PART I

Creative World

Ellie Rennie

INTRODUCTION CREATIVE WORLD

Two hundred years of American technology has unwittingly created a massive cement playground of unlimited potential. But it was the minds of 11 year olds that could see that potential.

(Craig Stecyk, 1975)

The creative world might not look much different from what came before. But as with the skateboarders who took the architecture of the city and saw in its shapes the potential for speed and style, creativity is about repurposing, subverting, and improving what is already there. Even now, as those who invented the term "creative industries" review its boundaries and inroads, the creative world is doing what it likes with the idea.

The general introduction outlines how creative industries came about through the forces of globalization, including changes in national and international economic patterns and in culture and communication. "Creative industries" suggests a new organizing principle to fit a reorganized world, where opportunity is located in unusual spaces: in knowledge, ideas, relationships; and in both local and global communities. Policies that seek to expand and mobilize creativity display a new awareness of how selection, sanctioning, and control can inhibit as well as support. It is therefore as much about the creative conditions and spaces where creative participation can occur as it is about the products themselves. It is also an acknowledgment that creativity is not just for the talented few, but a dynamic being picked up and pursued by people and groups in a range of contexts. So although "creative industries" deals with the vast, prevailing ramifications of a global economy, it also an idea that works from the ground up.

Craig Stecyk, whose words open this chapter, is an artist, writer, and photojournalist who took his passion for skateboarding and used it to redefine youth culture. The quotation is from a series of articles that he put together under a range of pseudonyms about a group of street kids called the Z-Boys. In the mornings the Z-Boys (and girl) would surf through the ruins of the dilapidated theme park at Pacific Ocean Park, Venice Bay – the "sea-side slum" they called Dogtown. When the ocean died down they would find concrete waves to ride: asphalt ditches around schools and later swimming pools in wealthy areas (locate, drain, skate, and clear out before the cops arrived). When they formed their team in the 1970s, skateboards were the equivalent of the hoola-hoop. Now the graffiti images that marked out their territory and decorated their boards are the symbols a massive skateboarding industry. You can learn more about the Z-Boys in a documentary that tells their story from their origins as surfers and skateboarders of "aggressive localism and outcast behavior" to their fame as leaders in sport and style (*Dogtown and Z-Boys* 2002). Not only did they "employ the handiwork of the government/corporate structure in a thousand ways that the original architects could never dream of" (Craig Stecyk in *Dogtown and Z-Boys* 2002) they also turned it into a creative industry. The documentary *Dogtown and Z-Boys*, which won two awards at Sundance film festival, was directed and co-written by Stacy Peralta, one of the original Z-Boys members. The production designer and co-writer? Craig Stecyk, of course.

Open Source

Innovation from the ground up is an idea that is catching on. Charles Leadbeater has written in one of his thought-pieces:

> Open source is a new model of networked, citizen-led innovation, which could have huge implications for other fields, in which knowledge and creativity are highly dispersed. For instance, the Natural History Museum, working with Lancaster University, is recruiting a small army of citizen naturalists to help it monitor biodiversity among unfashionable invertebrates, algae, ferns and lichen. The Museum has 350 scientists. Its aim is to augment its capacity by working with a fieldwork force of several thousands. (Leadbeater 2003b: 25)

Suddenly, amateurs – even people who love algae – are important: "In the future, people will have to work with, learn from and sometimes compete with them" (Leadbeater 2003b: 25). This activity arises from non-commercial spaces, where people engage in activities for personal fulfillment or community involvement. "Creative industries" recognizes that there is a "world" of ideas out there where creative pursuits are born. Leadbeater takes inspiration for this new creative arrangement in the phenomenon of open source. Open source is a design feature that enabled participation in the internet by as many people as possible, resulting in its rapid growth. To some it is a technological device, for others a hobby, for many a movement. For thinkers such as Leadbeater who are concerned with the future of innovation it is also a blueprint for economic and social advancement.

Commons on the Wires: Lawrence Lessig

In the first reading of this section, Lawrence Lessig explains how the internet was designed so that control of the network lay with the end-users rather than at the center, allowing for greater participation in the innovation of the technology. He endorses a model of innovation that is not dictated by big business or the talented few, but by the interest and time of enthusiasts everywhere. The architecture Lessig describes, End-to-end (e2e), is made possible through "packet-switching" – a system of protocols that shift packets of data by labeling and routing them to their destination along whatever pathway is most convenient at the time (see also Froomkin 1997). No central machine is necessary as control lies at the "ends" of the network (with the end user via the protocols they send). As a result, no one needs to get permission to participate in the internet (Leiner et al. 2000).

"Open source" means that the code used to construct a computer program is visible to all users, not just the creators. During the internet's early development, open source code was used to construct the architecture of the internet, which in turn inspired openness at the applications layer. If code is visible to all, then anyone can build new layers of protocols, and thereby develop new applications or produce new versions of existing ones. If, on the other hand, the code is hidden (as it is in much proprietary software), users are unable to copy, refigure or adapt the application. Lessig argues that a degree of "openness" is necessary for a technology to grow and develop. The two principles of e2e and open source meant that the internet could develop in any direction its users wanted it to develop. As David Reed puts it in Lessig's excerpt, the design was about presuming as little as possible, and not "running a bake-off." The internet needed to be "out of control" so that as many people as possible could contribute to its growth. For Lessig this is a key to innovation.

Despite this simple idea, the internet's early contributors had a hard time persuading the telecommunications industry to pay attention. What the experts *didn't get* was that the old business model of centralized network control, managed by a core group of highly trained professionals, was not necessarily the best path to technological growth. Business and government were still important in the rise of the internet, as Lessig explains at length in the book. But his argument also represents a significant shift in the role of creativity that occurs outside of these structures. For the creative industries, activities that might have once been considered "below the radar" in terms of the economic wealth of cities and nations have become important.

Perhaps the best known example of open source innovation is the operating system Linux. Before Linux there was Unix, invented by computer scientists at the US telecommunications company AT&T. As Unix could not be sold owing to legislation that prevented AT&T from participating in the computing industry, its inventors convinced the company to give it away, retaining its open source design. However, when the legislation changed in 1984 and the barrier was lifted, the company decided to make Unix proprietary, removing the ability for others to distribute and develop it. By this time, "a generation had devoted its professional career to learning and building upon the Unix system" (Lessig 2001: 53). Understandably they felt betrayed. Computer programmer and free software advocate Richard Stallman decided to develop a free version of Unix, which was later linked to a concurrent project developed by Finnish computer science student Linus Torvalds, creating GNU/Linux (otherwise known as just Linux). Linux is now the fastest-growing operating system in the world and it is estimated to be the result of the efforts of over 100,000 independent and voluntary enthusiasts. Due to its transparency as open source code, Linux is considered by many to be a far more robust system than Windows. This is an example of one not-so-"small army of citizens" that contributed to the growth of a technology. As JC Herz (in part V) points out in relation to the games industry, credence should be granted "to the collective intelligence of the network – the fact that a million people will always be smarter than 20 people and that there is a business value in that differential."

Open Commons

If this is how creativity begins, then where does it end? The main argument of what has become known as "the commons" debate is that as much as innovation can be encouraged by providing spaces within which people are free to use technology, the same technology can also become restrictive.

Code can be built that hides the architecture of the technology, thereby restricting people's capacity to adapt it and build new technologies. For Lovink, "the image of World Wide Web ghost town pops up, abandoned home pages, bored avatars, broken links, switched-off servers, controlled communities, spam-flooded email lists, and newsgroups. The freedom is there, but no one cares and no one will be able to find the counter-information through the corrupted portals and search engines at any rate" (Lovink 2002: 239). The fear within Lovink's dark vision is that the inter-creativity that defined the early internet will be dramatically reduced as old business structures take over – hardly what we could call a "new economy."

One example of this prospect is given by Lawrence Lessig and Mark Lemley (1998) in their submission to the FCC's investigation into the AT&T/MediaOne merger. Their concern was over the bundling of internet service providers (ISPs) with access to broadband infrastructure (note what is not allowed):

> The consequence of this bundling will be that there will be no effective competition among ISPs serving residential broadband cable. The range of services available to broadband cable users will be determined by one of two ISPs – @Home and RoadRunner, both of whom would be allied with the same company. These ISPs will control the kind of use that customers might make of their broadband access. They will determine whether, for example, full length streaming video is permitted (it is presently not); they will determine whether customers might resell broadband services (as they presently may not); it will determine whether broadband customers might become providers of web content (as they presently may not). These ISPs will have the power to discriminate in the choice of Internet services they allow, and customers who want broadband access will have to accept their choice. Giving this power to discriminate to the owner of the actual network wires is fundamentally inconsistent with End-to-End design. (para. 52, p. 11)

By determining the way in which the technology would be used and developed through the ISP, the end-to-end principle is potentially compromised. This suggests that in the new economy, fundamental concepts such as property and ownership require rethinking. Owning an idea (copyright) can add value, encourage distribution of an idea, and provide recompense and incentive for its author, but it can also restrict use of that idea by others (see also Meikle's discussion on open source licences in this section). The "creative industries" idea requires thinking about how intellectual property can create value and reward, but it also requires a conscientious approach to corporate enclosure that may restrict the emergence of new creative ideas.

When it comes down to it, this necessary balancing act alerts us to the fact that the new economy is not a natural equality. This is despite the fact that involvement in creative pursuits, individual flair, and good judgment seem as if they *should* benefit everybody; and even if recipes and software can be used again and again without being depleted (see Leadbeater's piece in part II). Castells (1996) writes of the nodes in the network where it is good to be and where success seems to strike out of nothing, but he also tells of nodes which are unlit, disadvantaged places. We also know that where knowledge plays a significant part in the economy, the gaps between rich and poor are likely to be greatest (Leadbeater 2003a). In such an economy success comes down to the provision of education and an environment for good ideas to flourish. It needs to be asked what else can be done for the creative world to be realized.

For Lessig, government has a role to play in ensuring that some spaces provide access to anyone – a portion of US broadband cable bandwidth, for instance. David Bollier writes that:

> Any sort of creative endeavour – which is to say progress – requires an open "white space" in which experimentation and new construction can take place. There must be the *freedom* to try new things and an unregimented workspace in which to imagine, tinker and execute new ideas. When all the white space is claimed and tightly controlled through commercial regimes that impose quantitative indices and quarterly profit goals, creativity is bureaucratized into narrow paths. There is no room for the visionary ideas, the accidental discoveries, the serendipitous encounters, the embryonic notions that might germinate into real breakthroughs, if only they had the space to grow. An argument for the commons, then, is an argument for more "white space." (Bollier 2001: 5)

For Lovink, it is about economic competence and diverse participants: "if we still have the naïve idea that an open and diverse cyberculture can somewhat influence the course technology is taking," the best approach is to "start up businesses and pollute the concepts used under the umbrella of the term New Economy" (in Meikle 2002: 177). The "creative industries" brings new ideas to old issues of inequality, primarily centered on what can be done to provide opportunities for creative participation. Even if it is the case that creative success is a luxury of the talented or entrepreneurial minority (see Howkins, in part II), a larger creative ecology can be seen to be emerging. The "white space" within which this can flourish is not just a matter of technology, but of resources and culture.

Self-Creating (not "Developing") Countries

At the Opening of New Media Center Sarai, Delhi: Geert Lovink

Geert Lovink's article tells of a creative space in Delhi, India, where media tools are provided to assist the creative efforts of artists, activists, and theorists. The media producers at Sarai draw their content from the city itself – for them a place that is much more than a node in a global economy. This is a picture of creativity that is intimately tied to a space and culture, where the Delhi's stories are told and listened to, but where any predictable form of cultural representation gets challenged. It shows how creative development involves more than policy choices around sustainability, resource provision, or information dissemination. Creative production allows the people involved to determine their own discourses – to describe themselves and articulate that locally and globally. As a creative center, Sarai is actively engaged in the intellectual and policy debates that it finds itself implicated in, including the role of new technologies, the crisis of development and the discourse of the digital divide.

The development aspects of the creative industries endorse creative participation as a means towards post-scarcity. This promotes an open-ended view of culture in which people can pursue opportunities through self-expression and creative production. Critique and inventiveness are likely outcomes of such a development rationale – where people and groups such as Sarai have the means to articulate and promote their own views and ideas. Where development policies in the past have focused on the alleviation of need – implying lack and helplessness in the receiver – groups such as Sarai actively refute (and disprove) connotations of powerlessness. Lovink quotes Jabeesh Bagchi, a Sarai member also from Raqs Media Collective:

> Development often implies the notion of victims of culture. I don't think in those terms. People live, struggle, renew, invent. Also in poverty people have a culture. I feel a little lost in this terrain, knowing that Sarai, to a large extent, is financed through development aid programs. I would never use a term like "digital divide." We have a print divide in India, an education divide, a railway divide, an airplanes divide. The new economy in India is definitely not

> conceived as a divide. It is a rapid expansion of digital culture. The digital divide is a "social consciousness" term, born of guilt. We should interpret the media in different terms, not just in terms of haves and have nots. (Jabeesh Bagchi in Lovink 2002: 210)

The assertion of local culture and ideas in groups such as Sarai, through the use of communications technologies, challenges the way that Asia, Africa, and Latin America are understood in the global trajectory. This is not about bringing technology to the poor (although Sarai and other such groups admit to being dependent upon development funds for their survival); it is about the creative expression of places and communities – something which, as Bagchi puts it, is not a matter of "haves and have nots." If a new economy is about ideas, knowledge, and creativity, then new configurations of development arise that can be constructed upon existing knowledge rather than need. Post-development theorist Arturo Escobar has written that

> We need to consider how a postdevelopment frame of communication practice may be linked with the idea of place as project, that is, with the potential to elevate local knowledges into different constellations of knowledge and power through enabling networks. (Escobar 2000: 171)

Sarai, Delhi, is one such place.

These ground-up initiatives can be helped at a national level through policy initiatives. Creative industries approaches seek policy solutions that allow communities to assert their cultural uniqueness. This involves both small-scale, community-empowering projects as well as nationally implemented schemes (see also Stuart Cunningham's introduction to part V of this volume).

Multicultural Policies and Integration via the Market: Néstor García Canclini

Néstor García Canclini argues that the failure to achieve greater political and cultural integration in Latin America has been a failure of policy. In particular he argues that cultural policy in Latin American nations has remained confined to officially sanctioned monuments, heritage, and fine arts and has remained largely nationally based. Therefore, attempts to achieve cultural policy on a continent-wide scale have remained confined to high culture and to monuments and folkloric heritage, giving preference to "a conservationist vision of identity and to an integrationist view based on traditional cultural goods and institutions."

At the same time, there has been a rapid expansion and uptake of electronic communications media in all of its forms – both US-based transnational media organizations as well as the expansion of Latin American-based conglomerates. This means that, "for the first time in history, the majority of commodities and messages received in each nation has not been produced in their own territory, do not result from the particular relations of production, and do not convey messages connected exclusively with given regions." They operate "according to a transnational, deterritorialized system of production and diffusion."

Importantly, Canclini does not propose a defensive cultural nationalism or a reversion to the "strong state" as an alternative to globalized media culture disseminated by transnational media corporations. The privatization of broadcasting and communications has weakened the capacity of Latin American states to intervene to ensure cultural diversity and equal opportunities for participation in communication. However, Canclini also recognizes that popular media are a source of cultural dynamism in Latin America, as are the large networks of independent educational, cultural, and communications organizations that operate largely outside of the domain of the nation-state. Canclini instead proposes the development of a "Latin American audiovisual space" that could enable the expansion of production and markets for domestic producers while having some capacity to regulate the flows of capital and product from outside of Latin America – most notably, of course, from the United States – and which could enable a greater degree of harmonization of development of the corporate, state-funded and independent sectors in forms compatible with the development of democratic citizenship in multicultural societies.

Both Canclini and Sarai see the possibilities for creativity in India and Latin America as existing beyond the binary opposition between the state and the market. It lies instead in a newly constructed idea of public space – not unlike the commons – spaces where civil society initiatives can flourish: "social movements, artists' groups, independent radio and television stations, unions, ethnic groups, and associations of consumers, radio listeners, and television viewers. Only the multiplication of actors can favour a democratic cultural development and the representation of multiple identities."

Resistance is Fertile

Spaces that once did not fit comfortably with ideas of progress and development are now where new ideas of progress and development can be found. Creativity can be the result of resistance and local culture.

Open Publishing, Open Technologies: Graham Meikle

Creative producers sometimes work in isolation, but more often as part of a group, and sometimes for a cause. Graham Meikle's *FutureActive* tells of the development of Indymedia (an open publishing web forum) from its origination as a calendar of events and actions for Sydney's community groups to a network of over 70 locally based sites around the world. Meikle shares Lessig's concern over the future of the internet, advocating open and non-proprietary spaces (what he calls "Version 1.0") over market enclosure ("Version 2.0"). His focus on the activist uses of the internet gives the sense of an information commons that is inhabited by groups and individuals with something to say – be that through electronic art, radio, open publishing, blogging, or hacking. Meikle grounds the technological achievement of the internet with a sense of place, demonstrating how innovation arises out of locally significant communities and cultures.

The emphasis on the local has ramifications for the way we see cultures of resistance and community-based media more generally. Compare this to a 1984 report by the UK research and development agency, Comedia, which found that the alternative press failed because it was not aggressive enough in positioning itself within the market, either through unwillingness or inability (discussed in Atton 2002: 33). The survival of the alternative press, in their opinion, was only due to subsidy in the form of "self-exploited labor" and benefit gigs donated by the music industry. The report found that the volunteerism that sustained the publications was merely the result of a "commitment to squatting and claiming as a way of life" (Comedia in Atton 2002: 36). The implication was that volunteerism was a pursuit to be associated only with extremist behavior. What is interesting about Comedia's findings in retrospect is the way in which "squatting and claiming" has become a creative pursuit that has been legitimated by the open source/commons movement. It is a shift towards acknowledging the places where creativity arises and seeing within that "natural" activity an opportunity for new ideas, and sometimes industries, to grow. Moreover, it suggests that the traditional binaries around work and personal fulfillment are shifting (see Angela McRobbie in part VI of this volume).

The creative industries approach is interested in how a creative world works as a seed-bed for new ideas. The idea of increased participation sits better with creative policies that seek to expand creativity than with previous cultural policies more concerned with its improvement and status (high art and culture). But although the radical, oppositional, and local are recognized as sites from which creative industries might develop, does "creative industries" leave room for dissent and critique? Is it true, as McRobbie (2001) argues, that creative industries is a policy that wants to turn "angry social critics" into successful commercial artists with little time for thinking about other matters? The creative spaces described by Meikle – alternative or tactical media – are a field of creative production that fits uneasily with ideas of commercially driven creativity intended to exploit global flows of capital and culture. These cultures pose a direct challenge to the idea that knowledge should be commodified, often to the idea of property itself. The argument that alternative culture inspires the music industry to new heights or gives the design industry a new (street) style, will grate with some. It appears that "creative industries" and its governmental, industry, and intellectual frameworks co-opt alternative culture into a commercial box that it resists by its very nature. Geert Lovink argues that:

> Alternative has been effectively reduced to style. In the media context, this means that we can no longer sell a certain forum – website, radio station, zine – as subversive or even revolutionary. It will have the immediate danger of being turned into a fashion, a lifestyle item. (Lovink in Meikle 2002: 112)

The reality is that the alternative and mainstream are getting more difficult to delineate, at least in an aesthetic sense, and this is part of the forces of the new economy. But this does not cancel out resistance. What is left is a series of critiques and alternative visions that seek to interrupt and challenge presumed knowledge and authority, described increasingly in alternative media theory as the "tactics" of the weak (de Certeau 1984; Couldry 2000; Klein 2000).

Linux, for instance, is not just an operating system, it is also a movement – best summed up by writer Neal Stephenson (1999) through his analogy of car yards. As he explains it, 90 percent of people will go to the biggest car dealer and buy the Microsoft station wagons, passing by the group who are camped by the road giving away free Linux tanks "made of space-age materials and sophisticated technology from one end to the other" (p. 7) – tanks that never break down and can be used on any street. But despite the tanks' superior qualities, most people won't go near a bunch of "hackers with bullhorns" trying to give stuff away by the side of the road. So what is the value in this activity? As much as alternative cultures can produce new ideas and systems, they will still present *alternatives* even within a creative approach to culture. In this way, the issue is bigger than Lessig's point about innovation. It is important to remember that the culture of this activity is alluring to volunteers and alienates others. There is a dialogue going on within those choices about where our creative pursuits should take us.

Creative industries would be a one-dimensional concept if it did not take this activity into account. Moreover, it would be incomplete in its understanding of creativity, denying the creative invention that exists in dissent and the imagining of alternative futures. Just as creative industries are a response to globalization, so are the new social movements that Meikle is describing. Joshua Karliner (2001) from the group Corpwatch describes the anti-globalization movement:

> the vast majority of our movement can be characterized as engaged in a debate with the corporate globalizers as to the direction modernity should take, rather than across the board opposition to it and advocacy for a return to strict, traditional values.

Alternative media is the site upon which much of this critique occurs. It is the creative manifestation of new economy critique that is not separate from it but a part of the self-reflexiveness of a knowledge society.

Creative production that arises out of amateur and alternative spaces (the "third" space beyond industry and government) has generally received little attention in cultural policy settings outside of the confines of community development. Some of the readings in this volume maintain that it cannot fit well with the creative industries framework either – at least where that is defined as government-designed initiatives directed at intellectual property go-getters (see Howkins in part II). Under such a remit there is not much room for creativity that is less capitalistic in its pursuit, and especially not for that which is outright disruptive. Indeed, it is doubtful that some of the radical and oppositional groups that are included under the banner of "alternative" would sit down to a government review on cultural policy anyhow. But in terms of the broader fate of innovation and creativity, this activity does figure. At one level it generates ideas, images, and styles that have commercial significance, as well as new methods of organization, collaboration, and training. Perhaps more importantly, the alternative and the amateur have signaled some critical themes in the social and economic dynamics that the creative industries term was invented to respond to and describe – tensions between property and freedom, work and personal fulfillment, dissent and government.

References

Atton, C. (2002) *Alternative Media*. Sage, Thousand Oaks.

Bollier, D. (2001) *Public Assets, Private Profits: Reclaiming the American Commons in an Age of Market Enclosure*. New America Foundation, Washington, DC.

Castells, M. (1996) *The Rise of the Network Society*, vol. 1. Blackwell, Massachusetts.

Couldry, N. (2000) *The Place of Media Power: Pilgrims and Witnesses of the Media Age*. Routledge, London.

de Certeau, M. (1984) *The Practice of Everyday Life*. University of California Press, Berkeley.

Dogtown and Z-Boys (motion picture; 2002), Sony Pictures Classics. Director Stacy Peralta, producer Agi Orsi, written by Stacy Peralta and Craig Stecyk.

Escobar, A. (2000) Place, Power and Networks in Globalization and Postdevelopment. In K. G. Wilkins (ed.), *Redeveloping Communication for Social Change*. Rowman & Littlefield, Lanham, 163–74.

Froomkin, A. M. (1997) The Internet as a Source of Regulatory Arbitrage. In B. Kahin and C. Nesson (eds.), *Borders in Cyberspace: Information Policy and the Global Information Infrastructure*. MIT Press, Cambridge Mass., 129–63.

Karliner, J. (2001) Where Do We Go From Here? *OpenDemocracy*. <http://www.opendemocracy.net> (accessed November 6, 2001).

Klein, N. (2000) *No Logo: No Space, No Choice, No Jobs*. Flamingo, London.

Leadbeater, C. (2003a) Seeing the light. *RSA Journal* 5505 (February), 28–33.

Leadbeater, C. (2003b) Amateurs: a 21st-century remake. *RSA Journal* 5507 (June), 22–5.

Leiner, B. M., V. G. Cerf, D. D. Clark, R. E. Kahn, et al. (2000) *A Brief History of the Internet*. ISOC, <http://www.isoc.org/internet/history/brief.shtml> (accessed October 2, 2001).

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

Lessig, L. and M. A. Lemley (1998) "In the Matter of AT&T/Media One."

Lovink, Geert (2002) *Dark Fiber: Tracking Critical Internet Culture*. MIT Press, Cambridge Mass. and London.

McRobbie, A. (2001) "Everyone is Creative": Artists as New Economy Pioneers? *OpenDemocracy*, <www.opendemocracy.net> (accessed August 30, 2001).

Meikle, G. (2002) *Future Active: Media Activism and the Internet*. Pluto Press, Annandale NSW.

Stephenson, N. (1999) *In the Beginning was the Command Line*. Avon Books, New York.

1 *Lawrence Lessig*

COMMONS ON THE WIRES

The internet is a network of networks. In the main, these networks connect over wires. All of these wires, and the machines linked by them, are controlled by someone. The vast majority are owned by private parties – owned, that is, by individuals and corporations that have chosen to link to the Net. Some are owned by the government.

Yet this vast network of privately owned technology has built one of the most important *innovation commons* that we have ever known. Built on a platform that is controlled, the protocols of the Internet have erected a free space of innovation. These private networks have created an open resource that any can draw upon and that many have. Understanding how, and in what sense, is the aim of this chapter.

[. . .]

The Internet is not the telephone network. It is a network of networks that sometimes run on the telephone lines. These networks and the wires that link them are privately owned, like the wires of the old AT&T. Yet at the core of this network is a different principle from the principle that guided AT&T. [. . .]

First described by network architects Jerome Saltzer, David Clark, and David P. Reed in 1981, this principle – called the "end-to-end argument" (e2e) – guides network designers in developing protocols and applications for the network. End-to-end says to keep intelligence in a network at the ends, or in the applications, leaving the network itself to be relatively simple.

"Commons on the Wires" from Lawrence Lessig (2001), *The Future of Ideas: The Fate of the Commons in a Connected World*. Random House, New York, pp. 25, 34–7, 39–48, 275–8 (notes). Reprinted by permission of Random House, Inc and International Creative Management, Inc.

There are many principles in the Internet's design. This one is key. But it will take some explaining to show why.

Network designers commonly distinguish computers at the "end" or "edge" of a network from computers within that network. The computers at the end of a network are the machines you use to access the network. (The machine you use to dial into the Internet, or your cell phone connecting to a wireless Web, is a computer at the edge of the network.) The computers "within" the network are the machines that establish the links to other computers – and thereby form the network itself. (The machines run by your Internet service provider, for example, could be computers within the network.)

The end-to-end argument says that rather than locating intelligence within the network, intelligence should be placed at the ends: computers within the network should perform only very simple functions that are needed by lots of different applications, while functions that are needed by only some applications should be performed at the edge. Thus, complexity and intelligence in the network are pushed away from the network itself. Simple networks, smart applications. As a recent National Research Council (NRC) report describes it:

> Aimed at simplicity and flexibility, [the end-to-end] argument says that the network should provide a very basic level of service – data transport – and that the intelligence – the information processing needed to provide applications – should be located in or close to the devices attached to the edge [or ends] of the network.[1]

The reason for this design was flexibility, inspired by a certain humility. As Reed describes it, "we wanted to make sure that we didn't somehow build in a feature of the underlying network technology . . . that would restrict our using some new underlying transport technology that turned out to be good in the future. . . . That was really the key to why we picked this very, very simple thing called the Internet protocol."[2]

It might be a bit hard to see how a principle of network design could matter much to issues of public policy. Lawyers and policy types don't spend much time understanding such principles; network architects don't waste their time thinking about the confusions of public policy.

But architecture matters. And arguably no principle of network architecture has been more important to the success of the Internet than this single principle of network design – e2e. How a system is designed will affect the freedoms and control the system enables. And how the Internet was designed intimately affected the freedoms and controls that it has enabled. The *code* of cyberspace – its architecture and the software and hardware that implement that architecture – regulates life in cyberspace generally. Its code is its law. Or, in the words of Electronic Frontier Foundation (EFF) cofounder Mitch Kapor, "Architecture is politics."[3]

To the extent that people have thought about Kapor's slogan, they've done so in the context of individual rights and network architecture. Most think about how "architecture" or "software" or, more simply, "code" enables or restricts the things we think of as human rights – speech, or privacy, or the rights of access.

That was my purpose in *Code and Other Laws of Cyberspace*. There I argued that it was the architecture of cyberspace that constituted its freedom, and that, as this architecture was changed, that freedom was erased. *Code*, in other words, is a *law* of cyberspace and, as the title suggests, in my view, its most significant law.

But in this book, my focus is different. The question I want to press here is the relationship between architecture and innovation – both commercial innovation and cultural innovation. My claim is that here, too, code matters. That to understand the source of the flourishing of innovation on the Internet, one must understand something about its original design. And then, even more important, to understand as well that changes to this original architecture are likely to affect the reach of innovation here.

So which code matters? Which parts of the architecture?

The Internet is not a novel or a symphony. No one authored a beginning, middle, and end. At any particular point in its history, it certainly has a structure, or architecture, that is implemented through a set of protocols and conventions. But this architecture was never fully planned; no one designed it from the bottom up. It is more like the architecture of an old European city, with a central section that is clear and well worn, but with additions that are many and sometimes confused.

At various points in the history of the Net's development, there have been efforts at restating its principles. Something called "RFC 1958," published in 1996, is perhaps the best formal effort. The Internet was built upon "requests for comments," or RFCs. Researchers – essentially grad students – charged with the task of developing the protocols that would eventually build the Internet developed these protocols through these humble requests for comments. RFC 1 was written by Steve Crocker and outlined an understanding about the protocols for host ("IMP") software. Some RFCs specify particular Internet protocols; some wax philosophical. RFC 1958 is clearly in the latter camp – an "informational" document about the "Architectural Principles of the Internet."[4]

According to RFC 1958, though "[m]any members of the Internet community would argue that there is no architecture," this document reports that "the community" generally "believes" this about the Internet: "that the goal is connectivity, the tool is the Internet protocol and the intelligence is end-to-end rather than hidden in the network."[5] "The network's job is to transmit datagrams as efficiently and flexibly as possible. Everything else should be done at the fringes."[6]

This design has important consequences for innovation – indeed, we can count three:

- First, because applications run on computers at the edge of the network, innovators with new applications need only connect their computers to the network to let their applications run. No change to the computers within the network is required. If you are a developer, for example, who wants to use the Internet to make telephone calls, you need only develop that application and get users to adopt it for the Internet to be capable of making "telephone" calls. You can write the application and send it to the person on the other end of the network. Both of you install it and start talking. That's it.
- Second, because the design is not optimized for any particular existing application, the network is open to innovation not originally imagined. All the Internet protocol (IP) does is figure a way to package and route data; it doesn't route or process certain kinds of data better than others. That creates a problem for some applications (as we'll see below), but it creates an opportunity for a wide range of other applications too. It means that the network is open to adopting applications not originally foreseen by the designers.
- Third, because the design effects a neutral platform – neutral in the sense that the network owner can't discriminate against some packets while favoring others – the network *can't* discriminate against a new innovator's design. If a new application threatens a dominant application, there's nothing the network can do about that. The network will remain neutral regardless of the application.

The significance of each of these consequences to innovation generally will become apparent as we work through the particulars that follow. For now, all that's important is that you see this design as a *choice*. Whether or not the framers of the network understood what would grow from what they built, they built it with a certain philosophy in mind. The network itself would not control how it would grow. Applications would. That was the key to end-to-end design. As the inventor of the World Wide Web, Tim Berners-Lee, describes it:

> Philosophically, if the Web was to be a universal resource, it had to be able to grow in an unlimited way. Technically, if there was any centralized point of control, it would rapidly become a bottleneck that restricted the Web's growth, and the Web would never scale up. Its being "out of control" was very important.[7]

[. . .]

The Internet isn't the only network to follow an end-to-end design, though it is the first large-scale computer network to choose that principle at its birth. The electricity grid is an end-to-end grid; as long as my equipment complies with the rules for the grid, I get to plug it in. Conceivably, things could be different. In principle, we might imagine that every device you plug into a grid would register itself with the network before it would run. Before you connected, you would have to get permission for that device. The owner of the network could then choose which devices to prohibit.

Likewise, the roads are end-to-end systems. Any car gets to enter the highway grid (put tolls to one side). As long as the car is properly inspected, and the driver properly licensed, whether and when to use the highway is no business of the highway. Again, we could imagine a different architecture: each car might first register with the grid before it got on the highway (the way airlines file flight plans before they fly).

But these systems don't require this sort of registration, likely because, when they were built, such registration was simply impracticable. The electronics of a power grid couldn't handle the registration of different devices; roads were built stupid because smart roads were impossible. Things are different now; smart grids, and smart roads, are certainly possible. Control is now feasible. So we should ask, would control be better?

In at least some cases, it certainly would be better. But from the perspective of innovation, in some cases it would not. In particular, when the future is uncertain – or more precisely, when future uses of a technology cannot be predicted – then leaving the technology uncontrolled is a better way of helping it find the right sort of innovation. Plasticity – the ability of a system to evolve easily in a number of ways – is optimal in a world of uncertainty.

This strategy is an attitude. It says to the world, I don't know what functions this system, or network, will perform. It is based in the idea of uncertainty. When we don't know which way a system will develop, we build the system to allow the broadest range of development.

This was a key motivation of the original Internet architects. They were extremely talented; no one was more expert. But with talent comes humility. And the original network architects knew more than anything that they didn't know what this network would be used for.

As David Reed describes, "[T]here were a lot of experiments in those days," and "we . . . realized that [there] was very little in common [other] than the way they used the network. There were sort of interesting ways that they used the network differently from application to application. So we felt that we couldn't presume anything about how networks would be used by applications. Or we wanted to presume as little as possible. . . . We basically said, 'Stop. You're all right' as opposed to running a bake-off."[8] These designers knew only that they wanted to assure that it could develop however users wanted.

Thus, end-to-end disables central control over how the network develops. As Berners-Lee puts it, "There's a freedom about the Internet: as long as we accept the rules of sending packets around, we can send packets containing anything to anywhere."[9] New applications "can be brought to the Internet without the need for any changes to the underlying network."[10] The "architecture" of the network is designed to be "neutral with respect to applications and content."[11] By placing intelligence in the ends, the network has no intelligence to tell which functions or content are permitted or not. As RFC 1958 puts it, the job of the network is simply to "transmit datagrams." As the NRC has recently concluded:

> Underlying the end-to-end argument is the idea that it is the system or application, not the network itself, that is in the best position to implement appropriate protection.[12]

[. . .] We can now see how the end-to-end principle renders the Internet an *innovation commons*, where innovators can develop and deploy new applications or content *with out the permission of anyone else*. Because of e2e, no one need register an application with "the Internet" before it will run; no permission to use the bandwidth is required. Instead, e2e means the network is designed to assure that the network cannot decide which innovations will run. The system is built – constituted – to remain open to whatever innovation comes along.

This design has a critical effect on innovation. It has been, in the words of the NRC, a "key to the explosion of new services and software applications" on the Net.[13] Because of e2e, innovators know that they need not get the permission of anyone – neither AT&T nor the Internet itself – before they build a new application for the Internet. If an innovator has what he or she believes is a great idea for an application, he or she can build it without authorization from the network itself and with the assurance that the network can't discriminate against it.

At this point, you may be wondering, So what? It may be interesting (at least I hope you think this) to learn that the Internet has this feature; it is at least plausible that this feature induces a certain kind of innovation. But why do we need to worry about this feature of the Internet? If this is what makes the Internet run, then as long as we have the Internet? won't we have this feature? If e2e is in the Internet's nature, why do we need to worry about e2e?

But this raises the fundamental point: The design the Internet has *now* need not be its design *tomorrow*. Or more precisely, any design it has just now can be supplemented with other controls or other technology. And if that is true, then this feature of e2e that I am suggesting is central to the network now can be removed from the network as the network is changed. The code that defines the network at one time need not be the code that defines it later on. And as that code changes, the values the network protects will change as well.

The consequences of this commitment to e2e are many. The birth of the World Wide Web is just one. If you're free from geekhood, you are likely not to distinguish the WWW from the Internet. But in fact, they are quite distinct. The World Wide Web is a set of protocols for displaying hyperlinked documents linked across the Internet. These protocols were developed in the late 1980s by researchers at the European particle physics lab CERN – in particular by Tim Berners-Lee. These protocols specify how a "Web server" serves content on the WWW. They also specify how "browsers" – such as Netscape Navigator or Microsoft's Internet Explorer – retrieve content on the World Wide Web. But these protocols themselves simply run on top of the protocols that define the Internet. These Internet protocols, referred to as TCP/IP, are the foundation upon which the protocols that make the World Wide Web function – HTTP (hypertext transfer protocol) and HTML (hypertext markup language) – run.[14]

The emergence of the World Wide Web is a perfect illustration of how innovation works on the Internet and of how important a neutral network is to that innovation. Tim Berners-Lee came up with the idea of the World Wide Web after increasing frustration over the fact that computers at CERN couldn't easily talk to each other. Documents built on one system were not easily shared with other systems; content stored on individual computers was not easily published to the networks generally. As Berners-Lee writes:

> Incompatibility between computers had always been a huge pain in everyone's side, at CERN and anywhere else. . . . The real world of high-energy physics was one of incompatible networks, disk formats, and character-encoding schemes, which made any attempt to transfer information between computers generally impossible. The computers simply could not communicate with each other.[15]

Berners-Lee thus began to think about a system to enable linking among documents – through a process called "hypertext" – and to build this linking on top of the protocols of the Internet. His ideal was a space where any document in principle could be linked to any other and where any document published was available to anyone.

The components of this vision were nothing new. Hypertext – links from one document to another – had been born with Vannevar Bush,[50] and made famous by Bill Atkinson's HyperCard on the Apple Macintosh. The world where documents could all link to each other was the vision of Robert Fano in an early article in the *Proceedings of the IEEE*.[16] But Berners-Lee put these ideas together using the underlying protocol of the Internet. Hyperlinked documents would thus be available to anyone with access to the Internet, and any document published according to the protocols of the World Wide Web would be available to all.

The idea strikes us today as genius. Its success makes us believe the idea must have been obvious. But what is amazing about the story of the birth of the World Wide Web is how hard it was for Tim Berners-Lee to convince anyone of the merit in the plan. When Berners-Lee tried to sell the plan at CERN, management was unimpressed. As Berners-Lee writes:

> What we hoped for was that someone would say, "Wow! This is going to be the cornerstone of high-energy physics communications! It will bind the entire community together in the next ten years. Here are four programmers to work on the project and here's your liaison with Management Information Systems. Anything else you need, you just tell us." But it didn't happen.[17]

When he went to a meeting of hypertext fans, he could get few to understand the "ah-ha" of hypertext on the Net. For years he wandered from expert to expert, finding none who understood the potential here. And it was only after he started building the Web out, and started informing ordinary people on a hypertext mailing list about the protocols he was developing, that the Net started to grow.

The experts didn't get it. Someone should put that on a bumper sticker and spread it around. Those controlling the resources of the CERN computer lab wouldn't support the technology that would give the world the Web. Only those innovators outside of the control of these managers saw something of the potential for the Web's growth.

Berners-Lee feared that competing protocols for using the Internet would wipe away interest in the WWW. One protocol built about the same time was called Gopher. Gopher enabled the easy display of a menu of options from a site. When you went to a Gopher-enabled site, you would see a list of links that you could then click on to perform some function. Gopher was extremely popular as an Internet application – running on the Internet protocols – and use of Gopher took off in the early 1990s.[18]

But for the purposes that Berners-Lee imagined, Gopher was extremely limited. It would not enable the easy construction of interlinked documents. It was closer to a universal menuing system than a system for linking ideas. Berners-Lee was afraid that this inferior standard would nonetheless stick before the new and better WWW became well known.

His fear, however, was not realized, both because of something Berners-Lee did and because of something the creators of Gopher did – and both are lessons for us.

Berners-Lee was no bully. He was not building a protocol that everyone had to follow. He had a protocol for displaying content on the World Wide Web – the HTML language that Web pages are built in. But he decided not to limit the content that one could get through a WWW browser to just Web pages. Instead he designed the transfer protocol – HTTP – so that a wide range of protocols could be accessed through the WWW – including the Gopher protocol, a protocol for transferring files (FTP), and a protocol for accessing newsgroups on the Internet (NNTP). The Web would be neutral among these different protocols – it would in this sense interconnect.[19]

That made it easy to use the Web, even if one wanted to get access to Gopher content. But the second doing was much more important to the death of Gopher as a standard.

As Berners-Lee describes it, high off its success in populating the world with Gopher, the University of Minnesota – owner of the right to Gopher – suggested it might exercise its rights to charge for the use of the Gopher protocol.[20] Even the suggestion of this terrified developers across the world. (It was, Berners-Lee writes, "an act of treason."[21]) Would developers be hijacked by the university once they depended upon their system? How much would they lose if the platform eventually turned against the developers?

Berners-Lee responded to this by convincing CERN to release the right to the Web to the public. At first he wanted to release the protocol under the GPL, or General Public License (the "GNU General Public License"). But when negotiations over that bogged down, he convinced CERN simply to release the rights into the public domain. Anyone had the right to take and use the protocols of the WWW and build anything upon them that they wanted.[22]

The birth of the Web is an example of the innovation that the end-to-end architecture of the original Internet enabled. Though no one quite got it – this the most dramatic aspect of the Internet's power – a few people were able to develop and deploy the protocols of the World Wide Web. They could deploy it because they didn't need to convince the owners of the network that this was a good idea or the owners of computer operating systems that this was a good idea. As Berners-Lee put it, "I had designed the Web so there should be no centralized place where someone would have to 'register' a new server, or get approval of its contents."[23] It would be a "good idea" if people used it, and people were free to use it because the Internet's design made it free.

Thus two networks – the network built by AT&T and the network we call the Internet – create two different environments for innovation. One network centralizes creativity; the other decentralizes it. One network is built to keep control of innovation; the other constitutionally renounces the right to control. One network closes itself except where permission is granted; the other dedicates itself to a commons.

How did we get from the one to the other? What moved the world governing our telecommunications system from the centralized to the decentralized?

This is one of the great forgotten stories of the Internet's birth. Everyone knows that the government funded the research that led to the protocols that govern the Internet. It is part of the Internet's lore that it was the government that pushed network designers to design machines that could talk to each other.[24] The government in general, and the Defense Department in particular, had grown tired of spending millions for "autistic computing machines."[25] It therefore wanted some system for linking the systems.

Yet we are practically trained to ignore another form of governmental intervention that also made the Internet possible. This is the regulation that assured that the platform upon which the Internet was built would not turn against it.

The physical platform on which the Internet took off came prewired. It was the telephone wires that linked homes to homes. But the legal right to use the telephone wires to link to the Internet did not come preordained. That right had to be earned, and it was regulation that earned it. Nothing guaranteed that modems would be permitted on telephone lines. Even today, countries in Asia regulate the use of modems on telephone lines.[26] What was needed before the revolution could begin was permission to connect the Net to this net.

And what made that permission possible? What made it possible for a different use to be made of the telephone wires from that which AT&T had originally imagined?

Here a second kind of regulation enters the story. Beginning in force in 1968, when it permitted foreign attachments to telephone wires, continuing through the 1970s, when it increasingly forced the Bells to lease lines to competitors, regardless of their purpose, and ending in the early 1980s with the breakup of AT&T, the government increasingly intervened to assure that this most powerful telecommunications company would not interfere with the emergence of competing data-communications companies.

This intervention took many forms. In part it was a set of restrictions on AT&T's permissible businesses. In part it was a requirement that it keep its lines open to competitors.[27] In part it was the general fear that any effort to bias communications more in its favor would result in a strong reaction from the government.

But whatever the mix, and whichever factor was most significant, the consequence of this strategy was to leave open the field for innovation in telecommunications. AT&T did not control how its wires would be used, because the government restricted that control. By restricting that control, the government in effect created a commons on AT&T's wires.

In a way analogous to the technical requirements of end-to-end, then, these regulations had the effect of leaving the network open and hence of keeping the use of the network neutral. Once the telephone system was used to establish a circuit, the system was kept free for that circuit to send whatever data across it the user wished. The network thus functioned as a resource left open for others to use.

This is end-to-end operating at a different layer in the network design. It is end-to-end not at the layer determining the connection between two phones on the telephone system. That connection may well be formed by a system that does not comply with the end-to-end rule.

But once the circuit is connected, then the environment created by the mix of technical principles and legal rules operating upon the telecommunications system paralleled an end-to-end design at the network layer. This mix of design and control kept the telephone system open for innovation; that innovation enabled the Internet.

Are there costs to the e2e design? Do we lose something by failing to control access to the resources – the bandwidth – of the network?

Certainly the Internet is not without its weaknesses. The capacity of the Net at any one moment is not infinite, and though it grows more quickly than the demand, it does at times get congested. It deals with this congestion equally – packets get transported on a first-come, first-served basis. Once packets leave one end, the network relays them on a best-efforts basis. If nodes on the network become overwhelmed, then packets passing across those nodes slow down.[28]

For certain applications, "best efforts" is not enough. Internet telephony, for example, doesn't do well when packets carrying voice get delayed. Any delay greater than 250 milliseconds essentially makes the system unusable.[29] And as content on the Net moves to real-time, bandwidth-demanding technology, this inability to guarantee quality of service becomes increasingly costly.

To deal with this problem, technologists have begun to propose changes to the architecture of the Net that might better enable some form of guaranteed service. These solutions generally pass under the title "Quality of Service" (QoS) solutions. These modifications would enable the network to treat different "classes" of data differently – video, for example, would get different treatment from e-mail; voice would get different treatment from the Web.

To enable this capacity to discriminate, the network would require more functionality than the original design allowed. At a minimum, the network would need to be able to decide what class of service a particular application should get and then treat the service accordingly. This in turn would make developing a new application more complex, as the programmer would need to consider the behavior of the network and enable the application to deal with that behavior.

The real danger, however, comes from the unintended consequences of these additional features – the ability of the network to then sell the feature that it will discriminate in favor of (and hence also against) certain kinds of content. As the marketing documents from major router manufacturers evince, a critical feature of QoS solutions will be their ability to enable the network owner to slow down a competitor's offerings while speeding up its own – like a television set with built-in static for ABC but a clear channel for CBS.

These dangers could be minimized depending upon the particular QoS technology chosen. Some QoS technologies, in other words, are more consistent with the principle of end-to-end than are others.[30] But proponents of these changes often overlook another relatively obvious solution – increasing capacity. That is, while these technologies will certainly add QoS to the Internet, if QoS technologies like the "RSVP" technology do so only at a significant cost, then perhaps increased capacity would be a cheaper social cost solution.

Put differently, a pricing system for allocating bandwidth solves certain problems, but if it is implemented contrary to end-to-end, it may well do more harm than good.

That is not to argue that it *will* do more harm than good. We don't know enough yet to know that. But it raises a fundamental issue that the scarcity mentality is likely to overlook: The best response to scarcity may not be a system of control. The best response may simply be to remove the scarcity.

This is the promise that conservative commentator George Gilder reports. The future, Gilder argues, is a world with "infinite" bandwidth.[31] Our picture of the Net now – of slow connections and fast machines – will soon flip. As copper is replaced with glass (as in fiber optics) and, more important, as electronic switches are replaced by optical switches, the speed of the network will approach the speed of light. The constraints that we know from the wires we now use will end, Gilder argues. And the end of scarcity, he argues, will transform all that we do.

There is skepticism about Gilder's claims about technology.[32] So, too, about his economics. The economist in all of us can't quite believe that any resource would fail to be constrained; the realist in all of us refuses to believe in Eden. But I'm willing to believe in the potential of essentially infinite bandwidth. And I am happy to imagine the scarcity-centric economist proven wrong.

The part I'm skeptical about is the happy progress toward a world where network owners simply provide neutral fat (or glass) pipe. This is not the trend now, and there is little to suggest it will be the trend later. As law professor Tim Wu wrote to me about Gilder's book:

> I think it is a "delta dollar sign" problem as we used to say in chemistry (to describe reactions that were possible, but not profitable). Private actors seem to only make money from infrastructure projects if built with the ability to exclude. . . . [H]ere in the industry, all the projects that are "hot" are networks with built-in techniques of exclusion and prioritization.[33]

Here *is* a tragedy of the commons. If the commons is the *innovation* commons that the protocols of the Net embrace, e2e most important among them, then the tragedy of that commons is the tendency of industry to add technologies to the network that undermine it. [. . .]

The Internet was born on a controlled physical layer; the code layer, constituted by the TCP/IP, was nonetheless free. These protocols expressed an end-to-end principle, and that principle effectively opened the space created by the computers linked to the Net for innovation and change. This open space was an important freedom, built upon a platform that was controlled. The freedom built an innovation commons. That commons, as do other commons, makes the controlled space more valuable.

Freedom thus enhanced the social value of the controlled: this is a lesson that will recur.

Notes

1 National Research Council, *The Internet's Coming of Age* (Washington, DC: National Academy Press, 2000), 30.
2 Telephone interview with David P. Reed (February 7, 2001), who contributed to the early design of the Internet protocols – TCP/IP – while a graduate student at MIT.
3 Lawrence Lessig, *Code and Other Laws of Cyberspace* (New York: Basic Books, 1999), 243, n. 19 (citing Kapor).
4 Network Working Group, "Request for Comments: 1958, Architectural Principles of the Internet," Brian E. Carpenter, ed. (1996), available at http://www.ietf.org/rfc/rfc 1958.txt.
5 Ibid.,2.1.
6 Ibid.
7 Tim Berners-Lee, *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web by Its Inventor* (San Francisco: HarperSanFrancisco, 1999), 99.
8 Telephone interview with David Reed.
9 Berners-Lee, 208.
10 National Research Council, 138.
11 Ibid., 107.
12 Ibid., 36–7.
13 Ibid., 37.
14 Douglas E. Comer, *Internetworking with TCP/IP*, 4th edn., vol. 1 (Upper Saddle River, NJ: Prentice-Hall, 2000), 691 (HTTP stands for "hypertext transfer protocol" and is "[t]he protocol used to transfer Web documents from a server to a browser"), 713 (TCP stands for "transmission control protocol"), and 694 (IP stands for "Internet protocol).")
15 Berners-Lee, 35.
16 See Robert M. Fano, "On the Social Role of Computer Communications," *Proceedings of the IEEE* 60 (September 1972), 1249.
17 Berners-Lee, 46. See also James Gillies, *How the Web Was Born: The Story of the World Wide Web* (Oxford and New York: Oxford University Press, 2000); Hafner and Lyon, *Internet Dreams: Archetypes, Myths, and Metaphors*, Mark J. Stefik and Vinton G. Cerf, eds. (Cambridge, Mass.: MIT Press, 1997).
18 Berners-Lee, 40 (describing Gopher and WAIS growing faster).
19 Ibid. (interconnect).
20 Ibid., 72–3.
21 Ibid.
22 Ibid., 74.

23 Ibid., 99.
24 John Naughton, *A Brief History of the Future* (London: Weidenfeld & Nicolson, 1999), 83–5.
25 Ibid., 84.
26 See, e.g., <http://www.asiapoint.net/insight/asia/countries/myanmar/my_spedev.htm>.
27 Steve Bickerstaff, "Shackles on the Giant: How the Federal Government Created Microsoft, Personal Computers, and the Internet," *Texas Law Review* 78 (1999), 25.
28 National Research Council, 130–1, n. 18 (describing best efforts as consequences of uniformity).
29 "Humans can tolerate about 250 msec of latency before it has a noticeable effect," <http://www.dialogic.com/solution/Internet/4070Web.htm>.
30 They are described in Mark Gaynor et al., "Theory of Service Architecture: How to Encourage Innovation of Services in Networks" (working paper, 2000, on file with author), 14.
31 See G. Gilder, *Telecosm: How Infinite Bandwidth Will Revolutionize Our World* (New York: Free Press, 2000), 158–64.
32 See, e.g., Bill Frezza, "Telecosmic Punditry: The World Through Gilder-Colored Glasses," *Internet Week* (December 4, 2000), 47; Rob Walker, "The Gildercosm," *Slate* magazine (September 11, 2000), available at <http://slate.msn.com/code/MoneyBox/MoneyBox.asp?Show=9/11/2000&idMessage=6030> (Gilder runs "against the current wisdom that sees bandwidth shortage as a problem"); Julian Dibbell, "From Here to Infinity," *Village Voice*, September 5, 2000, 65 ("It takes either profound sloth or transcendent faith to persist in voicing such breathless sentiments."). For a more favorable review, see, e.g., Blair Levin, "Review, *TELECOSM: How Infinite Bandwidth Will Revolutionize Our World,*" *Washington Monthly* (September 1, 2000), 54.
33 E-mail from Timothy Wu to Lawrence Lessig, February 16, 2001.