely recognized among participants. More than 50 people made important contributions to the standard e-mail design in the 1970s and 1980s, and by the end of the decade virtually all participants in the Internet made use of this design. Another lesson from the e-mail application innovation is that its usefulness was apparent at the time to the innovators but not to the sponsoring federal agency. As stated by Bob Kahn, DARPA “...would never have funded a computer network in order to facilitate e-mail” because other goals were more paramount, and person-to-person communication over telephones appeared sufficient.¹⁵

The spread of e-mail highlights the essential paradox of a skunk works: protecting wild ducks leads to long term benefits if the inventions get pointed in useful directions from an early stage. However, at an early stage virtually nobody in an organization except the most technically sophisticated is able to assess whether the wild ducks have succeeded in moving in a useful direction or in achieving even the most basic milestones!

How did participants make such assessments then, particularly into the late 1970s, after the basic science was done but considerable room was left for implementing new improvements? The integration of innovations into immediate operation shaped the consensus about innovations and helped determine whether suggestions for new protocols merited attention. As improvements arose, those improvements became gradually embedded in routine processes. If installation administrators did not think the innovations useful, they did not get installed nor used. If they did get used, the inventions got refined and began to accumulate additional improvements.

One additional aspect of this experience deserves attention: the DARPA skunk works was a technical meritocracy. In a technical meritocracy, individuals advance in standing through commonly recognized technical achievements, rather than by external credentialing. The technical meritocracy for the Internet survived as an informal consensus process in the 1970s and much of the 1980s. The meritocracy survived for several reasons. First, virtually all participants came out of an academic and research background. They found it natural to work within a technical meritocracy, developing consensus about improvements worth keeping. Second, most of the program officers shared this research background, and they justified their actions on a similar basis. Third, in any given area the group of researchers and administrators tended to be small enough that a technical consensus could emerge comparatively quickly. Fourth, and crucially, the top managers in the Department of Defense protected DARPA against other influences on decision making, such as promotion of researchers or their projects based on favoritism, political connections, seniority, fame, or other criteria different from that of a

¹⁴ See Partridge (2008), Crocker (2008a), and Living Internet (2009) for documentation of how subsequent technical improvements built on one another, beginning with an early project at the RAND Corporation in Los Angeles. These passages draw heavily on accounts by Partridge and Crocker.

¹⁵ See Segaller (1998), page 105.

technical meritocracy.

Note, however, that this technical meritocracy was pragmatic in its orientation because innovation was put into use. If installation administrators did not find it useful, then the invention did not get used. If it did not get used, the invention did not get refined, or accumulate additional improvements.

The Internet under the NSF

The example of the DARPA Internet project contains lessons of certain behavior to avoid. First and foremost, the DOD restricted participation in the use of the results from the DARPA experiments. These restrictions truncated the range of uses to which the technology could be put by truncating the set of users who could experiment with it. DARPA partially recognized this limitation, and it eventually permitted its contractor BBN to spin off a division and start a packet-switching service in the early 1970s.¹⁶

The issue kept returning, however, in part because DARPA sponsored experiments that succeeded more than anticipated. By the end of the 1970s, the DARPA network was operational, and though far from perfect, the key pieces of the engineering insights had moved far beyond their status as working prototypes. It connected a network of research contractors and university researchers who wanted to continue to collaborate with each other. The inventions were portable to others, who could (and did) independently design and operate their own networks. In fact, frustrations with gaining access to the DARPA network motivated some participants to start their own networks. For example, both Bitnet and CSNet began in the early 1980s, partly as a response to restricted access to DARPA. Both of these spinoff networks provided the functionality that users desired, enabling them to move data between computers in different locations, supporting file sharing, and enabling electronic mail. Each had a different architecture and rules for participation, however. CSNet aspired to provide connectivity only to computer science departments, while Bitnet connected computing systems between various researchers and universities. A third network at the time was more informally organized, and went by the label, UUCPNet or Usenet, and involved numerous participants both from inside universities and outside.

The increasing growth of alternative networks showed that such connectivity interested numerous participants other than the military. The network had grown beyond the core concerns of the military. Eventually, more researchers wanted to participate than DARPA had an interest in supporting. DARPA also worried about compromising the security of its own network by allowing non-military users to participate.

By the early 1980s, the limits of participation became a source of tension. DARPA finally in 1985 handed over control of part of the network to the National Science Foundation (NSF) in order to open it to the many civilian researchers interested in using it.¹⁷ By then the community of innovators had evolved into a loose confederation of researchers from many locations, so this administrative change simply ratified what had already begun to happen informally.

Innovation under NSF funding differed in several respects from innovation under DARPA funding. This is not surprising, because the missions of the two organizations differ. DARPA is part of the defense department, while the NSF supports civilian research. Just as with DARPA, no requirement about an immediate civilian application shaped activities other than a general understanding that the NSF's needs could be met more easily and cheaply if the U.S. computing and communications industries remained healthy. And as with DARPA's motivation, aspirations for resource sharing shaped NSF's investment. Much investment was therefore aimed at packet switching and the creation of an electronic communication network among researchers. The packet switching would enable the movement of files between supercomputer centers and many universities. Supercomputers were expensive fixed investments with no geographic mobility. The NSF aimed to use the Internet to permit many researchers to connect with those supercomputers, making greater use of the capacity and sharing the huge computing power

¹⁶ This company became known as Telenet, and grew into a very large commercial provider of packet-switching service. Eventually Telenet became part of the Sprint data network.

¹⁷ These issues are described in great detail in Abbate (1999).

they embodied.

Another aspiration for the NSF concerned a scaling issue that DARPA had not yet faced. The U.S. research community increasingly took to using the communication network for file sharing and electronic communication, and throughout this period traffic grew. The NSF aimed to build a routine and reliable network infrastructure, making it easy to adapt and spread to every place of higher learning in the United States – universities, community colleges, and research institutes.¹⁸ Over time, the investment aimed to give a wide range of participants – students, faculty, and administrators – a taste for what the Internet could do to help them in their work, namely, transmit electronic communication, data files, news, and other types of messages over long distances. The goal required a system that would handle traffic of many orders of magnitudes greater than anything accomplished to that point.

However, as with DARPA, NSF's management also came with some restrictions on participation, thus perpetuating the limitations of experimentation – only users connected to civilian research institutions could make use of the NSFNET, not, for example, commercial interests (except those that supplied services to NSF). However, restrictions due to this “acceptable use policy” were less binding than they had been with DARPA, and for a few years NSF's managerial control reduced many of the tensions over participation.¹⁹

The transfer to NSF had several more consequences. A new source of funding introduced a new budgetary process, a new outlook about the future, and new set of priorities for operations. In particular, NSF managed the backbone of the network, but gave discretion to many universities to modify their installations as they saw fit. NSF also differed from DARPA in its more relaxed approach to outsourcing equipment supplies, which had later consequences for transitioning from NSF administration into wide commercial use. By the time the Internet was commercialized, the surrounding industry was already in place to meet the needs of the new commercial market. For example, by the early 1990s there already was an industry building routers consistent with widely employed software protocols.

In the late 1980s, NSF presided over another seminal design choice—the switch to a routing protocol that allowed for more than one backbone.²⁰ Until the NSFNET came into existence, there was only one network and one backbone, and BBN operated it. The scale was limited, and NSF anticipated supporting a much large network. The NSFNET therefore introduced additional backbones and regional carriers. In due time, NSF worked with others to introduce routing protocols that no longer presumed NSF would be the sole manager for the backbone. This was the beginning of the technical design changes necessary for evolution to a commercial Internet with multiple commercial carriers.

By the beginning of the transition to commercialization in the late 1980s, the Internet was a large-scale and reliable data communications network with a well-documented code base upon which any participant could build additional layers of applications. While no serious networking engineer thought the Internet's technical capabilities had stopped evolving, insiders generally acknowledged that the research-oriented Internet had matured, moving beyond its “nuts and bolts” stage of development.²¹ At that point, Steve Wolff, then director of the NSFNET, recognized that there was no technical reason why the government had to solely operate the Internet. He also asserted that private firms could provide services as efficiently, or more so, than government-managed entities or sub-contractors. He therefore initiated a long series of steps, aimed at what would be a transfer of technology out of exclusive government management and use.

Wolff's decision in itself illustrates another extremely important lesson. When a technology reaches a point where private firms can commercialize it, the transfer does not necessarily happen on its

¹⁸ Frazier (1995), Leiner et al. (2003).

¹⁹ The NSF's “acceptable use” policy restricted the use of the NSFNET to any university research faculty, student, or institute that contributed to furthering the development of science in the United States.

²⁰ NSF switched from the routing protocol Exterior Gateway Protocol (EGP) and replaced it with Border Gate Protocol (BGP). The EGP protocol presumed a known pathway for connecting systems. BGP enables fully decentralized routing. To Internet veteran David Clark, making this change was one of the earliest technical signs of the pending arrival of commercial network and the retirement of NSFNET (Clark 2008).

²¹ This is the phrase used by Mandelbaum and Manderbaum (1993). See also Leiner et al. (2003).

own. It requires government managers who recognize this opportunity, and it may even require active nurturing from government officials, as it did in this case.

In the case of the Internet, this transition was quite early in some sense and quite late in another. By the time it was turned over to commercial use, the Internet had acquired most of the attributes that would lead to the transformation of every part of information and communications markets around the world. However, because of the NSF's "acceptable use" policy, there had been little experimentation with using the Internet for commerce. There also was little understanding about its cost structure outside of an academic environment. Few of the participants had incentives to fully explore how a wide range of inter-firm procedures would accommodate pricing, such as how interconnected networks would settle payments for exchanging traffic.

All in all, NSF's managers invested in numerous innovations that contributed to easing the transition to commercial markets. However, the limited experience with a variety of users undermined the ability of Wolff and his managers, as well as managers elsewhere, to forecast the appeal of new applications aimed at new commercial users.

The cost of innovation

The total cost to the government of creating the Internet is difficult to ascertain. It is known that during NSF's management (approximately 1985–1995) the agency invested \$200 million dollars in Internet technology.²² However, this figure does not include the DARPA funding that paid for most of the early invention in the 1970s and early 1980s. While DARPA's financial commitment to what became the Internet was undoubtedly considerable, to my knowledge no historian of these events has made a precise estimate of its size.²³ In addition, the cost tally of the Internet is further complicated because both DARPA and the NSF relied on distributed investments—the agencies paid for investments in backbone facilities and facilities for data-exchange, but offered only minimal support for investing in installations at universities. Most universities invested heavily in their own computing facilities, paid for by university funds.

The cost of the Internet would also include the substantial number of failures that were part of DARPA's broad portfolio of investments in computing science more generally. For example, it would include DARPA funding for a range of computer science efforts that did not work out as well as planned, such as in artificial intelligence. It also does not include a range of other experiments in computer science that NSF paid for and from which the general community of researchers learned.

The Internet also benefited from improvements in a wide range of computing equipment that would have occurred with or without government funding. Like any other IT-intensive activity, research on the Internet gained benefits from what was happening to all equipment based on advances in solid state circuitry. It was easier to make innovative gains when many of the other complementary inputs into the effort improved at the same time.

The cost of the Internet was not of interest to the government. DARPA quite explicitly did not use economic rationales to fund projects. DARPA funded high risk projects that "dealt with the development of fundamental or enabling technologies with which to reach newly defined DOD objectives....When DARPA judged success, it applied technical rather economic measurement standards."²⁴ Likewise, the NSF invested in developing Internet technologies to meet its agency mission, not with the intent of producing large economic gains.

In summary, the early Internet succeeded because of the mix of managerial wisdom, pragmatism, and technical meritocracy of those involved, and because those players kept their efforts trained on

²² Leiner et al. (2003).

²³ The entire expenditure for the IPTO, the agency within DARPA that funded most of the Internet, did not exceed approximately \$500 million over its entire existence (1963–1986), and the funding for what became the Internet was but one of many IPTO projects. (Norberg et al. 1996).

²⁴ Norberg et al. (1996).

scientifically worthwhile projects. The federal institutions sustained those efforts over a long period of time, building a community of researchers invested in innovating and refining attributes of the network. It eventually accumulated many attributes that today we recognize as the Internet, and which today we recognize as valuable. None of this was easy, automatic, or necessarily inexpensive.

The Commercial Internet

Once commercialized, the Internet was unlike any commercial communications network that came before it. While it still could be described as a packet-switching network for moving data between computing clients (as had been envisioned from its inception), this description does not fully describe the early commercial form the Internet. It accumulated more capabilities and functions as a range of firms began to use pieces of it to enhance services provided to paying customers. Over time, “the Internet” became a label for not only the Internet but also for all the applications that accumulated around the Internet, used pieces of the Internet, and commercialized new functions for the Internet. Together they delivered an impressive array of services to a wide range of users.

Supply of the Internet did not simply create its own demand. Rather, after years of development, a few applications were built that provided compelling value for tens of millions of decision makers. The size of the Internet access economy in the United States gives a sense of how big demand for the Internet became, once it started to commercialize. For example, the revenue associated with providing Internet access is one of the largest categories of revenue out of the value chain²⁵ for Internet services, and it grew quite large in only nine years (Table1). By 2006, total revenues reached \$39 billion; that is extraordinary for a technology that had almost no commercial service providers prior to 1989.²⁶

Table 1. Revenue for access markets (\$ millions)

Year	1998	1999	2000	2001	2002	2003	2004	2005	2006
Dial-up	5499	8966	12345	13751	14093	14173	14081	12240	10983
DSL	n.a.	228	1245	2822	4316	6954	10240	12034	15066
Cable	138	274	903	2600	4117	7372	9435	11139	13156
Wireless						n.a.	668	1140	n.a.

Source: *Greenstein and McDevitt (2009)*.

Note: n.a. = Not available.

These revenue levels are important to stress, because access fees generated most of the revenue during the first decade of the commercial Internet. The typical household spent more than three-quarters of its online time at free or advertising-supported sites.²⁷ Although subscription-based services and advertising services started growing rapidly after 2003, the amount spent on access fees each year has far exceeded advertising revenue. For example, the \$39 billion in access revenue in 2006 compared with \$9.7

²⁵ A value chain is a set of interrelated activities that produce a final product for end users.

²⁶ The closest commercial precursor to the Internet existed in the bulletin board industry, which generated several hundred million dollars of revenue before the commercial Internet blossomed and replaced it. The sentence uses 1989 as a marker because this is the year of entry for the first carriers for Internet traffic, PSINET and UUNET.

²⁷ See Goldfarb (2004). This discussion follows norms at the U.S. Census, as expressed in the Annual Service Survey. Most households devoted most of their Internet budget to access fees (largely for services provided by wi-fi hot spots, or dial-up, broadband, and wireless carriers) as opposed to subscription fees for content (largely provided by services such as Lexus-Nexus, the New York Times archive, and Wall Street Journal archive). AOL sought to blur the distinction between access and content with a “Walled Garden” strategy and successfully did so for a few years with its dial-up service. Later it reduced the importance of its access fees, relying on advertising for most of its revenue. This distinction does not count electronic commerce revenue, namely, use of electronic channels to support purchase of a good (e.g., clothing) or what had been a non-digital good (e.g., music).

billion in Web Search Portal revenue²⁸ (which includes advertising) and \$12.8 billion in Internet Publishing and Broadcasting revenue, of which \$2.9 is advertising revenue.²⁹ Advertising revenue is now growing at a more rapid pace than subscription and access fees, and it may exceed access revenue soon, but not as of this writing.

Widely dispersed market decisions lie behind this revenue growth, shown by the diffusion of Internet access to U.S. households (Table 2).³⁰ Starting with fewer than 20 percent of households in 1997, the Internet diffused to more than 73.1 percent of households by 2006. Similar results obtain for the diffusion of the Internet to business.³¹

Table 2: Internet Access in U.S. Households

Year	1997	199	1999	200	200	200	200	200	200	2006
Households*	103.	104.	105.0	106.	107.	108.	109.	110.	111.	112.0
Internet Adopters*	19.1	27.2	35.5	44.0	53.8	56.7	59.5	66.0	73.3	81.8
Broadband*	n.a.	n.a.	0.9	3.2	9.6	13.0	18.5	27.5	41.1	47.0
Dial-up*	19.1	27.2	34.5	40.8	44.2	43.7	41.0	38.5	32.2	34.7
% adopters **	18.6	26.2	33.8	41.5	50.2	52.5	54.6	60.0	66.0	73.1

Source: *Greenstein and McDevitt 2009*.

n.a.: Not Available.

* Millions of Households

** Percent of total households

Straightforward economic factors determined these trends: dial-up became available first and diffused to more than half of U.S. households. Broadband emerged later as a higher quality and more expensive alternative, albeit one available in only a few places and from a limited set of providers. Over time, however, broadband became more reliable and more widely available, and as that happened, many households paid to upgrade their Internet service.

Initial wave of entry and exit

A closer examination of the historical record shows that this market arose in distinct waves of entry and exit—the first wave of entry occurred after NSF opened the Internet to commercial users, coupled with the invention of the World Wide Web and the creation of the World Wide Web Consortium, which came into creation in 1994–1995. The World Wide Web, invented by Tim Berners-Lee, is an Internet application that links documents together through the use of hypertext and viewed by a web browser. The commercial browser began to diffuse in 1995, enabling new functionality and new businesses built around this new technology. The Web quickly became the software platform for many creations thereafter, motivating further experimentation and magnifying the potential for value creation from the first wave of entrants.

The creation of the commercial browser caused a change in expectations about what was possible to accomplish on the newly privatized Internet. The browser began as an academic project, but even that was sufficient to demonstrate an entirely new range of applications affiliated with linking various pages, displaying multi-media, and supporting a whole new interface for human-computing interactions.

²⁸ See Table 3.0.1 Information Sector (NAICS 51) or 3.4.2 Web Search Portals (NAICS 518112) in U.S. Census Bureau (2007).

²⁹ See Table 3.3.5 Internet Publishing and Broadcasting (NAICS 516) in U.S. Census Bureau (2007). There is negligible advertising listed for Internet Service Providers (other than Cable, Telephone or Wireless carriers) in Table 3.4.1.

³⁰ These data sources are described in more detail in Greenstein and McDevitt (2009).

³¹ As measured by the CPS supplement and the Pew Survey of the Internet and American Life. See Forman et al. (2003a, 2003b).

Following creation of the browser, participants expected an explosion of commercial activity by established firms, venture capitalists, Wall Street analysts, and entrepreneurs. Indeed, an immense entrepreneurial response did occur, which extended across a broad array of activities and applications, media, travel, commercial transactions, communications, and so on. The wave was a market response to new opportunity. Many different market participants sought to figure out how to apply the new technology to improve services to users. The effect was for unrestricted and entrepreneurial markets to apply and reapply these technologies over and over again to a wide range of problems and new applications.

During the wave of entry, new knowledge and lessons were shared at low cost.³² Several distinct models emerged taking advantage of the demand for electronic services. One prominent model subsidized the delivery of text and other visual media with advertising. Many of the adherents to on-line news, entertainment, and other information-based commerce found this to be attractive. Another prominent model used the Internet for the delivery of a service, such as the creation of on-line retailers like Amazon and the addition of an online counterpart to other branded catalogue retailers. Other models included developing a subscription service, such as for the NYT crossword puzzle; organizing a place for buyer and seller to conduct a transaction, such as an auction, and charging a fee for the service; organizing a fee-based listing service, such as an on-line help-wanted listing; providing a fee-based matching service, such as for singles; and providing a location for aggregating information from users (e.g., blogs, recommendations, wikis), supported by advertising.³³

The first generation of browser and Internet Service Providers (ISPs) diffused extraordinarily rapidly. For example, the fraction of households on-line jumped from 18.6 percent in 1997 to 41.5 just three years later (Table 2). The fast uptake of several popular applications of the 1990s (e.g., Hotmail, ICQ, and Yahoo!) reinforced this rapid diffusion.

However, the late 1990s saw more entry than actual demand would support a few years later. And so a shake-out ensued, first affecting access providers in 2000 (popularly known as the “Telecom Meltdown”), and then eventually many on-line retailers (popularly known as the “dot-com crash”). Investment uncertainty after the events of September 11, 2001, magnified the downturn affiliated with this adjustment. Sellers with high debt and low revenue exited. This occurred at all levels of the value chain for Internet services, as well as infrastructure building.

While the mass exit led to widespread losses for many entrepreneurs and investors in entrepreneurs, it should not come as a surprise with the benefit of hindsight. Much of the activity was exploratory in nature, and, as is typical, some explorations fail while others succeed. Moreover, historians of technology had described investment booms and busts for other episodes of technological innovations and commercialization, such as followed the growth of railroads, the growth of the steel industry, the growth of the automobile industry, the growth of electrical power, and so on. Finally, computing markets also had experienced booms and busts during the development and deployment of the personal computer, the minicomputer, and client-server systems, albeit at a smaller scale than what followed the commercialization of the Internet.

The drama of the decline obscured another trend – how first wave of entry and exit changed the economic landscape. A large array of on-line activities survived, including large providers (e.g., AOL, Yahoo!, E-Bay, Google), as well as a wide array of niche products, services, and productivity enhancements. Many catalogue retailers successfully transitioned into on-line retailers, such as Victoria’s Secret, and LLBean, and thrived just fine with their existing brand names and efficient order fulfillment. In short, even with excess entry, markets have a way of rapidly creating thriving businesses that take advantage of the opportunities enabled by the new technology. The results are hard to foresee until supply, demand, and prices plays itself out in all its glorious unpredictability.

³² The primary cost to society were the “co-invention costs,” that is, the expenses incurred by suppliers and buyers in the pursuit of customizing the general purpose technology to the unique needs and idiosyncratic circumstances their market participants faced. These costs arose for users trying to apply the technology and suppliers trying to sell it. For a discussion of co-invention, see Bresnahan and Greenstein (1997).

³³ For a summary of the diversity of models, see Hanson (2008) and Hirsch and Goldfarb (2008).

Accumulating innovation in the Internet³⁴

Accumulation of innovation in a market setting differs substantially from that in a skunk works. In a market setting, there are common signs of healthy innovative behavior, even in a quickly evolving industry such as the Internet. Commercial behaviour resides inside a complex value chain. No single firm controls the value chain. The quality, price, and user experience arise from the interactions between participants in the value chain.

Even when there is no agreement about which criteria observers should use to assess the performance of the commercial Internet, there are patterns of healthy conduct, that is, commercial behavior indicative of an innovative industry. Such healthy behavior correlates with desirable market-wide outcomes, such as improvement in products, lower prices, new capabilities, or other innovations that lead to productivity improvements among business users.

Three general features of the market foster accumulation of innovation. These are economic experimentation, entrepreneurial initiative, and vigorous standards competition. *Economic experimentation* is a market-oriented action designed to help a firm learn or resolve uncertainty about an unknown economic factor. Usually such lessons cannot be learned in a laboratory or controlled environment, either because they involve learning about the nuances of market demand or learning about sets of procedures for providing new services at a lower cost.³⁵

Not all economic experiments come with the same orientation or learning goal. Some focus on learning about the profitability of incremental changes in business processes. Some seek to learn about the restructuring of organizations and the profitability that may result from the simultaneous alteration of many processes. Some even seek to learn about the profitability of restructuring the relationship among many organizations within an industry.

Internet markets have been full of economic experiments in the last 15 years. That was especially so in the latter part of the 1990s, when firms took a wide variety of bets to learn about unknown aspects of customer demand and the costs for meeting them using Web technologies, such as the browser, server software to support it, and a range of other innovations. These experiments covered all parts of the value chain for delivering services—Internet access, client-server platforms, contracting among business partners, and so on. Carriers conducted them and so did content providers.³⁶

Entrepreneurial initiatives involve an organization in a risky and challenging business in pursuit of a new economic opportunity. These firms are the market “participants” that make the first brave attempts at deploying, distributing, or servicing a new good to a wide range of customers. Small start-ups take entrepreneurial action and so do large firms. Sometimes small businesses that take such risks are bought by large organizations. Sometimes small start-ups go public and grow into large firms themselves. The increasing presence of entrepreneurs in communications markets has brought rapid change to many sub-markets.

Yet, entrepreneurial activity can increase and decrease for distinct reasons. Experimentation and competition between leads to innovative entrants, or it may enhance the products of one particular firm. It forces incumbents to react, or, even better yet, anticipate the entrant and innovate in advance. This fosters incentives to lower prices and sponsor more innovative products, and sooner. Users benefit from all of those.

Vigorous standards competition also played a role in innovation in the Internet. That is because leading-edge technologies often cannot deploy on a wide scale without some routines or processes, and/or coordination of activities across many firms. Thus, the ratification of new standards generally acts as a

³⁴ Much of this section and the next provide a synopsis of arguments in Greenstein (2009c).

³⁵ *Economic experiments* pertain to any market experience that alters knowledge about the market value of a good or service (Rosenberg 1994; Stern 2005). Firms engage in economic experiments to reduce uncertainties about market value.

³⁶ See Greenstein (2008a) for an examination of the role of economic experiments in the evolution of Internet access.

leading indicator of impending technological progress and serves as another sign of a healthy innovative industry. While new standards and upgrades to existing standards may not arrive at a regular rate, a slow pace for development or a slow arrival of new standards usually sets off alarms.

To be sure, this benchmark is particularly challenging to put into practice, because some standards are more important than others. The Internet Protocol Suite TCP/IP has played a central role for decades, for example, and any alteration to them receives considerable attention, deservedly more attention than other standards. The same is true for protocols that govern the World Wide Web, which are handled at the World Wide Web Consortium (W3C). This is also so for important components of the Internet, such as upgrades to wire-line Ethernet. That topic is discussed at the Institute of Electrical and Electronics Engineers (IEEE) Standards Association committees assigned to new standards. In the case of wire-line Ethernet, for example, it tends to be sub-committees of the Working Group for Wireless Local Area Network Standards.

Standards design needs competition. Although the process of standards design in which market competition has played a role can be a messy process and often frustrating, confusing, and messy, this mess is necessary. Standards designed in the absence of competition usually have been orderly, infrequent, and simplified. Such standards have been more likely to lead an industry down an unhealthy and innovative path as it can go.

If a firm with market power designs a new standard, it will face strong incentives to roll it out slowly to protect the firm from cannibalizing its own monopoly rents. For example, in the days when IBM controlled a large part of the mainframe market it could not bring itself to abandon Extended Binary Coded Decimal Interchange Code (EBCDIC), its standardized proprietary language, or, for that matter, to help others migrate up from EBCDIC to the many other superior languages available. Despite plenty of improvements IBM could have made, its managers refused to deploy them, preferring instead to exploit locked-in users.³⁷

Monopolies also face strong incentives to have a “quiet life,” to paraphrase Sir John Hicks.³⁸ That is, monopolies may exert less effort when they choose standards, or design them to castrate user choices in such a way that leads to less inconvenience for the monopolist at the expense of the user (e.g., trimming product line breadth, or trimming away complex attributes of the product). For example, until the mid-1970s, AT&T held a monopoly over residential customer telephone handsets. Most households faced a limited menu of (over-engineered and excessively rigid) choices. Well engineered or not, there were too few choices in comparison with what a competitive market would have done.

With the breakup of AT&T’s monopoly, multiple providers began to match the offerings of its nearest rivals. In a short time, the heated and urgent competitive behavior familiar to consumer electronics eventually overtook the market, leading to a plethora of choices at a wide range of prices. In other words, in the absence of restraining limitations on discretion, monopolies have designed selfish standards. An antidote to the selfish standards of monopolies has been competition between standards. In the history of the Internet, massive entrepreneurial entry drove innovation, and accessible standards contributed to it.

Negotiations between open and proprietary standards

One feature of the competitive Internet is probably the most crucial for accumulating innovation from dispersed market participants. Not surprisingly, it is the most controversial. *Negotiations over interdependent processes* shaped how the market accumulated services and built on each other. These negotiations took on importance because every participant, the innovative and not so innovative, operated within a system of technically interrelated components and services, where these processes interoperated. The failure or reduction in performance of any of these activities could lead to degradation of the quality

³⁷ Brock (1975).

³⁸ “...People in monopolistic positions...are likely to exploit their advantage much more by not bothering to get very near the position of maximum profit, than by straining themselves to get very close to it. The best of all *monopoly* profits is a *quiet life*.” Hicks (1935).

of outcomes for many users.

In a network with a high degree of technical interrelatedness, there are general gains to all parties from bringing routines into business processes and activities, much like there are gains to adopting standards and platforms to coordinate activities. While there may be no better way to reduce complexity, adopting such routines may require negotiation between multiple parties. For example, even the simplest of activities, such as sending e-mail, involves many participants, and efficient delivery of services depends on advanced agreement about how their business activities will interrelate.

To reduce the uncertainty about how such services interoperate, commercial firms take one of two approaches, either they negotiate arrangements (contractual norms) in advance with all relevant participants, or they do it all themselves by offering a platform (a bundle of standards) that accomplishes the same task, internalizing the contracts within one firm's decision making.

Although the inception of the early internet was a “network of networks,” today leading firms and their business partners view the commercial Internet as a “network of platforms.” This seemingly small change in definition is far from insignificant. The rise of a plethora of platforms on the Internet is a source of both celebration and consternation. Platforms perform functions that firms and/or users value. Their presence usually suggests that some firms/users are better off with them than without, and it usually suggests they have replaced an inadequate non-proprietary standard inherited from the era prior to commercialization. At the same time, large or dominant platform leaders usually possess market power, and that occasionally gives them the ability to resist non-proprietary standards that serve the interests of some rivals.

Which is better, proprietary or open? Such debates inevitably boil down to restrictions on the discretion of incumbent management to determine standards. Proprietary and open standards contrast most sharply in their respective approaches to transparency and participation.

With standard proprietary platforms, leading firms retained discretion and guided participation within strict rules. Generally, strong platform leaders, such as IBM, Microsoft, and Intel, retained their authority by owning assets on which others depended, and by not being transparent about how such assets would change in the near future. Such practices came into direct conflict with the transparent and participatory processes for standards development in the Internet, particularly as practiced at the Internet Engineering Task Force (IETF), which used these processes to support group decision making.³⁹

Conversely, transparent processes are those in which decision makers alert participants to imminent change—sometimes well in advance—when their change will diminish the returns on others’ innovative investments. In many Internet standards forums, such as the IETF, the organizations take considerable effort to remain transparent and embed such norms in the operations of the group.⁴⁰ Such transparency is one of the reasons why standards processes have become a strong indicator of the imminent release of leading-edge technologies in Internet equipment. Interested parties monitor the designs (because they can attend IETF meetings) and know that their near rivals do the same (because the data are available to anyone). All parties plan to match each other along the dimension of the standard and differentiate along the dimensions in which each has competitive advantage. Competition ensues once the standard is upgraded from its Beta to an endorsed and official standard.

Transparency is also a feature found frequently in open source projects with importance to the Internet value chain, such as Linux, Apache, Firefox, and the W3C. In the experience of Internet networks, a minimal level of transparency has been a necessary element of an open value chain operating for a large number of users. Defining such minimal levels is an important economic issue if other firms will not make long-term investments unless they understand at a fine level of detail how their software must interact with another’s.

Transparency is distinct from participatory rules.⁴¹ Participatory processes are those in which

³⁹ These trade-offs are discussed at length in Greenstein (2009b), which contrasts the operations of Microsoft, as it developed Windows 95, with the operations of the IETF, as it operated at the beginning of commercialization.

⁴⁰ Bradner (1999).

⁴¹ West and O’Mahoney (2008).

sponsoring organizations invite comment, discussion, and input from others affected by their actions. Such organizations solicit input through public forums, e-mail lists, blogs, community sites, and a range of other activities. Standards organizations vary considerably in their policies for encouraging or discouraging participation. Some organizations charge fees, some require participants to meet certain technical qualifications, and others allow any observer to attend but not vote. For example, the organization that designed and updates Wi-Fi does not allow unrestricted participation; firms must pay a fee in order to send representatives.

Wide participation is found quite frequently in open source projects, particularly those without sponsorship. Often technical skill determines participation. For example, the Firefox browser community has quite diverse participation from numerous corners, though participants tend to self-select on the basis of technical skill simply because they would be lost otherwise. Similar observations hold for Linux and Apache. In both examples, most participants are quite technically skilled, and in the latter case such skill acts as an explicit qualification. Wikipedia is perhaps the best-known example of an online project that encourages wide participation from a community of contributors, and where no skill test is applied to contributors.

Wide participation is probably the least common attribute among standards consortia sponsored by commercial firms. Most managers prefer to retain decision making authority, guarding investment decisions in the name of stockholders. There is concern that giving up such discretion risks having participants take investment in directions that do not serve firm interests.⁴² In addition, accommodating wide participation normally comes at a cost, such as slower decision making and more onerous managerial challenges in coming to consensus. Hence, even some ostensibly open standards processes chose to restrict participation. For example, Tim Berners-Lee established the W3C with a less participatory structure than found in the IETF, where he had personally experienced the drawbacks of slow decision making when he first tried to standardize the core inventions behind the World Wide Web.⁴³

Although the Internet experience does not give precise directions towards the best choice for participation rules, in the past, wider participation has tended to beat out no participation. Thus, every proprietary platform adopts some degree and form of participation (although, to be sure, with varying degrees of transparency). Perhaps the biggest surprise from the Internet experience is the persistence of standards-making institutions with wide participation and transparent processes. Once established, these institutions have persisted, co-existing alongside proprietary platforms, sometimes as competitors and sometimes as complements. Firms have learned to live with these institutions, and many firms have learned to thrive alongside them. Cisco, Intel, IBM and many Wi-Fi firms are active participants in these standards forums. Even (previously reluctant) firms such as Microsoft, AT&T, and Verizon have found it useful to participate and fund such activities.

Policy and governance

Direct government support for R&D in the creation of the Internet had two potential effects. It accelerated the arrival of the technology, and it influenced its direction. During the pre-commercial era of the Internet, the creation of a general purpose technology for exchanging packets of data between many firms was risky. It had no immediate obvious commercial payoff. DARPA's program officers intended to fund radical technological progress that otherwise would not have been funded by private firms. They intended to develop research communities in those areas where almost none had existed. The investment aimed to develop fundamental scientific understanding and engineering experience, accelerating the arrival of actual products and services at some point. DARPA accomplished its goals and then some. Long before the Internet arrived in particular, packet switching was but a theoretical idea and expensive to implement. There is no doubt the initial work funded by DARPA in the 1970s accelerated the arrival of the technology. No other private firm at the time (such as IBM or AT&T) had projects in the area coming

⁴² West and O'Mahoney (2008).

⁴³ This frustration is described in detail in Berners-Lee and Fischetti (1999).

close to DARPA's efforts in size and scope.⁴⁴

DARPA also funded the building of research prototypes and the building of a prototypical system. Arguably that went beyond the aspirations of the DOD, but it was an immediate byproduct of its success and a natural extension in terms of extending the scientific/engineering frontier. It had enormous value too, demonstrating the feasibility of what had been a theoretical idea.

In an industry where many potential future technologies vie for attention, there was value in demonstrating that one of these forecasts was viable. In this case, DARPA Net, and later NSFNET, showed that the successful operation of the entire system could have great value. This illustrates how demonstrations can serve as a focal point for further development, particularly in the face of widespread industry resistance prior to such demonstration.

However, government funding came with a drawback. Restrictions on participants and “acceptable use” truncated experimentation and entrepreneurial initiatives. This truncation was a detriment to understanding the potential for the network outside of the limited uses to which DARPA's community put it.

The government's role differed in the later time period, particularly the era just prior to the blossoming of the commercial Internet. In the private sector by the late 1980s, most savvy commercial observers anticipated the arrival of mass market electronic commerce, but few anticipated that the Internet would be it, and, for related reasons, few guessed that any would arrive as soon as did the Internet. Moreover, private firms largely ignored the Internet. Why? The skepticism at many corporations can be interpreted in three (overlapping) ways:

- a misunderstanding of the potential for the Internet, perhaps due to the commitment to an alternative technological vision or forecast;
- a situation in which the technology lacked internal “champions” inside leading organizations, perhaps because of the expectation that it would cannibalize too many revenue streams at existing business;
- a situation in which a technology benefited many users at once, perhaps because no single firm had incentive to nurture adoption that seemingly did not directly contribute to their own bottom line.

For example, neither AT&T's management nor that at any of the “baby Bells” expressed any strategic interest in commercializing services related to TCP/IP in the late 1980s. Most preferred to invest in services such as the Integrated Services Digital Network (ISDN),⁴⁵ which the managers considered the technical direction worth exploring.

To be fair to AT&T's managers, they were not alone. For example, despite employing some of the best researchers in the world on this topic and despite the involvement of its research division in the NSFNET, IBM's strategic planning at the corporate level also ignored the Internet. The company did not aggressively commercialize related services. Instead, its corporate plans called for commercializing a proprietary set of networking technologies, built around its Systems Network Architecture (SNA). These plans ultimately led to some of the most high-profile product development failures in IBM's history.⁴⁶ Likewise, Digital Equipment Corporation, then the second largest computer company in the world, was strongly committed to DECNet, a proprietary network service.

Even many sponsors of the NSFNET ignored its commercial potential. Only a few participants took entrepreneurial actions, such as the carriers PSINet and UUNet, which entered in 1989, far before many others. Only a few bulletin board providers, such as Prodigy and CompuServe, made early switches to the Internet.⁴⁷ For example, none of the other mid-level carriers made a switch to for-profit status until

⁴⁴ Abbate (1999).

⁴⁵ ISDN is a circuit-switched telephone network system that can carry voice as well as data over the same line.

⁴⁶ Gerstner (2002).

⁴⁷ Though, for a number of reasons, it largely did not help them gain market share, or thrive during the first wave of entry. This story is told particularly well by Banks (2008).

after the NSF commercialization neared completion in 1994—five years later. As it turned out, AT&T did not start to offer Internet service until 1995, which is about the same time as many other mainstream firms. First it offered service to business and then to homes a year later. It continued to do well with business users, as well as briefly with home users before it faded later in competition with AOL. The mid-level networks started to convert to for-profit status about the same time. Most of the baby bells were even later than that.

Shaping direction

The rise of the Internet shaped the direction of technical change. It contained two important attributes: it comprised largely or non-proprietary features/protocols/standards coupled with an open organization to support existing standards and to update them (e.g., the IETF and W3C). Its existence was unexpected, in part because its leadership structure differed from every other alternative considered plausible by contemporary executives – such as a network dominated by any established firm (e.g., IBM), carrier (e.g., AT&T), equipment manufacturer (e.g., Lucent), quasi-government agency (e.g., the ITU), or industry consortium (e.g., ISO).

For example, although many firms had e-mail services for their own computer networks, none of them had incentives to combine their systems with others'. No single firm had incentives to aggregate innovative suggestions from a vast array of contributors at the early and risky stage of developments. The Internet's non-proprietary features acted as an attribute around which many participants could agree because none of them individually risked too much nor benefitted too much.

Coupled with the open and non-proprietary nature of the Internet was a surprising set of technological leaders. Although government support made no difference to the stature of early innovators, whose reputations would have been high with or without DARPA's projects, that was not true for the first generation of Internet developers. They often expressed surprise at the discretion they had, many of them becoming leaders as graduate students. Their historical recollections refer to many moments when they wondered when they would be displaced by "a professional crew," that is, more senior researchers in the field of computer science.⁴⁸

Perhaps the biggest change in direction came from the structure of governance that came along with the Internet. The IETF and W3C, among others, were open processes, in the sense that they fully documented their activities, did not restrict participation, and never actively sought to exclude any innovator from building applications on to the installed base of accumulated protocols.

Not all aspects of that open structure were important.⁴⁹ For example, plenty of industry experience suggested that commercial organizations producing proprietary hardware and software designs could be as innovative as open communities. Sometimes open communities have been more innovative and sometimes proprietary firms have been, and both have co-existed. Rather, open institutions had two key structural features: not withholding information and not restricting its use. These features enabled the World Wide Web to commercialize so quickly. More pointedly, the IETF's leadership was unwilling to withhold information from anyone, effectively not excluding outsiders, such as Tim Berners-Lee at the time, even though his inventions potentially displaced so many established processes and existing technologies, even technologies supported by the IETF.⁵⁰ Once the World Wide Web began to diffuse, no established firm could stop others from building on it and bringing about a massive change in many aspects of economic activity.

⁴⁸ Crocker (1987).

⁴⁹ This argument is developed more fully in Greenstein (2009b).

⁵⁰ Even though the vast majority of participants inside the IETF viewed Tim Berners-Lee's proposal with indifference or hostility, all were perfectly willing to let him use all of the IETF's tools. Berners-Lee, thus, built on top of existing IETF approved protocols with full freedom and discretion. Thus, he was able to take action quickly. See Berners-Lee and Fischetti (1999).

Governance of the rules of the game

Experience in this industry highlights the importance of good governance – simply spending federal R&D money or adopting policies from a check list, by itself, would not have been sufficient to achieve success. Rather, successful public support for innovation has been embedded in an institutional structure that provided checks and balances, counterbalancing the risk of any effort from degenerating into pork barrel spending and coddling of existing incumbents. Creating this kind of system has required time, judgment, and (sometimes) strong political will.

During the Internet's pre-commercial era, many issues required sound judgments by public servants who were focused on executing a vision of what they thought would benefit the technological development of the Internet. Indeed, a crucial feature of DARPA's success resided in stating a clear mission for its effort.⁵¹ Another involved choosing managers with extraordinary intelligence and competence, giving them funds and discretion, and allowing them to work with minimal oversight.⁵² Managers played a crucial role at both DARPA and NSF, but they did not act alone. They had support from their direct supervisors and their co-workers, all of whom could articulate their general mission and understand how that translated into short purposeful managerial action. The pre-commercial era of the Internet also received political support from those in the defense department committed to DARPA's autonomy.⁵³ Political actors did not intervene in the research involved in the Internet, although the Mansfield Amendment did influence a number of other related projects funded by DARPA. Political management also supported NSF's stewardship and beyond (e.g., support from Senator Al Gore and Congressman Rick Boucher).⁵⁴

During the commercial era, government played a role in setting the “rules of the game” by shaping negotiations among participants. In particular, legal questions covering intellectual property, monopoly powers, and other limitations and protections have shaped the Internet landscape. Market actors are sensitive to persistent and unresolved legal uncertainty over liability, ownership, and other legal rules that shape returns on investment. Hence, crucial parts of the value chain for the Internet have stalled as participants awaited legal or regulatory rulings settling boundaries.⁵⁵ Recently, for example, YouTube was founded in an era when there were multiple plausible definitions for the precise legal safe-harbor for including copy-righted material on a web site for user-supplied video. These definitions today still remain ambiguous. Google acquired You-Tube in spite of the shadow of the legal risk, and its investments (worth hundreds of millions of dollars) will most likely change as court decisions change.

After the retirement of the NSFNET and during the massive investments in the commercial Internet, it was fashionable to claim that the government's role was minimal in fostering innovative incentives. Such a claim is fatuous at best. For example, government actors were involved in determining rules for and resolving disputes about the minimal technical requirements telephone companies had to follow when interconnecting with a dial-up ISPs. Government actors also were involved when WorldCom sought to merge with MCI, thus aspiring to assemble a large fraction of the Internet's backbone under one organization. And, perhaps better known, government actors were involved in a wide array of issues that arose around Microsoft's behavior during the browser wars.

Such debates quickly reached back into the institutions that governed standards for the Internet. For example, legal precedents were set at a wide array of government organizations with jurisdiction over

⁵¹ Licklider's three criteria for funding research still sound prescient today: “1. The research must be excellent research as evaluated from a scientific or technical point of view; 2. The research must offer a good prospect of solving problems that are of interest to the Department of Defense; 3. The various sponsored efforts must fit together into one or more coherent programs that will provide a mechanism, not only for execution of the research, but also for bringing to bear upon the operations in the Defense Department the applicable results of the research *and* knowledge and methods that have been developed in the fields in which the research is carried out” Norberg et al. (1996). See also Waldrop (2001) for a wider discussion.

⁵² Norberg et al. (1996).

⁵³ Norberg et al (1996).

⁵⁴ Wiggins (2000) provides an overview of Al Gore's role in securing funding for NSF, and Segaller (1998) partially recounts Boucher's role in opening the Internet to commercial use. For the latter see also Shah and Kesan (2001).

⁵⁵ See the analysis of such matters in the area of communications carriers in, for example, Goldstein (2005).

these disputes, such as at the Department of Justice, Federal Trade Commission, Federal Communications Commission in the United States, and at equivalent European Union Regulatory bodies.

The policies that have resulted from legal battles have been a source of regulatory tension and friction. For example, established regulations, known as Computer II, compelled the U.S. phone industry to accommodate the new ISP industry. Managers in the existing telephone firms did not want to accommodate dial-up ISPs, but did so at first because Computer II required it of them.

But policies such as Computer II were not there to support the Internet. Rather, they were the outgrowth of two long-standing principles: (1) common carrier regulation for telephones, which prevented the telephone company from being selective about who they served; and (2) antitrust regulations, which had led to the divestiture of AT&T, and, more importantly, to a series of regulations for governing carrier interactions with others, such as equipment firms, and providers of services over telephone lines, such as Bulletin Board providers.⁵⁶ Because of these legal actions, the United States had a less hostile approach to entrepreneurial entry of dial-up ISPs than did most of the world for reasons unrelated to the Internet in particular.

The United States also had a very nurturing legal regime for consortia and standard-setting bodies. At crucial moments these policies fostered a healthy dose of vigorous standards competition. Once again, these policies existed for their own reasons, and not because any policymaker was trying to encourage the Internet in particular. For example, such laws played a crucial role in Tim Berners-Lee's personal decision making. In the United States, he received a much more welcoming set of conditions than those he faced in Switzerland, motivating him to leave employment at CERN and establish the W3C at MIT.

Finally, Why the Internet Worked

The history of the Internet highlights two distinct ways of organizing a long term program for accumulating innovation in a complex interdependent system. One approach relies on autonomous research institutions (skunk works) to organize and nurture inventive employees (wild ducks). The other approach relies on commercial markets to aggregate dispersed initiatives from a wide array of entrepreneurial participants.

A skunk works faces a significant danger of innovating into areas where there is no demand, and therefore no economic value. How did the pre-commercial Internet create value in spite of the absence of demand? First, it avoided some dangers by keeping prototype and operations sufficiently close to one another. The first participants in the non-commercial Internet assessed value from their own experiences, and DARPA managers nurtured and permitted experimentation to blossom. That helped them create useful and innovative applications such as e-mail and packet switching. The DARPA skunk works worked within community norms that fostered accumulating technologies on the merits, avoiding technical dead ends. In addition, DARPA and the NSF played a pivotal role in becoming the element of "demand" for which innovation was supplied. The agencies' substantial funding to research institutions (as well as leveraged funds through distributed investments to universities) procured the innovation for what became the breakthrough technologies leading to today's Internet. This investment was not easy, automatic, or inexpensive, but many would argue it has been one of the most important innovative undertakings supported by the U.S. government.

On the other hand, the skunk works approach restricted participation and truncated experimentation by excluding innovation along lines that did not support the "acceptable use" requirements of the government agencies. Such restrictions limited learning to an artificially narrow range of issues and left a wide array of other applications untouched.

In contrast, the commercial era of the Internet played to the strength of market-based innovation. It permitted decentralized exploration from commercial firms facing a wide array of incentives and a wide

⁵⁶ For a summary of the consequences of these rules for dial-up ISPs in the United States, see Goldstein (2005), or Greenstein (2008b).

variety of idiosyncratic circumstances. Market-oriented exploration did a marvelous job of exploring the range of uncertain factors affiliated with satisfying demand, thus demonstrating the benefits conducting many economic experiments. Once released to commercial interests, the Internet became the springboard for a dizzying array of applications that were not envisioned by the sponsoring government agencies. These applications, particularly the World Wide Web and its associated browsing technology, quickly infiltrated nearly every aspect of U.S. business and domestic life, and their effect continues to grow.

However, these explorations came to fruition because they were built upon a backbone technology that no single player or group of players in the market was willing or able to undertake—or for that matter, were forward-thinking enough to even visualize it. Throughout the history of the Internet, standards, protocols, and other rules of governance have shaped the direction and rate of innovations emerging from it. Some of these guiding factors grew with the project, most notably standards such as TCP/IP and protocols that govern the World Wide Web, as well as standard-setting bodies such as the W3C and IEEE. Other influential forces were not specific to the Internet but shaped it markedly, such as Computer II and legal rulings against monopoly control over communications technologies.

Perhaps because the DARPA skunk works invested heavily in many different directions (and many of them ultimately not bearing fruit), and used the brain power of many researchers at an array of institutions, they themselves garnered the power of decentralization (which is usually affiliated with the marketplace), albeit in narrower more disciplined form. Similarly, because the commercial Internet relied on strict protocols and standards-setting bodies, they in effect demonstrated the discipline seen more typically in centrally funded efforts.

While the government-based approach to innovation and the market-oriented approach each have their strengths and challenges, in the case of the Internet, these two systems came together in a unique, phased, and ultimately complementary way.

References

Abbate, Janet (1999), *Inventing the Internet*, MIT Press; Cambridge, MA.

Allen, Robert C. (1983), "Collective Invention," *Journal of Economic Behavior and Organizations*, 4(1), pp. 1-24.

Banks, Michael A., (2008), *On the Way to the Web: The Secret History of the Internet and its Founders*, Apress; Berkeley, CA.

Berners-Lee, Tim, with Mark Fischetti (1999), *Weaving the Web, The Original Design and Ultimate Destiny of the World Wide Web by Its Inventor*, Harper Collins; New York.

Bradner, Scott (1999), "The Internet Engineering Task Force," in *Open Sources: Voices from the Open Source Revolution*, eds. C. DiBona, S. Ockman, and M. Stone, Sebastopol, CA; O'Reilly Media Inc.

Bresnahan, Timothy, and Shane Greenstein (1997), "Technical Progress and Co-Invention in Computing and in the Use of Computers." *Brookings Papers on Economics Activity: Microeconomics*, Pp. 1-78.

Brock, Gerald, 1975, "Competition, Standards and Self-Regulation in the Computer Industry," in *Regulating the Product: Quality and Variety*. Richard Caves and Marc Roberts (eds), Ballinger Publishing Company; Cambridge, MA.

Clark, David (2008). Personal communication between David Clark and the author, September 26.

Crocker, David (2008a), "A Personal View: The Impact of email work done at the RAND in the mid-1970s," Mimeo, <http://www.bbiw.net/articles/rand-email.pdf> (accessed September 8, 2009).

Crocker, David (2008b). Personal communication between David Crocker and the author, August 7.

Crocker, Steven D. (1987) "The Origins of RFCs," in RFC 1000 - Request For Comments Reference Guide, compiled by J. Reynolds and J. Postel, ISI, August, <http://rfc.sunsite.dk/rfc1000.html> (accessed August 5, 2009).

Forman, Chris, Avi Goldfarb, and Shane Greenstein 2003a “The geographic dispersion of commercial Internet use.” In *Rethinking Rights and Regulations: Institutional Responses to New Communication Technologies*, eds. Steve Wildman and Lorrie Cranor, 113-45. Cambridge, MA: MIT Press.

———. 2003b, “Which industries use the Internet?” In *Organizing the New Industrial Economy*, ed. Michael Baye, 47-72. Amsterdam: Elsevier.

Frazier, Karen (1995), *Building the NSFNet: A Partnership in High Speed Networking*, <ftp://nic.merit.edu/nsfnet/final.report/.index.html>, downloaded, August, 2009.

Gerstner, Louis, V. Jr., 2002, *Who Says Elephants Can't Dance? Inside IBM's Historic Turnaround*. HarperBusiness; New York.

Goldfarb, Avi (2004), “Concentration in Advertising Supported On-line Markets: An Empirical Approach,” *Economics of Innovation and New Technology*, 13 (6), 581-594.

Goldstein, Fred (2005), *The Great Telecom Meltdown*, Artech House; Boston.

Greenstein, Shane (2008a), “Economic Experiments and Industry Know-how in Internet Access Markets,” in (eds) Adam Jaffe, Josh Lerner and Scott Stern, *Innovation, Policy and the Economy, Volume 8*, MIT Press.

Greenstein, Shane (2008b), “Innovation and the Evolution of Market Structure for Internet Acces in the United States,” in in (eds) William Aspray and Paul E. Ceruzzi (2007), *The Internet and American Business*, MIT Press; Cambridge, MA. Pp 47 – 104.

Greenstein, Shane (2009a), “The Emergence of the Internet: Collective Invention and Wild Ducks,” *Industrial and Corporate Change*.

Greenstein, Shane (2009b), “Open platform Development and the Commercial Internet.” In (Ed) Annabelle Gawer, *Platforms, Innovation and Competition*, Edward Elgar.

Greenstein, Shane (2009c), “Glimmers and Signs of Innovative Health in the Commercial Internet,” *Colorado Law Review of High Technology*.

Greenstein, Shane, and Ryan McDevitt (2009), “The Broadband Bonus: Accounting for Broadband’s Impact on US GDP,” NBER Working paper 14758

Hanson, Ward (2008), “Discovering a Role Online: Brick-and-Mortar Retailers and the Internet,” in (eds) William Aspray and Paul E. Ceruzzi (2007), *The Internet and American Business*, MIT Press; Cambridge, MA. Pp. 233 – 258

Hicks, John. 1935. “Annual Survey of Economic Theory. The Theory of Monopoly,” *Econometrica*, V3 (1), pp. 1-20.

Hirsch, David, and Brent Goldfarb (2008), “Small Idea, Big Ideas, Bad Ideas, Good Ideas: Get Big Fast and Dot Com Venture Creation,” in (eds) William Aspray and Paul E. Ceruzzi (2007), *The Internet and American Business*, MIT Press; Cambridge, MA. Pp. 233 – 258

Leiner, Barry, Vinton Cerf, David Clark, Robert Kahn, Leonard Kleinrock, Daniel Lynch, Jon Postel, Larry Roberts, Stephen Wolff, (2003), *A Brief History of the Internet, Version 3.32*, Last revised, 10 December, 2003. <http://www.isoc.org/internet/history/brief.shtml> (accessed September 8, 2009).

Living Internet. 1996–2009. “Electronic Mail (Email),” <http://www.livinginternet.com/e/e.htm> (accessed September 8, 2009).

Mandelbaum, R., and P.A. Mandelbaum (1992), “The Strategic Future of Mid-Level Networks,” in Kahin, Brian, *Building Information Infrastructure, Issues in the Development of the National Research and Education Network*, McGraw-Hill Primis, Cambridge, MA.

Maney, Kevin (2003), *Thomas Watson Sr. and the Making of IBM*, John Wiley and sons; Hoboken, NJ.

Meyer, Peter (2003), “Episodes of Collective Invention,” US Bureau of Labor Statistics Working Paper

No. 368. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=466880 (accessed September 8, 2009).

Norberg, Arthur, Judy O'Neill, and Kerry Freedman (1996), *Transforming Computer Technology, Information Processing for the Pentagon, 1962-1986*, Johns Hopkins University Press; Baltimore, MD.

Partridge, Craig (2008), "The Technical Development of Internet Email," *Annals of the History of the Computing*.

Rich, Ben R, and Leo Janus (1994), *Skunk Works; A Personal Memoir of My Years at Lockheed*, Back Bay Books; Boston, MA.

Roland, Alex, and Philip Shiman, (2002), *Strategic Computing: DARPA and the Quest for Machine Intelligence 1983 – 1993*. MIT Press; Cambridge, MA.

Rosenberg, Nathan (1977), "The Direction of Technological Change: Inducement Mechanisms and Focusing Devices," In Nathan Rosenberg, *Perspectives on Technology*, Cambridge Press.

Rosenberg, Nathan (1994), "Economic Experiments," in *Inside the Black Box*, Cambridge University Press, Cambridge.

Segaller, Stephen (1998) *Nerds: A Brief History of the Internet*, TV Books LLC; New York.

Shah, Rajiv C., and Jay P. Kesan, (2001), "Fool us Once, Shame on you – Fool us Twice, Shame on us: What we Can Learn from the Privatizations of the Internet Backbone Network and the Domain Name System," *Washington University Law Quarterly*, 79, pp 89 – 220.

Stern, Scott (2005), "Economic Experiments: The Role of Entrepreneurship in Economic Prosperity," in *Understanding Entrepreneurship: A Research and Policy Report*. Kaufman Foundation.

U.S. Census Bureau. 2007. "2007 Service Annual Survey: NAICS 51, Information Sector Services." http://www.census.gov/svsd/www/services/sas/sas_data/sas51.htm (accessed September 8, 2009).

Waldrop, Mitchell (2001), *The Dream Machine: J.C.R. Licklider and the Revolution that Made Computing Personal*, Penguin Books; New York.

West, Joel, and Siobhan O'Mahoney, (2008), "The role of Participation Architecture in Growing Open Source Communities," *Industry and Innovation*, 15 (2), April, pp. 145-168.

Wiggins, Richard (2000), "Al Gore and the Creation of the Internet," *First Monday*, <http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/799/708> (accessed September 8, 2009).