Toward Next-Generation Networks: Consolidation, Integration & Network Control

François Bar
Stanford University

Issues of Risk and Responsibility in Contemporary Engineering and Science:
French and US Perspectives
France-Stanford Centre for Interdisciplinary Studies (FSCIS)
Inaugural Conference – Stanford April 7 & 8, 2003

During the late 1990's "internet boom", following the passage of the 1996 Telecommunications Act, it seemed that nothing could stand in the way of rapid development of a radically new telecommunications infrastructure. Doubts were occasionally raised about the wisdom of particular business plans for internet applications, but everyone accepted the basic premise that bandwidth-hungry applications would proliferate and require ever faster broadband networks. While some expressed reservations about the long-term survival chances of many dot-com plays, few questioned the future economic health of the 'plumbers' of the internet –whatever combination of e-businesses would survive, they would create massive business opportunities for those making and deploying the networks' switching gear, routers, and fiber optics.

The 2000 financial collapse the internet and telecommunications sectors forced a profound re-assessment of the previous years' euphoria and showed that the boom's "sure things" entailed substantial risks. The mighty telecom giants, it turned out, had feet of clay. The telecom bust that started late that year affected not only recent entries like 360 Networks and Global Crossings, but long-standing industry players such as AT&T, Nortel and MCI/Worldcom. Yet, despite the boom and bust cycle, information

[*] A previous version of this paper was prepared for a briefing on "The Internet: Boom, Bust and Beyond", University of Maryland, October 2002.

networking has already profoundly transformed the U.S. economy and is here to stay. The profound reorganization in business processes and consumer behavior made possible by new networks will endure. As we consider policies and business strategies for the next generation networks, it is therefore critical to sort through the rubble of the last few years and assess what worked and why, to examine the risks entailed in early predictions and assess the various actors' responsibilities.

This briefing paper begins with a description of the 'telecom bust' that begun in 2000 and explores some of the competing explanations for the bust. It then reviews the factors that account for the success of the internet in the 1990s (the heydays of its evolution to date), seeking generic lessons for network infrastructure development. As a baseline against which to assess particular features of the internet's environment, that section contrasts the internet with the case of second- and third-generation wireless networks. Three dimensions of the discussion emerge as critically relevant to network evolution: infrastructure, architecture and users. Each of these highlights particular risks that weighed on the sector, and attendant responsibilities. The last section of the paper then examines each of these three dimensions in greater detail, to assess their likely evolution in the transition toward next-generation networks. It then concludes by highlighting the three fundamental challenges that face business and policy players as they embark on the next round of information infrastructure deployment and use.

The Telecom Bust

2000 certainly was a bad year for the networking industries. It is however important to differentiate two very different events that took place during that year. The first was the dot-com debacle in March-April 2000, the second was the telecommunications industry collapse that began six months later, in the late Fall of 2000. The dot-com bust was largely the belated recognition that many internet ventures had been funded without sound business plans and embraced by speculators rushing into a new market that they did not understand. People realized that it didn't make much sense to buy furniture or pet food over the internet and have it delivered through the mail, that losing money by providing free information was unlikely to lead to long-term success. The ensuing downturn was obviously a major loss for these investors, but was not a major setback for the economy as a whole. These startups employed only a few people, many of the products they made were not critical for the rest of the economy. While the NASDAQ declined from its high of over 5000 in the first quarter to just above 3000 within a few weeks, it went on to recover through the summer and the downturn did not spread broadly throughout the economy.

In contrast, the telecom collapse that begun later that year, and continues today, was a much more serious problem for the economy. Early failures by some of the Competitive Local Exchange Carriers (CLECs) and new network players were quickly followed by sweeping problems in all sectors of the telecom industry, from equipment makers Lucent, Nortel, even Cisco, to network operators like AT&T, MCI and Sprint. The downdraft then engulfed the entire Information Technology sector. By contrast with the Spring's dot-com bust, this affected a much broader swath of the economy. The telecom industry employs more people than the dot-com sector and constitutes a large part of the economy. Further, telecom infrastructure was critical for the workings of many other kinds of businesses in the economy.

What did account for the collapse of the telecom sector? Some observers point to lack of demand for high-bandwidth applications. They cite in particular the slower than expected pick-up of broadband connections, DSL and cable modems, which explained the early collapse of the CLECs. While it is true that broadband uptake hasn't matched the most optimistic expectations, market growth in that area has nonetheless been quite healthy, with the number of broadband connections growing by 157% in 2000 and 81% in 2001 (FCC 2002). That certainly compares favorably with the growth of many past communication technologies.

Rather than 'lack of demand', the problem is perhaps better characterized as unrealistic expectations. In fact, Andrew Odlyzko (Odlyzko 2000) first identified a pervasive myth about internet growth: throughout the 1990s, policymakers and business strategists repeatedly claimed that “internet traffic was doubling every three months.”(Odlyzko 2000) This may have been true in the very early days of the internet, but as Odlyzko shows, during the 1990s traffic has doubled “only” every year. While this annual doubling constitutes spectacular growth, it still falls quite short of a quarterly doubling. The expectations derived from this flawed but widespread estimate certainly fueled network deployment, and account for part of the current over-supply. Rapid progress in the underlying technologies used to build telecommunications networks accentuated the problem of over-supply. Most notably, the advent of wave division multiplexing (WDM), boosted the capacity of fiber optics by a factor of 100, even as more new fiber miles were laid out.

At its root, the telecom bust is first and foremost a result of over supply. Too many market players, incumbents and newcomers, built too many new networks, with too much capacity. In addition to over-optimistic demand forecasts and technological progress, this over-supply was further fueled by the economics of fiber supply: The fixed costs of acquiring rights of way and digging trenches far outweigh the marginal costs of burying one extra strand of dark fiber. As a result, the various carriers borrowed heavily to finance the construction of long-haul fiber networks, banking on skyrocketing demand to generate the income flows that would allow them to pay back those loans. When actual demand failed to match these expectations, prices collapsed and they were unable to meet their financial commitments. They were forced to pursue various avenues to restructure this debt, including bankruptcy.

Andrew Odlyzko provides a succinct summary of the situation: "Technology successfully met the challenge posed by unrealistic business plans that were formulated in willful disregard of real demand" (Odlyzko 2002). That situation has two important consequences for the next stage. First the business models of all existing telecommunications carriers are under sever stress and are in need of serious re-evaluation. Second, the transition toward next generation networks will require some mechanism to absorb the extra capacity built into the current networks. The approaches adopted to deal with these two issues will play an important role in shaping the next regime.

1990's Internet vs. today's wireless

Much of the current policy debate centers on the creation of the right conditions and environment for successful deployment of next generation information networks. In preparation, it is worth reviewing the factors that played a role in the success of the internet. In this section, we contrast the Internet's situation with the case of the wireless communication industry, which remains organized much like the traditional telecommunications sector. Three categories of factors are especially relevant: industry structure and regulation, standards and network architecture, applications and lead users. In all three, the internet presents a sharp contrast with wireless networks.

Industry structure and regulation

The internet was created as a software overlay atop existing telecommunication networks. Internet service providers therefore did not have to build specific facilities and were able to secure cheap bandwidth. Combined with requirements that local exchange carriers allow them open access to local loops without access charges, this meant that Internet Service Providers faced low entry barriers. Explicit regulatory conditions imposed on the incumbent telecommunications carriers, in particular open access requirements, fostered the growth of a vibrant internet industry. As a result, throughout the 1990s, the internet service provision industry has seen many players, engaged in dynamic competition against one another through service quality, including network availability, services and geographic coverage.

By contrast, wireless network operators face considerable entry barriers. They need to obtain spectrum licenses and invest in the construction of their own network infrastructure. Regulation of the wireless industry through the allocation of spectrum to specific wireless applications, and the licensing of a limited number of competitors for each of these applications, has led to a much more concentrated industry structure. Furthermore, the absence of open access requirements for the wireless infrastructure effectively blocks the unfolding of an evolution trajectory similar to that of the internet, where service providers could compete through the use of existing wireless infrastructure. As a result, the industrial structure of the current wireless service provision industry is highly vertical, with single companies controlling the license, the deployment of a wireless infrastructure and all the services provided upon it (Noam 2001).

Standards and network architecture

The internet owed a further share of its success to a unique approach to the development and governance of standards. A bottom-up, democratic process led to the elaboration of standards emerging from working prototypes, exposed to comments and refinements by other members of the internet community (Abbate 1999). Internet standards have been 'open' in the sense that their specifications are published widely, on a timely basis as they evolve, and are freely available to all. Providers of internet services, applications or equipment then compete on the basis of their implementation of these open standards.

This approach to standardization constitutes an important basis for the internet's overall design, the so-called end-to-end architecture (Saltzer, Reed et al. 1984). According to this principle, applications should be pushed 'up and out'. It is best for the network to be as transparent as possible so applications can be pushed 'out' and reside at the ends of the network, where they can be modified directly by their users, without requiring changes in the network itself. When features must be implemented within the network, that should be done as far 'up' as possible in the network's functional layers, so that change in these features is again as independent as possible from the network's hardware. The end-to-end architecture has served the internet extremely well, enabling multiple sources of innovation, in particular by the network's users (Lessig 2001).

By contrast, wireless standards have traditionally evolved in a much more confined environment. Standards tend to be proprietary and participation in the standard setting process is highly restricted (Noam 2001). Unlike the internet world, competition in the wireless world is more likely to center on the standards themselves than on competing implementation of common standards. The architecture of wireless networks result from and further accentuates its standards' features. Transmission services and communications applications tend to be deeply imbedded in the network itself. While on the internet any end-user can think up a new service, implement it on his computer and share it with the world (think for example of peer-to-peer music sharing software), the current architecture of most wireless networks makes such evolution dynamic impossible. The fact that internet users have essentially relied on one common appliance to access the network (the personal computer) made it easier to support such sharing. By contrast, wireless networks provide access to a wide variety of devices (cell phones built to many different standards, personal digital assistants) running diverse, incompatible operating systems.

Applications and lead-users

Partly because of the flexibility that came with its end-to-end architecture, the internet has been able to support extremely dynamic user-driven innovation. Most of the internet applications have been invented and developed by users, not by network providers. This contrasts sharply with traditional telecommunications networks, where only the network provider (e.g. ATT, through its research arm at the Bell Laboratories) could conceive, develop and deploy new services and applications, leaving the user with little influence beyond the choice to buy or not new devices and services. On the internet, end-users have been able to create the services that respond to their needs, or adapt applications shared by previous inventors over the network. The result is the tremendously dynamic innovation process we witness on the internet (Werbach 1997).

This dynamic is further reinforced by the fact that the internet has expanded to include progressively broader circles of users: the military-sponsored researchers who used the initial Arpanet were soon joined by colleagues working on various non-military NSF research projects, then by much broader education and research communities, then by commercial users and finally by the crowds of residential users. The internet has proven flexible enough to thrive through these successive extensions (Hart 1992). By contrast, innovation on wireless networks remains largely dominated by the main providers of equipment and networks. Because the industry is tightly controlled by a limited number of players, because the wireless networks standards are closed, because the networks' architecture leads to embed services in the network's core, end users have had only a very limited role in shaping the evolution of the wireless networks and their applications.

Obviously in these three dimensions, the internet's success rests on characteristics starkly different from those of wireless networks. This is not to say that one approach is necessarily always superior to the other, but to point out the different evolutionary dynamics set out by the two different environments. The next section explores current trends in each of these three areas. In each, the factors that accounted for the initial success of the internet now come under pressure, creating three fundamental challenges for policy as we embark on the deployment of next generation networks. These challenges present decision-makers with an important set of decisions. One approach would be to shape policy so as to extend the dynamics that fostered internet success. On the other hand, it is also possible that the forces that fueled the internet's successful growth have run their course, and that we ought to explore possible new forces to carry us toward the next generation networks.

Challenges for next generation networks

Who (what) will finance the next Infrastructure?

Internet growth in the 1990s piggy-backed on the telephone infrastructure. To be sure, portions of the long-haul networks were added specifically to transport internet data and switching devices were retooled to accommodate TCP/IP traffic. Overall however, the internet continues to ride on the last generation's network. This infrastructure will not be able to support the transition to the next generation information networks. First, as data traffic overtakes voice traffic, it certainly makes sense to design the transport infrastructure specifically for the carriage of data. Second, the business model under which traditional telecommunication infrastructure has been provided is under serious attack, as evidenced by the recent bankruptcies in long distance and CLECs industries

The ‘telecom bust’ hasn’t touched all telecommunication networks equally. Most seriously touched have been the long-haul carriers, both traditional players like MCI/Worldcom, ATT, Sprint, and the new players like Global Crossing, 360 Networks or Level-3. Another category of victims has been the competitive local exchange carriers (CLECs). By contrast, the dominant local network players, the incumbents Bell Operating companies, have fared relatively well so far. Even with the opening of the local market after the 1996 Telecommunications Act, they have largely retained their hold on the local market. While they have been forced to open their networks and make them available to competitors (the so-called service based competition), they have managed not to suffer too much under this new regime.

Are the local networks inherently sounder, or will they prove to be equally vulnerable once competition takes hold? We should remember that it took over ten years after the widespread introduction of long-distance competition in 1984 for ATT’s share of the long distance market to fall below 50%. Competition however has proved ruinous to the long distance participants, because of the underlying economics of the network industry: enormous fixed costs attached to the construction of networks, combined with near zero marginal costs of serving one additional customer, inevitably lead to mutually destructive price wars. Perhaps the same transition will unfold at the local level, undermining the position of the incumbents. Three alternative scenarios present themselves for the future of the underlying infrastructure.

The first, and perhaps most likely based on current developments, is *re-verticalization*. The local network market has already been substantially consolidated since the ATT divestiture, with two large players remaining (Verizon and SBC) and two smaller companies (Qwest and SBC). Qwest has been vertically integrated since its acquisition of US West, one of the Baby Bells, in 2000. The three other ILECs have recently begun to offer long distance service in a growing list of states and, under this scenario, would merge with or acquire one of the leading long-distance carriers. The resulting industry structure would be very reminiscent of the pre-divestiture Bell System, this time as an oligopoly rather than a monopoly. On the upside, the emerging oligopolies would have considerable financial muscle and the capability to direct resources to the modernization of the telecom infrastructure. On the downside, they would lack incentives for dynamic innovation, much like the pre-divestiture Bell System (Economides 2001).

A second path could be labeled *competitive bankruptcy* (Varian 2002). MCI/Worldcom recently became the largest telecommunication company to file for bankruptcy protection, and as a result will be able to operate with a much lower cost structure. This certainly puts significant pressures on other telecommunications carriers to follow suit, and place themselves on a equal competitive basis with MCI/Worldcom. The unfolding of this scenario would result in an industry structure essentially unchanged from the current situation, but with players reduced to managing existing assets and unable to secure new financing to modernize their network.

A group of Internet and telecommunications experts has recently argued in an open letter to FCC Chairman Powell that both of these paths would only forestall the inevitable demise of the current telecom infrastructure[1]. These two paths, they say, amount to artificial life-support for a network whose underlying technology and architecture are fundamentally obsolete. The end-to-end architecture will ultimately win out, as the right way to proceed on to the next generation, and force incumbent telecommunications companies out of business. As a result, they advocate a third path of *fast failure*, arguing that we would be better off to let the existing telecommunications carriers fail quickly, to allow for the rapid rebuilding of the network along sound end-to-end principles. This, they claim, is the only approach that could unleash innovative forces and make it possible

---

[1] See *A Letter to FCC Chairman Michael Powell: Support "Fail Fast"*, October 21, 2002, available at http://netparadox.com/

to move on. The right way to proceed however is far from obvious, because "the best network is the hardest one to make money running" (Isenberg 2002). The best network is open, provides end-to-end connectivity, and is "stupid" – passing information through with a minimum of handling. Such a network is technically preferable to other approaches, but one for which there is no straightforward business model.

From a policy standpoint, these three trajectories point to fundamental choices about the transition to next generation networks.

1) Is it possible to build upon the current telecommunications framework or is the current arrangement so hopelessly flawed that an entire new structure should be created?
2) Should the owners of the infrastructure be able to control entirely how that infrastructure is used, or should they be required to make their network available to others on an 'open access' basis?
3) What investment incentives will there be for infrastructure investment in the next generation networks? .

A possible future would be one in which the current long-haul fiber networks are acquired out of bankruptcy court and re-deployed to offer a low-cost, high-bandwidth long-haul infrastructure. Add to that long-haul network a wireless local 'loop', based on next-generation wireless technology –a evolution of current Wi-Fi, ubiquitous, unlicensed, and high-bandwidth local data network. This scenario is one possible outcome for the fundamental rethinking of spectrum policy currently undertaken by the FCC (FCC 2002). Spectrum policy reform could result in a substantial amount of spectrum subject to unlicensed and flexible use, which could constitute a critical local access option for the next generation infrastructure.

It is important to note that each one of these paths out of the current telecom bust leads to a dramatically different infrastructure arrangement, relying on different technologies, with different architectures, different business models, different incentives for investment and innovation. The choices made in the short-term to emerge from the immediate crisis will therefore have significant implications for the ultimate characteristics of the next generation network infrastructure.

The future of end-to-end?

One question that is particularly central to the future arrangement relates to network architecture: is end-to-end fundamental to future success? Is it a 'natural' equilibrium for the network (that is, left to its own devices, would the network naturally evolve in that direction)? Or would it require active policy protection to endure, were we to determine that it is desirable?

It is now well accepted that the internet's success rested in large part on the virtuous innovation cycle made possible by the network's "end-to-end" architecture (Lemley and Lessig 2000). While in past communication networks, intelligence and control were built in the network's core, the internet offered a radically new approach where intelligence and control moved to the edge, in the devices connected to the network. In traditional 'core-intelligent' networks, the network owner (i.e. the phone companies) were the only actors who could innovate, with the internet's end-to-end those who reside at the edge can take part in the innovation process. The internet serves then as a neutral platform which supports innovation from multiple sources. Of course, no platform is ever absolutely "neutral" and even the internet embodies biases –for example against synchronous, low latency applications . However, while the biases embodied in traditional telecommunications networks overwhelmingly favored their owners, the internet is a lot less one-sided (Sandvig 2002).

Regardless of its starting point, today's internet is clearly moving away from pure end-to-end principles and evolving toward the placement more intelligence at its core. Partly this is a result of the its coming of age: as long as the internet was primarily a research network connecting a set of cooperating academics, the neutral network served them well. More fundamentally, as the internet became a ubiquitous commercial and mass medium, strong forces have prompted some of the actors to seek more control, bringing a departure from the end-to-end principles (Blumenthal and Clark 2001). These forces fall in two broad categories, economic and socio-political.

Economic forces lead Internet Service Providers (ISPs) to break from the end-to-end principle as a way to differentiate themselves and provide services levels distinctive from those of their competitors. In order to offer specific Quality of Service (QoS) levels for example, they need to build more intelligence in the core of their networks so that different data streams are given different priorities. Many of the new services they seek to offer, such as telephony, video conferencing, broadcasting, and virtual private network extensions to LANs, are demanding (synchronous) applications. This means that the traditional "best effort" packet delivery associated with end-to-end architecture is often inadequate. Many other strategic offerings, ranging from firewalls, to complex caching and replication services, also push ISPs to build increasing intelligence in the core of their networks. At the same time, their desire to cater to a mass market of less sophisticated users (such as the typical AOL subscriber), pushes ISPs to build more intelligent, easy to use networks. The increased commercialization of the Internet has created a large class of users who value usability more highly than control.

Socio-political forces have also been driving the internet further away from the pure end-to-end principle. For example as the number of users has grown, it has become impossible to know and trust all: in an untrustworthy world, defenses can be (and have been) built within the network. Similarly, it is very difficult for law enforcement agencies to monitor traffic over a pure end-to-end network. Increasing concerns about national security tend to lead to further departures from end-to-end.

This evolution creates a fundamental challenge to the model of an open, end-to-end network. It makes clear that the end-to-end architecture is not a stable, natural outcome for the network, but a state that requires protection if it is to survive. Left to its own devices, the internet would likely move away from the end-to-end principle and become an increasingly less neutral network. The biases it would take on would naturally reflect the political and economic interests of those who control the network. One avenue would be political, with the promotion of policies that preserve the end-to-end character of the internet. That would include some open access requirements, but would presumably require further measures (Lessig 2002).

Ultimately, the question that must be answered as we proceed with the deployment of next generation networks is what balance to strike between the interests of the networks owners and the interests of the users. It is not clear that market forces can resolve this tension with a socially optimal solution, because the interests of the users are both less organized (more diffused, less monetized) than those of the network owners, and less immediate. Indeed, and perhaps more importantly, the gains to be had from multi-source innovation are dynamic, unpredictable future benefits. It is difficult for the market to weigh them against the short- and medium-term benefits accruing to the network owners form increased control. The fundamental challenge will then be to reconcile market forces inevitably introduced by increasing reliance on competition to govern the telecom infrastructure and the maintenance of incentives for open networks that allow users to play an active role in the innovation process.

Will users continue to lead?

The recent history of telecommunication (since the Hush-a-Phone decision in 1958, which marked the official start of deregulation) is marked by the role users have played in shaping the evolution of information networks. The transition from analog to digital technologies as the foundation of networks has unleashed two dynamics of critical importance to the evolution of the communication infrastructure and its uses. First, control over network configuration has become separable from network ownership, so that network owners no longer determine alone what communication patterns and services their networks will support. Second, network users actively contribute to network innovation, now that they can shape the configuration of networks they do not own. As a result, new forces drive network transformation and new network uses arise.

With the advent of digital technologies, network control has become separable from network ownership. Configuring the old analog networks –defining who can use them to communicate with whom, along what communication patterns– often meant physically changing the layout of their hardware components. As a result, only the telephone companies, cable operators or broadcasters who owned these networks could control how they would be used. Because the configuration of today's digital networks is programmable, it has become possible for network users (end-users or corporations) or third parties (such as network service providers) to control the configuration of networks they do not own.

This simple fact has profound implications for the regulation of individual networks and their interconnection. It requires new rules for the definition of network architecture, for sharing access rights to the networks' hardware and software. This is the core of today's hottest policy debates, such as the definition of a framework governing competition for telecom, video and internet services, and the underpinning of today's most ambitious business strategies, such as efforts by corporate network users to articulate new work processes and new supply chains around network architectures they can now influence. The internet offers the best known and most vivid illustration of these evolutionary dynamics, but this virtuous cycle is certainly not limited to the internet. One of its strongest manifestations can be found in corporate networks. There, as corporate strategies, work organizations and networks co-evolve, we can find the most powerful effect of the information networks. They truly sit at the heart of productivity gains that can be had through the network-based re-organization of work.

Indeed network users, empowered by their newfound ability to control network configuration, have become key actors in the evolution of communication networks. Network deployment follows an iterative, cumulative, path-dependent cycle (Bar 2000). Throughout this evolutionary cycle, communication networks constitute both the support and the object of innovation. For example, they help companies articulate organizational change and the experimentation they support often results in the development of new network applications. Case studies of information network use in a variety of industries show that three phases typically mark the unfolding of these cycles. At first, end-users automate their communication activities, taking advantage of new communication technologies to perform old tasks in new ways. In a second phase, they experiment with the possibilities of the communication network that resulted from the initial automation. Network-based experimentation typically then provokes a third phase of re-configuration, where the network's configuration is transformed to incorporate the lessons learned during the previous phases, and the organization of usage is redefined to take advantage of the new infrastructure. As a result of this re-organization, users find themselves with a new communication network infrastructure, upon which they can start automating and experimenting anew, embarking on successive cycles.

Most critical in that respect will be the definition of the architecture of the emerging information infrastructure. Network architecture has profound implications for the fairness and efficiency of the economic activities that use the network (Bar 2001). From a static standpoint, network rules will define how those who do not own a network can draw upon its internal features and control its configuration to create competitive advantage. At issue here is the deployment of a seamless and coherent information infrastructure that allows cross-organizational applications, such as electronic commerce, to be implemented across separately owned networks.

From a dynamic standpoint, these new rules will largely determine who develops, and on what terms, the innovations that will drive future network evolution. Restrictive regimes will tend to favor existing providers as the source of innovation; more open frameworks will tend to sustain the dynamic that has shown such promise with the Internet -- a dynamic in which users are essential drivers of innovation. The user-driven innovation process was made possible by the open nature of the network platform they were using, the first-generation internet or the telecommunication network. The resulting communications services became available to other users through the seamless Internet infrastructure, for these other users in turn to refine further. This process relied on the open availability of standard components (e.g., agreed protocols, file transfer, remote access) within the Internet infrastructure.

Communication policy defines physical and logical network architecture -- it shapes and configures communication spaces that foster certain kinds of interactions and communication patterns. The third key challenge as we engage in the next phase toward next-generation networks will be to maintain that creative dynamic, sustaining user-driven innovation.

Conclusion: Toward Next Generation Networks

The recent financial bust in the information and networking industries forces a re-examination of the forces that led to the preceding boom, and provides an opportunity to reassess both business strategies and public policy as we engage in the next phase, moving toward next generation networks. As we proceed, one fact is certain: information networks have been centrally important to economic growth in the past, as they have supported a profound reorganization of economic processes of production and exchange. They are likely to be as important in the next phase. We are faced with three sets of fundamental choices. Each corresponds with one of the key factors that account for past success, each results from a challenge brought in part by the limits discovered through that success as the forces that drove the last revolution seem to reach some limits.

The first set of choices relates to the infrastructure itself: what strategies and policies will allow us to move past the current malaise and set of recent (and pending) bankruptcies? The basic challenge is to determine whether the past infrastructure model, its industrial structure and usage rules, can carry us forward or whether it is fundamentally flawed. In the first case, some reorganization will be necessary. Further horizontal consolidation among carriers is likely, as is some measure of re-verticalization. But ultimately in this scenario, we would decide to move forward with an infrastructure that largely mirrors the one we have had to date. On the other hand, it may be that the existing model is so fundamentally undermined by the forces unleashed during the boom/bust cycle (a combination of end-to-end architecture and ubiquitous free/open bandwidth through a new form or access to spectrum), that we would be better off quickly ending the agony and proceeding on a new basis, the economic basis of which remains to be determined.

The second set of choices has to do with the related issues of standards and architecture: What overall model should we adopt to govern the interconnection, interoperability and interaction of multiple actors within the next generation network environment? One approach would be to give ever greater independence to the owners of individual pieces of the infrastructure (physical and 'virtual', facilities and services). Each would determine what standards, interconnection modes and architecture best fits its economic and strategic interests. In that world, we would likely see a greater fragmentation and specialization of individual infrastructure segments, where players would compete against one another on the basis of the particular standards and architectures they adopt and the strategic interconnection deals they strike. On the other hand, we may decide that integration across the infrastructure's various segments is more important than the ability for each fragment to be independent and specialized. This would require a new set of policies that codify open access rules and interconnection standards. In this world, we would likely see competition articulated around particular implementations of commonly accepted standards and architectures.

The third set of choices relates to the relative roles of suppliers and users in this new environment: What is the optimal balance of power and control between those who own the next generation infrastructure and those who use it? We know from past experience that users have played a fundamental role in innovating and shaping the evolution of networks. We also know that some segments of the user population have neither the inclination nor the sophistication required to constantly tinker with the network, and would rather have others innovate for them.

Ultimately, the answers we give to these questions could lead our society and our economy in dramatically contrasted directions, supported by different network infrastructures headed on different evolutionary trajectories. They would have substantial impact on the structure of human activities that rely on communication networks, resulting in different distribution of economic power, political and social control. The decisions to be made as we move beyond the current telecommunications turmoil are therefore worth careful consideration.

Bibliography

Abbate, J. (1999). Inventing the Internet. Cambridge, MIT Press.

Bar, F. (2001). The Construction of Marketplace Architecture. Tracking a Transformation: E-commerce and the Terms of Competition in Industries. Washington, DC, Brookings Institution Press**:** 27-49.

Bar, F., with Neil Kane and Caroline Simard (2000). Digital networks and Organizational Change: The Evolutionary Deployment of Corporate Information Infrastructure. International Sunbelt Social Network Conference, Vancouver, British Columbia.

Blumenthal, M. S. and D. D. Clark (2001). Rethinking the Design of the Internet: The End-to-End Arguments vs. the Brave New World. Communications Policy in Transition: The Internet and Beyond. B. M. Compaine and S. Greenstein. Cambridge, MA, MIT Press**:** 91-140.

Economides, N. (2001). "Coming Apart, Coming Together: The AT&T Breakup (Round Three) and the Remonopolization of Telecommunications." SternBusiness.

FCC (2002). High-Speed Services for Internet Access: Status as of December 31, 2001, Industry Analysis and Technology Division - Wireline Competition Bureau.

FCC (2002). Spectrum Policy Task Force Presents Recommendations for Spectrum Policy Reform.

Hart, J., R. Reed and F. Bar (1992). "The Building of the Internet: Implications for the Future of Broadband Networks." Telecommunications Policy **16**(8): 666-689.

Isenberg, D., and David Weinberger (2002). The Paradox of the Best Network.

Lemley, M. A. and L. Lessig (2000). "The End of End-to-End: Preserving the Architecture of the Internet in the Broadband Era." The Berkeley Law & Economics Working Papers **2000**(2).

Lessig, L. (2001). The Future of Ideas: The Fate of the Commons in a Connected World. New York, Random House.

Lessig, L. (2002). A threat to innovation on the web. Financial Times.

Noam, E. (2001). The Next Frontier for Openness: Wireless Communications. 29th TPRC Conference, Alexandria, VA.

Odlyzko, A. (2002). What caused the telecom crash: Technology, business, or demand? ISSCC2002, San Francisco, California.

Odlyzko, A. M. (2000). "Internet growth: Myth and reality, use and abuse." iMP: Information Impacts Magazine.

Saltzer, J. H., D. P. Reed, et al. (1984). "End-to-end arguments in system design." ACM Transactions on Computer Systems **2**(4): 277-288.

Sandvig, C. (2002). Communication Infrastructure and Innovation: The Internet as End-to-End Netwokr that Isn't. A Symposium with the Next Generation of Leaders in Science and Technology Policy, organized by the AAAS, Columbia University and Rutgers University, Washington DC.

Varian, H. (2002). Restoring Financial Health to the Telecommunications Sector. En Banc Hearing on Financial State of the Telecommunications Sector, FCC.

Werbach, K. (1997). Digital Tornado: The Internet and Telecommunications Policy, Federal Communication Commission.