

Understanding Infrastructure: Dynamics, Tensions, and Design



Report of a Workshop on “History & Theory of Infrastructure: Lessons for New Scientific Cyberinfrastructures”

Paul N. Edwards
Steven J. Jackson
Geoffrey C. Bowker
Cory P. Knobel

January 2007




Table of Contents

EXECUTIVE SUMMARY	I
I. INTRODUCTION	1
BACKGROUND TO THE WORKSHOP	1
THE LONG NOW OF INFRASTRUCTURE	3
DEFINING CYBERINFRASTRUCTURE	5
BUILDING CYBERINFRASTRUCTURE.....	6
II. DYNAMICS	7
A HISTORICAL MODEL OF INFRASTRUCTURE DEVELOPMENT	8
SYSTEMS VS. NETWORKS AND WEBS.....	11
REVERSE SALIENTS	14
GATEWAYS.....	15
PATH DEPENDENCE	17
SCALE EFFECTS	19
APPLYING THE HISTORICAL MODEL TO CYBERINFRASTRUCTURE.....	21
III. TENSIONS	24
INTEREST AND EXCLUSION	24
OWNERSHIP AND INVESTMENT MODELS	29
DATA CULTURES, DATA TENSIONS	31
IV. DESIGN	33
STANDARDS AND FLEXIBILITY	36
V. CONCLUSIONS AND RECOMMENDATIONS	37
RECOMMENDATIONS.....	39
<i>Learning from cyberinfrastructure</i>	39
<i>Improving cyberinfrastructural practice</i>	41
<i>Enhancing resiliency, sustainability, and reach</i>	42
VI. BIBLIOGRAPHY	43
APPENDIX A. CONFERENCE PARTICIPANTS	48

Executive Summary

National Science Foundation support for scientific cyberinfrastructure dates to the 1960s. Since about 2000, however, efforts in cyberinfrastructure development have gathered momentum, guided by an increasingly comprehensive vision. Yet assembling the range of NSF-sponsored projects into a genuine infrastructure — highly reliable, widely accessible basic capabilities and services supporting the full range of scientific work — remains an elusive goal. Close study of other infrastructures, from railroads and electric power grids to telephone, cellular services, and the Internet, provides insights that can help guide and consolidate the NSF vision.

Since the 1980s, historians, sociologists, and information scientists have been studying how and why infrastructures form and evolve; how they work; and how they (sometimes) disintegrate or fail. In September 2006, a three-day NSF-funded workshop on “History and Theory of Infrastructure: Lessons for New Scientific Cyberinfrastructures” took place at the University of Michigan. Participants included experts in social and historical studies of infrastructure development, and domain scientists, information scientists, and NSF program officers involved in building, using, and funding cyberinfrastructure. The goal was to distill concepts, stories, metaphors, and parallels that might help realize the NSF vision for scientific cyberinfrastructure. This report summarizes the workshop findings, and outlines a research agenda for the future.

Social and historical analyses reveal some base-level tensions that complicate the work of infrastructural development. These include:

- **Time**, e.g. short-term funding decisions vs. the longer time scales over which infrastructures typically grow and take hold
- **Scale**, e.g. disconnects between global interoperability and local optimization
- **Agency**, e.g. navigating processes of planned vs. emergent change in complex and multiply-determined systems.

Such complications challenge simple notions of infrastructure building as a planned, orderly, and mechanical act. They also suggest that boundaries between technical and social solutions are mobile, in both directions: the path between the technological and the social is not static and there is no one correct mapping. Robust cyberinfrastructure will develop only when social, organizational, and cultural issues are resolved in tandem with the creation of technology-based services. Sustained and proactive attention to these concerns will be critical to long-term success.

Dynamics. Historical infrastructures – the automobile/gasoline/roadway system, electrical grids, railways, telephony, and most recently the Internet – become *ubiquitous, accessible, reliable, and transparent* as they mature. The initial stage in infrastructure formation is system-building, characterized by the deliberate and successful design of technology-based services. Next, technology transfer across domains and locations results in variations on the original design, as well as the emergence of competing systems. Infrastructures typically form only when these various systems merge, in a process of consolidation characterized by *gateways* that allow dissimilar systems to be linked into *networks*. In this phase, standardization and inter-organizational communication techniques are critical. As multiple systems assemble into networks, and

networks into webs or “internetworks,” early choices constrain the options available moving forward, creating what historical economists call “path dependence.”

Tensions. Transparent, reliable infrastructural services create vast benefits, but there are always losers as well as winners in infrastructure formation. Questions of ownership, management, control, and access are always present. For example:

- Who decides on rules and conventions for sharing, storing, and preserving data?
- Local variation vs. global standards: how do we resolve frictions between localized routines and cultures that stand in the way of effective collaboration?
- How can national cyberinfrastructure development move forward without compromising possibilities for international or even global infrastructure formation?

Design. These and other tensions inherent to infrastructure growth present imperatives to develop *navigation strategies* that recognize the likelihood of unforeseen (and potentially negative) path dependence and/or institutional or cultural barriers to adoption. Cyberinfrastructure seeks to enable a decentralized research environment that: 1) permits distributed collaboration; 2) provides incentives for participation at all levels; and 3) encourages the advancement of cross-boundary and interdisciplinary scholarship. Since all three of these goals are simultaneously social and organizational in nature *and* central to the technical base, designing effective navigation strategies will depend on strategic collaborations between social, domain, and information scientists. In particular, comparative studies of cyberinfrastructure projects can reveal key factors in success (and failure). Research on practices of standardization and modularity can help retain the openness, flexibility, and broad-scale usability of cyberinfrastructure, minimizing the path-dependent effects of standard-setting.

Recommendations: NSF should consider action in three broad areas.

- ***Learning from cyberinfrastructure.*** By applying well-understood evaluation tools, we can assess and compare existing cyberinfrastructure projects, both in the US and abroad. The resulting knowledge can be used to improve reporting mechanisms and incentive structures. Cyberinfrastructure projects can also be instrumented to collect social and organizational data.
- ***Improving cyberinfrastructural practice.*** Social science research can assist with NSF goals of training and enrolling professionals into the cyberinfrastructure-based research agenda. These goals may be achieved in part by improving diagnostics for current research environments, providing direct training for information managers, graduate students, and early-career faculty, and developing funding structures that support work on multiple time scales.
- ***Enhancing resilience, sustainability, and reach.*** Since infrastructures develop by creating links among varied systems, the NSF agenda may be promoted by forging and strengthening connections outside academic and governmental channels. Social scientists can help to recruit under-represented groups and institutions, as well as to create partnerships with organizations that have substantial existing expertise in areas complementary to scientific research, such as intellectual property standards and management.

I. Introduction

Background to the workshop

Academic scientists and funding agencies throughout the advanced industrialized world have recently embarked on major efforts to imagine, develop, and build new forms of “cyberinfrastructure” or “e-science”.¹ Separately and for several decades, historians and social scientists have studied the development of other kinds of infrastructure (railroads, waterworks, highways, telephony, business communication systems, the Internet, etc.). Reading across the body of this work produces two striking general results: first, there is a good deal of contingency, uncertainty, and historical specificity that attends any process of infrastructural development; second, despite these variations, there are shared patterns, processes, and emergent lessons that hold widely true across the comparative history and social study of infrastructure. This report represents a first attempt to bring these two fields of inquiry and practice together. In particular, it seeks to distill from the messy history and practice of infrastructure some general lessons and principles that might inform, guide, and in some cases caution the contemporary work of cyberinfrastructural development.

The report reflects the findings of the NSF-funded workshop, *History and Theory of Infrastructure: Lessons for New Scientific Cyberinfrastructures*. Hosted by the University of Michigan School of Information, and sponsored by the National Science Foundation’s Human and Social Dynamics Program, the Computer and Information Science and Engineering Directorate, and the Office of Cyberinfrastructure,² the workshop brought together more than thirty historians, social scientists, domain scientists, and cyberinfrastructure developers for three days of open, focused, interdisciplinary discussion around the patterns, perils, and possibilities of infrastructure (both cyber and other).³

The group was charged with three general tasks: first, to identify dynamics, tensions,

¹ While we use the NSF term “cyberinfrastructure” throughout this report, similar arguments can be made for the UK “e-science” program and other efforts to develop new computational infrastructure in support of innovative and collaborative science.

² NSF grant #0630263. We thank each of these entities for their generous support.

³ We thank the workshop participants for their many and invaluable contributions. In particular, Tineke Egyedi, Cal Lee, Erik van der Vleuten, and JoAnne Yates composed “think pieces” which we used as the basis for workshop sessions (and for parts of this report). These individuals, as well as Johan Schot, Jane Summerton, and Fran Berman, also advised us during a planning session held at the NSF on June 26, 2006. University of Michigan doctoral students Clapperton Mavhunga, Yong-Mi Kim, Charles Kaylor, and Trond Jacobsen served as rapporteurs and technical assistants, under the capable direction of Cory Knobel, who also organized the “live” workshop blog. Doctoral student Archer Batcheller prepared the workshop website (www.si.umich.edu/InfrastructureWorkshop/agenda.php). For a full list of all workshop participants including disciplinary and institutional affiliations, see Appendix A.

strategies, and design challenges that are common across the wider history and contemporary practice of infrastructural development; second, to begin to distill from this collected experience concrete lessons and principles that might shape and inform the activities indicated under the National Science Foundation's Vision for Cyberinfrastructure; and third, to propose a research agenda for cyberinfrastructure studies.⁴ Our work drew upon and benefited greatly from the valuable series of "Cyberinfrastructure for..." reports addressing applications of CI in various domains, but this workshop concerned a different question: rather than "What can cyberinfrastructure bring to the social sciences and humanities?" we asked "What can the findings and methods of social and historical analysis bring to the development of cyberinfrastructure?"⁵

The report that follows delivers the workshop's key findings in five sections. The Introduction provides a general overview and orientation to concepts of infrastructure, along with a brief overview of cyberinfrastructure and the NSF's cyberinfrastructure program. The Dynamics section surveys questions relating to the genesis, development, and scaling of infrastructure, and includes examples, patterns, and principal findings from historical studies of infrastructure. The Tensions section describes infrastructure as fundamentally contested, and samples the kinds of conflicts that developing infrastructures have frequently encountered. It pays particular attention to scientific data and data cultures as both focal objects and stumbling blocks in cyberinfrastructure development. The Design Strategies section explores a fundamental contradiction: if effective infrastructures are rarely "built" in an entirely top-down, orderly, and blueprint-like way (as we shall argue), can we nevertheless think of design as a reasonable and important aspiration for would-be infrastructure developers? What might count as legitimate and promising design strategies? The Conclusions and Recommendations section collects the broader findings and translates them into concrete recommendations for those charged with shaping and implementing the NSF cyberinfrastructure vision.

Beyond the workshop itself, the report signals the first fruits of what we hope will be an ongoing and mutually beneficial collaboration between those with expertise in the social and historical analysis of infrastructure and those tasked with developing it in contemporary settings. As the workshop resoundingly demonstrated, a good deal rides on the front-end phases of infrastructure development. Building a robust, empirical, and broad-based analytic capacity to support cyberinfrastructure development should be an NSF priority of the highest order.

⁴ See, e.g. the NSF Cyberinfrastructure Council's *NSF's Cyberinfrastructure Vision for 21st Century Discovery* (Ver 7.1; July 20, 2006), available at www.nsf.gov/od/oci/ci-v7.pdf. For an important earlier account of this vision, see the *Report of the Blue-Ribbon Advisory Panel on Cyberinfrastructure* ("the Atkins Report"), available at www.nsf.gov/od/oci/reports/toc.jsp.

⁵ For answers to the first question, see, *inter alia*, *Report of the Commission on Cyberinfrastructure for the Humanities and Social Sciences*, and the *Final Report of NSF SBE-CISE Workshop on Cyberinfrastructure and the Social Sciences*. These and other domain-specific cyberinfrastructure reports can be accessed through the NSF Office of Cyberinfrastructure homepage, at: <http://www.nsf.gov/od/oci/reports.jsp>.

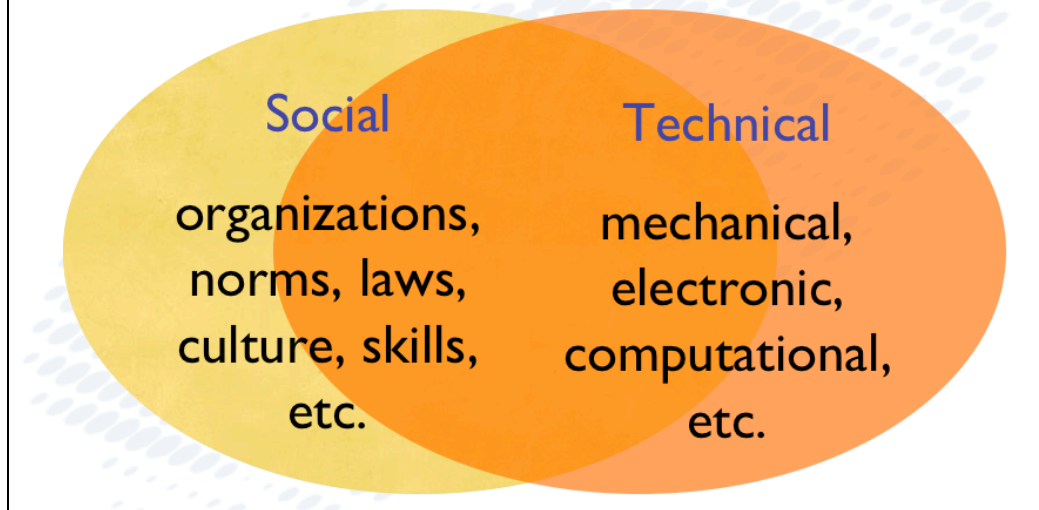
The long now of infrastructure

Stewart Brand's "clock of the long now" will chime once every millennium: a cuckoo will pop out (Brand 1999). Accustomed as we are to the "information revolution," the accelerating pace of the "24/7" lifestyle, and the multi-connectivity provided by the World Wide Web, we rarely step back and ask what changes have been occurring at a slower pace, in the background. For the development of cyberinfrastructure, the long now is about 200 years. This is when two suites of changes began to occur in the organization of knowledge and the academy which have accompanied – slowly – the rise of an information infrastructure to support them: an exponential increase in information gathering activities by the state (statistics) and knowledge workers (the encyclopedists) on the one hand and the accompanying development of technologies and organizational practices to sort, sift and store information.

When dealing with infrastructures, we need to look to the whole array of organizational forms, practices, and institutions which accompany, make possible, and inflect the development of new technology. JoAnne Yates made this point beautifully in describing the first commercial use of punch card data tabulators, in the insurance industry. That use became possible because of *organizational* changes within the industry. Without new forms of information management, heralded by such low status technologies as the manila folder and carbon paper, accompanied by new organizational forms, there would have been no niche for punch card readers to occupy (Yates 1989). Similarly, Manuel Castells argued that the roots of contemporary "network society" are new organizational forms created in support of large corporate organizations, which long predate the arrival of computerization (Castells 1996). James Beniger described the entire period from the first Industrial Revolution to the present as an ongoing "control revolution" in which societies responded to mass production, distribution, and consumption with both technological and organizational changes, designed to manage ever-increasing flows of goods, services, and information (Beniger 1986). In general there is more continuity than cleavage in the relationship of contemporary "information society" to the past (Chandler and Cortada 2003).

The lesson of all these studies is that organizations are (in part) information processors. People, routines, forms, and classification systems are as integral to information handling as computers, Ethernet cables, and Web protocols. The boundary between technological and organizational means of information processing is mobile. It can be shifted in either direction, and technological mechanisms can only substitute for human and organizational ones when the latter are prepared to support the substitution.

Boundaries between social and technical action can often be shifted *in either direction*



In the “long now,” two key facets of scientific information infrastructures stand out. One clusters around the nature of work in the social and natural sciences. Scientific disciplines were formed in the early 1800s, a time Michel Serres felicitously describes as the era of x-ology, where “x” was “geo,” “socio,” “bio” and so forth (Serres 1990). Auguste Comte classified the division of labor in the sciences, placing mathematics and physics as the most developed and best models, and sociology as the most complex and least developed, more or less where Norbert Wiener placed them 130 years later in *Cybernetics and Society* (Wiener 1951). This was also the period during which the object we now call the database came to be the lynchpin of the natural and social sciences. Statistics etymologically refers to “state-istics,” or the quantitative study of societies (states); it arose along with censuses, medical records, climatology, and other increasingly powerful techniques for monitoring population composition and health (Porter 1986). Equally, the natural sciences – moved by the spirit of the encyclopedists – began creating vast repositories of data. Such repositories were housed in individual institutions, such as botanical gardens and museums of natural history. Today they are increasingly held in electronic form, and this is fast becoming the norm rather than the exception. For example, the Ecological Society of America publishes digital supplements, including databases and source code for simulation models, for articles published in its journals (www.esapubs.org/archive/), and a researcher publishing a protein sequence must also publish his or her data in the (now worldwide) Protein Data Bank.

The second facet clusters around scientists’ communication patterns. In the 17th and 18th centuries scientists were largely “men of letters” who exchanged both public and private correspondence, such as the famous Leibniz/Clarke exchange. From the early 19th century a complex structure of national and international conferences and publishing practices developed, including especially the peer-reviewed scientific journal. Communication among an ever-broader scientific community was no longer two-way, but *n*-way. New forms of transportation undergirded the development of a truly international scientific community aided also by *linguae francae*, principally English and French.

New scientific cyberinfrastructures must be understood as an outgrowth of these developments. Databases and n -way communication among scientists have developed embedded in organizational and institutional practices and norms. There is far more continuity than many recognize. However, as scientific infrastructure goes cyber, there is also genuine discontinuity. The social and natural sciences grew up together with communication and data-processing technology. Changes in these latter will have ripple effects throughout the complex web of relations that constitutes scientific activity.

Defining Cyberinfrastructure

Most often, cyberinfrastructure is defined by jotting down a laundry list. The reference Atkins Report for the National Science Foundation defines it as those layers that sit between base technology (a computer science concern) and discipline-specific science. The focus is on value-added systems and services that can be widely shared across scientific domains, both supporting and enabling large increases in multi- and interdisciplinary science while reducing duplication of effort and resources. According to the Atkins Report, cyberinfrastructure consists of “hardware, software, personnel, services and organizations” (p. 13). This list recognizes from the outset that infrastructure is about more than just pipes and machines. The more recent cyberinfrastructure vision document is similarly diffuse, though it regrettably somewhat sidelines the social and organizational in the definition:

Cyberinfrastructure integrates hardware for computing, data and networks, digitally enabled sensors, observatories and experimental facilities, and an interoperable suite of software and middleware services and tools. Investments in interdisciplinary teams and cyberinfrastructure professionals with expertise in algorithm development, system operations, and applications development are also essential to exploit the full power of cyberinfrastructure to create, disseminate, and preserve scientific data, information, and knowledge (*NSF CI Vision ver. 7.1*, p. 6).

Both these definitions do, however, draw attention to the dynamic, complex nature of cyberinfrastructure development.

While accepting this broad characterization, this report’s long-now perspective invites a discussion of first principles. For this we return to Star and Ruhleder’s now classic definition of infrastructure (Star and Ruhleder 1996), originally composed for a paper on one of the early scientific collaboratories, the Worm Community System. Here we show how their definitions can be ordered along two axes, the social/technical and the local/global:




Figure 1. Cyberinfrastructure as distributions along technical/social & global/local axes (diagram courtesy of Florence Millerand).

In building cyberinfrastructure, the key question is not whether this is a “social” problem or a “technical” one. That is putting it the wrong way around. The question is whether we choose, for any given problem, a social or a technical solution, or some combination. It is the *distribution* of solutions that is the object of study. An everyday example comes from the problem of email security. How do I distribute my trust? I can delegate it to my machine, and use Pretty Good Encryption for all my email messages. Or I can work socially and organizationally to make certain that sysops, the government, and others who might have access to my email internalize a value of my right to privacy. Or I can change my own beliefs about the need for privacy – arguably a necessity with the new infrastructure.

For our purposes, cyberinfrastructure is the set of organizational practices, technical infrastructure and social norms that collectively provide for the smooth operation of scientific work at a distance. All three are objects of design and engineering; a cyberinfrastructure will fail if any one is ignored.

Building Cyberinfrastructure

At first glance, the term “building” seems apposite. After all, infrastructures are composed of interoperating systems, each of which had a builder. But complex structures have different types of builders and are not always the result of intentional planning (Dennett 1996). As we will see in the Dynamics section of this report, the eventual growth of complex infrastructure and the forms it takes are the result of converging histories, path dependencies, serendipity, innovation, and “bricolage”

(tinkering). Speaking of cyberinfrastructure as a machine to be built or a technical system to be designed tends to downplay the importance of social, institutional, organizational, legal, cultural, and other non-technical problems developers always face. Axelrod and Cohen's idea of *harnessing complexity* cautions against seeking tight control over technologically-enabled organizational structures; even if it were a good idea, it simply wouldn't work (Axelrod and Cohen 2001). By extension, the organizational aspects of science and the role of the social sciences in cyberinfrastructure should be integrated into the work of design. Here, one of Star and Ruhleder's observations is key:

Infrastructure is fixed in modular increments, not all at once or globally.
Because infrastructure is big, layered, and complex, and because it means different things locally, it is never changed from above. Changes take time and negotiation, and adjustment with other aspects of the systems involved.

Hence this report turns away from a language of design and engineering, reframing the discussion in a more organic lexicon. Since infrastructures are incremental and modular, they are always constructed in many places (the local), combined and recombined (the modular), and they take on new meaning in both different times and spaces (the contextual). Better, then, to deploy a vocabulary of "growing," "fostering," or "encouraging" in the evolutionary sense when analyzing cyberinfrastructure.

Adopting an alternative framework for cyberinfrastructure analysis frees us from presuming that the final working form of scientific cyberinfrastructure will look much like the initial vision. Further, this framework is responsive to the findings of science studies that science, theory, and inquiry are created locally, and build out from these local contexts. As cyberinfrastructure grows and takes shape by drawing in new communities, each with its distinctive histories, needs, and practices, we can expect a common sense and (partially) shared understandings of cyberinfrastructure to emerge. Such processes may be aided by the crafting of a shared functional lexicon or "pattern language" (Alexander 1979) for cyberinfrastructure.

II. Dynamics

This section outlines an historical model of infrastructure development, one which has been repeatedly confirmed across numerous cases from the 19th century to the present.

For cyberinfrastructure projects, this model leads to three significant conclusions. First, true infrastructures only begin to form when locally constructed, centrally controlled *systems* are linked into *networks and internetworks* governed by distributed control and coordination processes. Second, infrastructure formation typically starts with *technology transfer* from one location or domain to another; adapting a system to new conditions introduces technical variations as well as social, cultural, organization, legal, and financial adjustment. Third, infrastructures are consolidated by means of *gateways* that permit the linking of heterogeneous systems into networks.

The section then turns to three key dynamics of infrastructure development. *Reverse salients* — critical unsolved problems — may be technical, but are also frequently social or organizational in nature, particularly in the network/internetwork formation phase.

Gateways are defined as technologies and standards applied across multiple communities of practice. The transition from systems to networked infrastructures requires generic and meta-generic gateways, as opposed to the dedicated or improvised gateways used in systems. Third, as infrastructures grow they create *path dependence*; as organizations and individuals come to rely on an infrastructure, they adapt to it, coupling many small-scale and local elements to the larger commodity service. This phenomenon has positive and negative aspects.

These concepts explain why it is difficult to alter infrastructures once they have become established, and thus why choices in the early phases of development — as in the case of cyberinfrastructure today — really make a difference. The section ends by comparing the historical trajectories of electric power, computing, and cyberinfrastructure.

A historical model of infrastructure development

The time scale in historical studies of infrastructural change is decades to centuries — considerably longer than most research projects in cyberinfrastructure! Historians' principal model of infrastructure development draws on Thomas Parke Hughes's *Networks of Power* (1983), on the evolution of electric power. Hughes's model was adapted and extended over two decades by a loose-knit group of historians and sociologists studying "large technical systems" (LTS), including telephone, railroads, air traffic control, and other major infrastructures (Bijker and Law, 1992; Braun and Joerges, 1994; Coutard, 1999; Coutard et al., 2004; La Porte, 1991; Mayntz and Hughes, 1988; Bijker et al., 1987; Kaijser et al., 1995; Summerton, 1994). The Hughes model conceptualizes invention and innovation in terms of systems rather than isolated devices.

System builders and systems. System builders create and promote *systems*, i.e. linked sets of devices that fill a functional need. Hughes's paradigmatic example of a system builder is Thomas Edison. Other inventors had already hit upon light bulbs; what set Edison apart was his conception of a lighting system including generators, cables, and light bulbs. The system delivered a service (lighting), rather than a commodity (electricity) or an isolated device (the light bulb). Similarly, digital computing did not achieve commercial success until manufacturers such as Univac and IBM supplied not just CPUs, but complete data processing systems, including mass storage (magnetic tape and disks), input devices (keyboards, punch cards), and output devices (printers, card punches). They also rapidly found that technical systems alone were insufficient; they had to supply training, software, and other kinds of support as well. Historians concur that IBM's rise to dominance in the late 1950s was based as much on the services it supplied to customers as on the technical features of its products; it also built on its large installed base of punch card and other equipment. Indeed, IBM's research group at Almaden is currently trying to establish the discipline of service science.

Successful system building always includes organizational, financial, legal, and marketing elements. Historians have noted the common phenomenon of system-builder teams made up of one or more technical "wizards" or "supertechs," who handle system conception and innovation, working together with a "maestro," who orchestrates the organizational, financial, and marketing aspects of the new system. Such teams can also include a charismatic "champion" who stimulates external interest in the project, promoting it against competing systems and generating widespread adoption (McKenney et al., 1995). Well-known examples of such teams in information and

communication infrastructure are Alexander Graham Bell and Theodore Vail (AT&T); Thomas Watson Sr. and James Bryce (IBM); Robert Taylor, Lawrence Roberts, Robert Kahn and Vint Cerf (ARPANET/Internet); Steve Jobs and Steve Wozniak (Apple); Bill Gates, Paul Allen, and Steve Ballmer (Microsoft); and Tim Berners-Lee and Robert Cailliau (World Wide Web).

“Wizard,” “maestro,” and “leader” label roles, not people; they may be held by individuals, groups, or organizations, as well as in various combinations. Our emphasis here is not on heroic individuals — whose powers and importance are almost always exaggerated — but on the social features of this pattern. First, system building typically begins as a social act (even a dyad is a social system). Second, the wizard-maestro-leader combination reflects the spectrum of crucial capabilities: technical, organizational, and social.

Government agencies have sometimes played key roles in the system-building phase of major infrastructures. During and after World War II, for example, the principal sources of support for US digital computing research were military agencies, especially the Office of Naval Research and the Air Force. Very large contracts for the SAGE air defense system helped IBM take the lead in the American computer industry (Edwards, 1996). The government has the ability to plan for the long term; the Dutch government in the sixteenth century, for example, planned forestry growth over the subsequent two hundred years as part of its naval construction infrastructure. Similarly, government has the ability to shepherd research projects over long periods of time – as witness the successful creation of the Internet.

Technology transfer and growth. Once an LTS has been successfully constructed in one location, *technology transfer* to other locations (organizations, cities, nations) follows. Because conditions at the new locations differ, this process always produces variations on the original system design as well as new organizational support. This adaptation leads to a phenomenon Hughes called “technological style”: the distinctive look and feel of the “same” technical system as it appears in differing local and national contexts. As it develops, a new LTS not only requires further technical innovation, but also increasingly incorporates heterogeneous components. Finance capital, legal representation, and political and regulatory relationship management become indispensable elements of the total system. Relevant economic forces include economies of scale and scope, and economies of reach (Kaijser, 2003).

As the LTS spreads from place to place, competing systems may be introduced with dissimilar, frequently incompatible properties. In the early days of electric power, for example, competition occurred among dozens of systems using different line voltages, as well as both direct and alternating current, all with their advantages and defects.

Issues of *scaling* become crucial during the growth phase. Systems that worked well in a small local area, with a few hundred users, typically require substantial redesign in order to function in many places with thousands or millions of users. Concerns such as technical support, billing, capital investment, management of user expectations, marketing, and many other issues come to the forefront during this phase.

During growth, attention to *users and user communities* can become critical to success or failure. A key problem is that the development process builds expertise among the developers; as a result, developers can lose their ability to see how novices, or users in

a different field, perceive and use their system. Information technology projects frequently founder when they attempt to transition rapidly from a small, close-knit developer community to a larger, more diverse community of novice users, such as consumers or scientists in unrelated fields. Rapid growth can make this especially difficult to manage; users often take the process into their own hands, leading to divergent norms, practices, and standards implementation (Fischer, 1992; Hanseth and Monteiro, 1998; Kahin and Abbate, 1995; Abbate, 1999).

Consolidation: network formation. In *consolidation*, the final stage of the LTS model, competition among technological systems and standards is resolved in one of two ways. In rare cases, one system wins total victory over the others. More often, developers create *gateways* that allow previously incompatible systems to interoperate. The rotary converter, for example, allowed AC power to be converted to DC on a large scale, permitting competing electrical distribution systems to be connected (David and Bunn, 1988). Today, gateways such as AC/DC power converters for consumer electronics and telephone adapters for international travel. Platform-independent software, languages, and presentation formats such as Java, HTML, and PDF are information technology examples. By allowing heterogeneous technical systems to interoperate, gateways (see below) permit the creation of *networks* such as power grids, railroad, telephone, and the ARPANET, NSFNET, and Internet.

As in the system-building phase, the goal of network formation is to deliver a service. For example, distributed packet switching computer networks were first developed in the mid-1960s at the UK's National Physical Laboratory (NPL) and at various ARPA research contractors, especially Bolt Beranek and Newman, in the United States. In each case, the critical step was not the technical development of packet switching itself, but the conception of the network as a way to share data and programs among expensive mainframe computers under a timesharing regime (Abbate, 1999; Hafner, 1996; Hauben and Hauben, 1997). The network delivered not only a physical connection or a communication technique (packet switching), but a service (shared programs and data).

The consolidation phase can be seen as complete when the service in question has become, from the point of view of users, a commodity resource, i.e. an undifferentiated good such as electricity, telephone switching, or IP connectivity. The “computer utilities” envisioned in the late 1960s — giant timesharing computers that would provide “powerful and reliable systems capable of serving large communities” (Fano, 1992, [original 1967] p. 39) — represent a vision of computing as a commodity infrastructure. Ultimately this was supplied not through systems (timesharing), but through the Internet and grid computing.

Recently, historians have begun attending to another aspect of consolidation: the role of infrastructure in transnational linking. Transborder bridges and tunnels; power grids; national telegraph and telephone systems; containerized international shipping and road/rail transport; airports; the Internet and WWW; and many other infrastructure projects involve resolution of political, legal, and financial issues simultaneously with technical standards. At the same time, they alter the nature of national boundaries, especially at the level of culture and national sovereignty (Held et al., 1999; Schot et al., 2006; Vleuten et al., 2006). Delinking occurs, particularly in wartime; transnational infrastructural links are usually among the first objects of military engagement. Clearly, scientific cyberinfrastructure will play a role in transnational linking as well. Coordinating

now with cyberinfrastructure projects worldwide may reduce the difficulty of consolidation, if and when it occurs.

Governments have played important roles in the consolidation phase of infrastructure development. The general approach during the 1850-1975 period has been called the “*modern infrastructural ideal*”: universal service by a single provider (Graham and Marvin, 2001). Following this logic, many national governments provided most or all of these services; national rail, road, electric power, and PT&T (post, telephone & telegraph) networks are the prime examples. The United States kept some of these services private, but in many cases allowed formation of monopoly providers or “public utilities” under oversight by regulatory agencies. Telephone, electric power, sewer, water, and natural gas are examples. (Friedlander, 1995a; Friedlander, 1995b; Friedlander, 1996). In other cases, such as rail and air transport, competition continued, but within a legal framework of public oversight. In its early phases the Internet was created with ARPA and NSF funding, and it was used principally for military, research, and educational activities; these features allowed its government sponsors to prohibit commercial activity until the Internet was privatized in the early 1990s. Today, the Internet and WWW are subject to national regulations of many kinds.

Splintering of the “modern infrastructural ideal.” Although early versions of the LTS model ended with consolidation, there is a further phase. Starting around 1975, in the United States, the UK, and to lesser extent elsewhere, the model of monopoly utilities was increasingly displaced by a deregulated, market-oriented approach, with reduced but still significant public oversight (as in air transport, telephone, television, and energy services). By increasing the ability of multiple suppliers to coordinate their operations, balance loads, and handle system breakdowns, new information technologies played a major role in this ongoing transition. Increased capacity for decentralized coordination (as opposed to centralized control) enabled a retreat from the logic of vertical integration.

The result in most infrastructure areas has been a pronounced *splintering* of the single-provider, monopoly utility model (Graham and Marvin, 2001). Frequently this has also meant service tiering, with wealthy customers and heavy users receiving premium, highly reliable services, while poor people and infrequent users must rely on low-grade services or be excluded altogether – an issue being revisited today under the broad rubric of internet neutrality, one which is highly relevant to the social goals of producing scientific cyberinfrastructure for disadvantaged communities.

Having emerged in an era of ideological opposition to large new government-funded projects, the cyberinfrastructure movement has sought new models in the success of open-source software development projects, especially Mozilla Firefox with the highly refined “Bugzilla” bug reporting system and Linux. The RFC process in the development of the Internet was a major precursor to these (Russell, 2006). Building the Internet involved massive long term investment by ARPA, the NSF and other government agencies – frequently, as the Hughes report to the NAS noted, involving making long term bets on particular key players irrespective of short-term pay-off (the continuing CNRS funding model).

Systems vs. networks and webs

The growth, consolidation, and splintering phases of the historical model mark a key transition from homogeneous, centrally controlled, often geographically local systems to

heterogeneous, widely distributed networks in which central control may be partially or wholly replaced by coordination.

In general, and specifically in the meaning of the cyberinfrastructure framework, *infrastructures are not systems*. Instead, they are networks or webs that enable locally controlled and maintained systems to interoperate more or less seamlessly. It is typically only in the consolidation phase, with the appearance of standardized, generic gateways, that most LTSs become genuine infrastructures, i.e. ubiquitous, reliable, and widely shared resources operating on national and transnational, scales. Thus we define a spectrum running from systems (centrally organized and controlled) to networks (linked systems, with control partially or wholly distributed among the nodes) to webs (networks of networks based primarily on coordination rather than control). Table 1 summarizes this distinction.

	Systems	Infrastructures	
		Networks	Internetworks or Webs
Key actors	System builders Users (adjustment roles)	Gateway builders Standards bodies Corporations & governments Users (transformative roles)	Gateway builders Standards bodies Corporations & governments Users (foundational roles)
Elements	Heterogeneous components and subsystems	Heterogeneous systems	Heterogeneous networks
Gateways and standards	Dedicated or improvised	Generic or meta-generic	Generic or meta-generic
Control vs. coordination	Control Central, strong	Control and coordination Partially distributed, moderate strength	Coordination Widely distributed, weak Reliant on other infrastructures
Boundaries	Closed, stable	Open, reconfigurable	Open, reconfigurable Virtual or second-order large technical systems
Examples	Local electric power company Enterprise computing (e.g. banks, insurance companies)	Railroad; electric power grids Grid computing NEES, GEON National weather services	Intermodal freight Global telephone system (fixed + mobile + VOIP) Internet and WWW World Weather Watch

Table 1. Systems vs. infrastructures (modified from Edwards, 1998a).

The last column in Table 1, “Internetworks or webs,” refers to integration *across* networks. Perhaps the best example is intermodal freight, in which ISO standard containers, which may be mounted on standard truck or rail wheelbases and lifted into container ship holds, smooth transfers among independent road, rail, and shipping infrastructures. In information technology, the Internet and WWW are obvious examples. However, integrated information internetworks long predate the Internet. Telegraph, telephone, and postal mail were linked into a 19th century “analog information internetwork” Another pre-Internet example is the World Weather Watch, which collects data from national weather services, sends it to global data processing centers, and returns processed global forecasts and data to the national services for their own use (Edwards, 2006).

The analog information internetwork

Greg Downey has described how 19th century business users effectively combined the available communication systems into a single information internetwork, a “...combination of character-transmission telegraph, voice-transmission telephone, and physical-transport Post Office networks. I call this an ‘analog’ internetwork because... information could only move over each component network in a single form, requiring repeated physical translations as it moved through the internetwork (handwriting to voice to dot-and-dash and back again). Although the telegraph itself was in some sense ‘digital’ — based as it was on three possible states: no pulse, a short pulse (dot), and a longer pulse (dash) — those states were conveyed at varying cadences through the physical actions of rapidly pressing telegraph keys and attentively listening to telegraph sounders, and so were still analog at the core. ...

Historical actors who used and studied the telegraph, telephone, and Post Office saw the three as an internetwork. Business texts from the 1910s through the 1930s instructed students that proper business practice when sending telegrams involved all three media: even when paying for the ‘report delivery’ and ‘repeat back’ options to make sure telegrams were accurately transmitted and received (with those reports coming by telephone), important telegrams were to be ‘confirmed immediately by a properly dated and signed letter’” (Downey, 2001:213-14).

The moments at which systems become linked into networks, and networks become linked to form internetworks, thus represent crucial phases in infrastructure formation.

Virtual infrastructures and second-order LTSs: email, WWW, cellular telephony.

Better information technology increases the capacity for distributing control from a central point to the nodes in a network. It also permits integration of services across networks. The contemporary phrase for this capability is “digital convergence.” As we have seen, this is in fact an old trend, not necessarily dependent on computer technology. However, the flexibility and power of computers has undeniably been the principal reason for the explosion of new basic services built upon existing infrastructures in recent decades. Analysts have named these “virtual infrastructures”

and “second-order large technical systems” (Braun, 1994; Edwards, 1998b). Cyberinfrastructure for the sciences is a specialized manifestation of this trend.

The principal historical examples of successful large-scale infrastructure formation since the 1970s are email, the World Wide Web, and cellular telephony. Since both email and the Web “sit on top of” the Internet, they are also outstanding examples of virtual infrastructures. Cellular telephony, requiring very large investments in cell towers and transmitters, followed a trajectory more like that of electric power or rail. All three cases, however, depended strongly on pre-existing infrastructures. However, its integration with the existing land-line telephone network makes it also a second-order LTS, combining existing services in new ways based largely on information technology. The WWW — often touted as a miracle of decentralized grassroots development — would not exist without the Internet, whose early history strongly resembles those of other infrastructures. Meanwhile, cellular telephony’s success is strongly linked to its integration with the pre-existing land-line telephone network, which provided a huge head start toward building a critical mass. (Even with this base to build on, however, cellular telephony’s growth curve resembles those of the other infrastructures charted in Figure 3, below; the first cellular phone call was made in 1973.) The explosive growth experienced by all three of these recent infrastructures would not have occurred without the slower growth of the older infrastructures that underlie them.

Together with the splintering phenomenon described above, our increasing capacity to build virtual infrastructures and second-order large technical systems using coordination mechanisms is important background to cyberinfrastructure formation. These phenomena have created a paradigm of increasingly articulated, fragmented and tiered service delivery, across all infrastructures. The decline of the monopoly public utility model suggests that forming new large-scale infrastructures may be more difficult than in the past. Indeed, NSF reports on cyberinfrastructure reflect a distaste for large-scale, long-term projects by single providers, and a corresponding enthusiasm for open source models and “federated” systems linked by coordination rather than control.

Reverse salients

The LTS tradition highlights “reverse salients” in system development. This is a military metaphor referring to points where an advancing front is held back. In terms of technological change, the analogy refers to engineering problems whose solution is required for the entire system to work or to grow. Examples are long-distance transmission of electric power; automatic switching of telephone calls; or linking computer networks which use different data packet sizes and addressing schemes. Since reverse salients are often widely understood to be the most significant problems in a field, they are normally a locus of intense research efforts. Typically multiple groups converge on solutions around the same time.

Reverse salients need not be technical; in fact, the most important reverse salients are often legal, political, social, or cultural. Government and other national-level institutions have played critical roles in identifying, and sometimes also in shifting, reverse salients in the sciences. For example, in its 1992 report *Computing the Future*, the National Research Council criticized the discipline of computer science for a narrow agenda that failed to engage applications areas and interdisciplinary arenas such as human-computer interaction (Hartmanis and Lin, 1992). Arguably the current NSF cyberinfrastructure initiatives represent a continuation of this effort, this time focused on

changing the culture and social relations of both computer science and the domain sciences to reduce duplication of effort while creating basic middleware services on which present and future inter- and multi-disciplinary research can rely.

Overcoming social reverse salients in computer networking

In the late 1960s era of expensive mainframe computers, jealously guarded by system operators and research groups, ARPA's idea of sharing programs and data across a network with researchers located elsewhere seemed like an invasion, particularly since such sharing often required assistance from local operators.

In an interview, Lawrence Roberts recalled how ARPA compelled its contractors not only to connect to, but also to use, the early ARPANET. "The universities were being funded by us, and we said, "We are going to build a network and you are going to participate in it. And you are going to connect it to your machines. By virtue of that we are going to reduce our computing demands on the office. So that you understand, we are not going to buy you new computers until you have used up all of the resources of the network.' So over time we started forcing them to be involved" (quoted in Abbate 1999, 55).

In a similar but less dramatic way, the NSF also compelled participation in the NSFNET by requiring its supercomputer centers to make network connections available to all qualified educational or research users. As in the case of the ARPANET, having provided very costly equipment, the funding agency was in a position to set conditions that strongly promoted broad-based, inexpensive access.

Examples of reverse salients relevant to cyberinfrastructure include: generating metadata (this is an unfunded mandate); the tangle of intellectual property rights; techniques for federating databases held at multiple institutions using different equipment and data formats; domain specific data sharing and publication cultures; reluctance of modelers who have been working with a given program to shift to a better one if the learning curve is too steep; lack on incentives in universities for infrastructure-building and data sharing work; inability to translate across different fields and so forth.

Gateways

As Tineke Egyedi has observed, gateway technologies confer differing degrees of flexibility on technical systems depending on the degree to which they are standardized. Gateways may be *dedicated* (improvised, or devised specifically for a particular system); *generic* (standardized sockets opening one system to interconnection with others); or *meta-generic* ("modeled," i.e. specifying a framework or protocol for the creation of specific generic standards, without specifying those standards directly). Table 2 outlines Egyedi's framework.

Degree of Standardization	Scope of Gateway Solution	Examples
High (modeled)	Meta-generic	OSI ⁶
Medium (standardized)	Generic	XML, Java, ISO container ⁷
Low ('improvised')	Dedicated	AC/DC rotary converter

Table 2. Relationship between degree of standardization and scope of gateway solution (from Egyedi, 2001).

Plug adapters (e.g. 3-pin to 2-pin AC power, Firewire 400/800) and AC/DC power converters are excellent everyday examples. Gateway technologies of all three types are manifested in software as well, e.g. document format converters, allow one document format to be converted into another; one operating system to emulate the properties of another; and so on.

Gateways represent a key principle of infrastructure development: **plugs and sockets** that allow new systems to be joined to an existing framework easily and with minimal constraint. Gateways are often wrongly understood as “technologies,” i.e. hardware or software alone. A more accurate approach conceives them as combining a technical solution with a social choice, i.e. a standard, both of which must be integrated into existing users’ communities of practice. Because of this, gateways rarely perform perfectly.

“Information technology standards have been touted as a means to interoperability and software portability, but they are more easily lauded than built or followed. Users say they want low-cost, easily maintained, plug-and-play, interoperable systems, yet each user community has specific needs and few of them want to discard their existing systems. Every vendor wants to sell its own architecture and turbo-charged features, and each architecture assumes different views of a particular domain (e.g., business forms, images, databases). International standards founder on variations in culture and assumptions — for example, whether telephone companies are monopolies — in North America, Europe, and Asia” (Libicki, 1995:35).

⁶ The Open Systems Interconnection Reference Model defines seven “layers” of computer network function, from physical links to applications. Within each layer, standards can evolve separately so long as they conform to the model (see Abbate, 1999, Chapter 5).

⁷ XML is the eXtensible Markup Language. Java is a cross-platform computer language. ISO (International Standards Organization) container refers to standard sizes, shapes, and connectors for shipping containers used for freight transport by ship, rail, and truck.

“Below the level of the work.” Neither the exact implementation of standards, nor their integration into local communities of practice, can ever be wholly anticipated. For this reason, gateways in information infrastructures work best when they interlock with the existing framework “below the level of the work,” i.e. without specifying exactly how work is to be done or exactly how information is to be processed (Forster and King, 1995). Most systems that attempt to force conformity to a particular conception of a work process (e.g. Lotus Notes) have failed to achieve infrastructural status because they violate this principle (Grudin, 1989; Vandenbosch and Ginzberg, 1996). By contrast, email has become fully infrastructural because it can be used for virtually any work task.

Path dependence

Path dependence refers to the “lock-in” effects of choices among competing technologies. It is possible, following widespread adoption, for inferior technologies to become so dominant that superior technologies cannot unseat them in the marketplace. Standard examples include keyboards (QWERTY vs. Dvorak), video (VHS vs. Betamax), and nuclear reactor designs (light water vs. gas-graphite). Factors that can reduce the ability to adopt an alternative include social investment (e.g. the ~100-hr. training required to learn QWERTY) and the difficulty of overcoming positive network effects (e.g. for the case of automobiles, the gasoline distribution network was well established long before rural electric grids). Individual habits and organizational routines are highly efficient modes of organizing behavior, but they can be strongly resistant to change.

Key elements of the path dependence concept are:

- **Localized learning.** Individuals and organizations “satisfice” rather than optimize. All possible technological choices cannot remain on the table forever. Once they have made an initial investment, people adapt themselves, their organizations, and their technological choices to that investment rather than (re)consider alternatives (Foray, 1997).
- **Irreversibility.** Beyond some tipping point of widespread adoption, choosing an alternative to the dominant system becomes too costly, not only in money but also in time, attention, retraining, and coordination.
- **Network effects.** The value of certain kinds of technology increases exponentially with widespread adoption. Telephones aren’t worth much if only a few people have them, but become indispensable when most people do.
- **Inefficiency.** For economists, true path dependence exists only if some alternative technological path would be substantially more efficient in some sense (usually cost, but also labor time, etc.). This effect is debated. Some economists argue that the claimed inefficiencies are improbable and in any case cannot be proven, since there is no way to determine all the real-world ramifications (including inefficiencies) of an alternative technological system if it was never widely implemented (David, 1985).

Whether or not path dependence leads to economic inefficiencies, the concept is a useful metaphor for cyberinfrastructure developers (Figure 2). Technological change is always path dependent in the sense that it builds on, and takes for granted, what has gone before. Today’s choices constrain tomorrow’s possibilities. Yet they also create new possibilities, i.e. directions that could not have been taken in the absence of

technology X. Thus — as workshop participants stressed — path dependence leads to many positive effects.

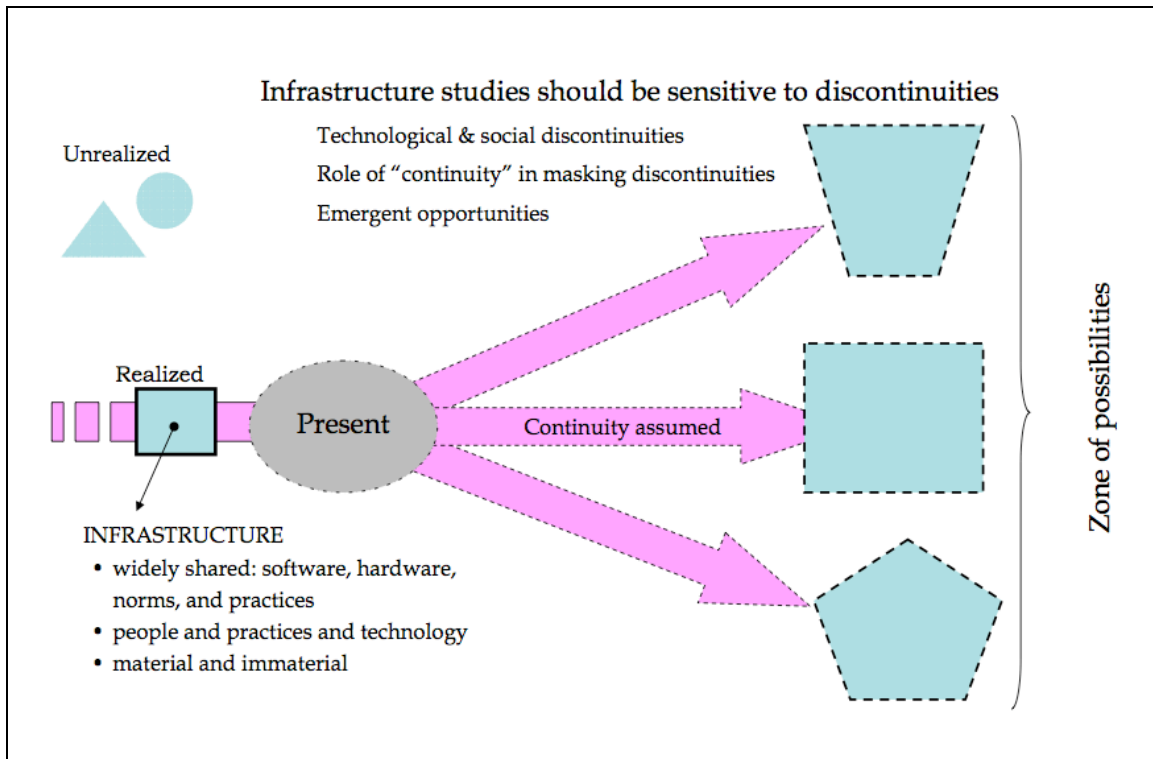


Figure 2. Visualizing path dependence and discontinuity. (Graphic prepared by Trond Jacobsen.)

As a metaphor, path dependence also applies to the practice of science in ways particularly relevant to cyberinfrastructure. Progress is possible precisely because new practices build upon old ones (positive path dependence), but this can also mean inheriting defects, entrenching them even more deeply. Climatologist Michael Oppenheimer coined the term "negative scientific learning" to describe this phenomenon.

A relevant example is data collection in the experimental sciences. Since the currency of scientists' careers is reputation, based on credit for new discoveries, the data produced by experiments were traditionally treated as the private intellectual property of the experimenter. Typically these data were closely guarded, at least until the results of data analysis were published, and often afterward as well. Publication of raw data in their entirety was a rare exception, not the rule.

In the last two decades or so, as data sets have grown ever larger and new techniques of data mining and reanalysis have improved, it has become clear that the private ownership model for scientific data represents an inefficient use of resources (as path dependence would predict). Much experimental data can and should be released for others to analyze and reuse. The NSF and other agencies now require public release of data after an appropriate waiting period (to allow experimenters to publish and receive credit for their work). Yet despite this requirement, changing practices based on the

private ownership model has proven much more difficult than anticipated, for both technical and social reasons. Decades or even centuries of private-ownership practice has led to a plethora of data collection practices and data formats, many of them idiosyncratic, as well as an absence of the metadata needed by other scientists to understand how the data was originally produced. The cultural norms of many experimental sciences devalue efforts to share or publish data, or even to record metadata beyond that needed by the original producers. The consequence is that much “shared” data remains useless to others; the effort required for one group to understand another’s output, apply quality controls, and reformat it to fit a different purpose often exceeds that of generating a similar data set from scratch.

Path Dependence: Notes from the Workshop

The integration of ad hoc systems into large networks is rife with examples of “bad choices.” Tight integration leads to a need for standards, which requires making choices — often those that in retrospect create inefficiencies. By the same token, there are manifold examples of felicitous choices.

We need to distinguish between technological and institutional paths. Humans can function as intermediaries across boundaries imposed by technological paths, linking projects, disciplines and institutions. How can we cultivate awareness of the social dimensions of integration?

Communication across fields, cultures, and institutions begins with pidgin languages. If this communication endures, the pidgin can become a full-fledged creole: a bridging lingua franca, spoken natively across divides. What incentives exist, or can be created, to enable or generate translation across entrenched practices and institutions?

“Scientific revolutions” can be seen as breakout moments when old, well-worn paths of theory, data collection, and analysis are overturned in favor of new ones. The cyberinfrastructure and e-science efforts, if successful, may represent such a moment. Intellectual path dependence implies that seeding ideas and work practices that view data as a fundamentally collective, shared resource, rather than as the private possession of individuals and work groups, could have enormous impact.

Scale effects

Cyberinfrastructure developers are focused on transforming small-scale, short-term, local projects into large-scale, functional infrastructures. The pattern of history suggests that this will take a long time — on the order of decades.

Scholars of the diffusion of innovation have demonstrated an “S-curve” pattern of adoption for successful large technical systems. In the case of major infrastructures, the duration of this curve is typically 40-50 years (Figure 3). After an initial period of linear growth, these infrastructures entered periods of exponential growth, falling into a slow linear growth period again upon reaching their maximum extent. Both the adoption rate

(Figure 4) and the 35-year period between about 1970 (early ARPANET) and the present fit this model well.

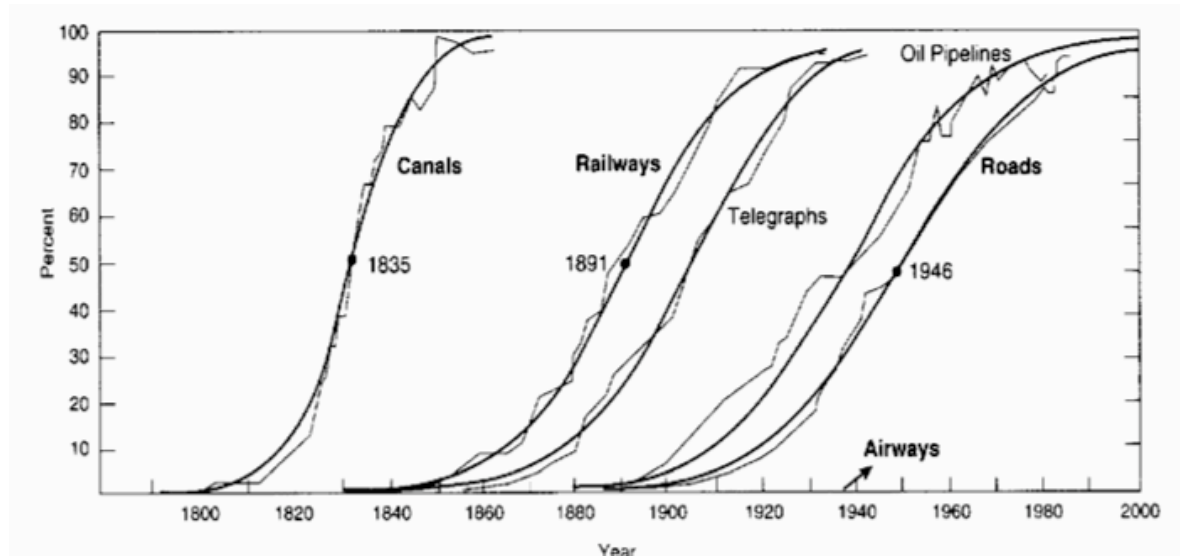


Figure 3. Growth of infrastructures in the United States as a percentage of their maximum network size (reproduced from Grüber and Nakićenović, 1991).




Figure 4. Technology adoption relative to total US population. B = birth of system (reproduced from Hannemayr, 2003).

Applying the historical model to cyberinfrastructure

Just as email has become infrastructural because it can be readily used for many tasks, the Excel spreadsheet has become the vehicle of choice for transfer of data among some scientific communities. It is easier for scientists in dissimilar specialties to use the pidgin they all know (Excel) than to learn each others' "languages," i.e. database designs (Borgman, 2007). The example demonstrates the basic principles of infrastructure formation:

- Reverse salients (no common data format or software);
- Attempts to bridge dissimilar systems using Excel as a gateway;
- Path dependence (entrenchment of a widely shared, but inefficient standard that arrived on the scene before better ones were available)

Finally, of course, the example illustrates the urgency of the cyberinfrastructure project.

Where does scientific cyberinfrastructure now stand along the infrastructure development path sketched in this section? Without further research no detailed answer to this question is possible, given the large number and wide variety of cyberinfrastructure elements. We think much could be learned from a more systematic attempt to do this. Still, we can sketch a very preliminary analysis, encapsulated here in Table 3 (p. 23), comparing electric power, electronic digital computing, and cyberinfrastructure.

The basic pattern to date, to which NSF initiatives are responding, has been one of stovepiped construction of specialized systems for domain sciences. Within a few domain areas, especially meteorology, purpose-built networks for distributed digital data collection, analysis, and distribution predate the Internet (World Weather Watch; Edwards 2006), but these were rare. Other sciences took advantage of the Internet and NSFNet in the 1980s. Numerous collaborative projects were initiated in a variety of domains during the first half of the 1990s (although predecessor projects date to the early 1970s).⁸ More recent major projects include Teragrid and NEES, both begun in 2000, and GEON (started in 2003).

Although many of these projects involve innovative collaboration among subdisciplines that may not normally connect, the LTS historical model would characterize virtually all of them as system-building — i.e. the first, early phase of competing and conflicting, locally based development (where "local" in this case references scientific domain rather than geographical location). They supply limited, specialized services to a predefined community, rather than providing ubiquitous, widely shared, commodity services across the sciences. In terms of the historical model, technology transfer would be the next step. This might occur, for example, if some set of tools built for one collaborative or distributed information environment were taken up by a similar project in another

⁸ For a substantial list by date, see the catalog developed by the Science of Collaboratories project, www.scienceofcollaboratories.org/Resources/colisting.php?startDate+asc.

scientific domain. Another step would be the emergence of gateways capable of connecting existing systems.

Comparative analysis of different cyberinfrastructure projects could — as we suggest in the Recommendations section of this report — illuminate the kinds of transfers, gateways, and other processes most likely to lead in the direction of true infrastructure.

Development phase	Electric power networks	Computing	Cyberinfrastructure
System building	Edison, Westinghouse	Univac IBM "Seven Dwarves": GE, Honeywell, etc.	Grid computing (GridPhyN, Teragrid) Collaboratories (SPARC) Domain science networks (NEES) E-science, e-Social Science (UK)
Reverse salients	Long distance power transmission AC/DC conversion Load factor	Machine- and vendor-specific languages, operating systems, programs Lack of backward compatibility (both hardware and software) High cost and fragility of main memory (tubes)	Middleware specifications Diverse data formats, dispersed databases Metadata: collecting, standardizing Disciplinary cultures Incentives Institutional cooperation Intellectual property rights
Technology transfer and growth	English and German Edison Paris and London exhibitions Thomas-Houston Company merges with Edison Electric	IBM International, ICL (UK), CII Honeywell Bull (France), etc. Packaged software	Geosciences Network (GEON) Earth Systems Modeling Framework (ESMF)
Consolidation	General Electric Westinghouse Public electric utilities	IBM System 360 IBM plug-compatibles and clones Wintel PCs (Microsoft Windows + Intel chips)	???
Gateways (between systems)	AC/DC rotary converter Standard AC voltages Load balancing Regional coordinating centers Computer modeling of loads	Compilers (machine-independent programs) Unix Interface Message Processors (ARPANET) Requirement for ARPA and NSF grantees to build network Ethernet IBM antitrust suits	Excel spreadsheets (common data format) Modeling frameworks Metadata standards ???
Network	National electric grid	Vendor-specific networks ARPANET, NSFNET, BITNET	???
Gateways (between networks)	International coordinating centers	TCP/IP; Unix with built-in TCP/IP support HTML	???
Internetwork	Deregulated electricity supply, integrating small suppliers International electric grids	Internet WWW	???

Table 3. Comparing infrastructure development: electric power, computing, cyberinfrastructure.

Understanding Infrastructure

III. Tensions

In addition to the patterns and dynamics noted above, infrastructures of all types have encountered, and often provoked, a series of deeply felt tensions. Once established, infrastructures may hide or disguise such tensions, so that once bitterly-contested decisions and design choices appear as unproblematic or even natural features (for example, that telephone networks should connect households as well as businesses). In this way, “achieved” infrastructures may take on an aura of inevitability that makes them seem uncontroversial and harmonious, lacking any conflict.

This impression is misleading. Even long-standing infrastructures turn out never to be quite as finished as we might suppose. Also, in their moments of emergence, infrastructures can be a site of intense conflict, through which relevant social actors (“legitimate stakeholders” in policy terms), the distribution of benefits and losses, and even the general “rules of the game” are worked out simultaneously. From this perspective, infrastructures, especially those in the making, are what political scientists term *agonistic* phenomena: imagined, produced, refined, and occasionally reassessed in a stratified and deeply conflictual field.

Because of its potential to upset or remake previously accepted relations and practices — as noted repeatedly in NSF’s Cyberinfrastructure Vision report — the development of new infrastructure may include a good deal of what economists have labeled “creative destruction,” as practices, organizations, norms, expectations, and individual biographies and career trajectories bend – *or don’t* – to accommodate, take advantage of, and in some cases simply survive the new possibilities and challenges posed by infrastructure. This section surveys a few such tensions, and points out the distinctive challenges and opportunities these may pose to cyberinfrastructure.

Interest and exclusion

Across virtually every type and class of emergent infrastructure we can identify provisional “winners” – those whose positions, programs, quality of work or general life experiences are enhanced by the developing infrastructure – and “losers”: 19th-century towns bypassed by the emerging networks of rail and road; neighborhoods without a public sewer system; categories of work and worker rendered obsolete by the shifting automation strategies of companies and industries; academic communities sidetracked by new waves of scientific funding; social groups further disadvantaged by the orientation and pace of high-speed infrastructural development. Emergent infrastructures function as redistribution mechanisms, reorganizing resource flows across scales ranging from the local workplace or research laboratory to the global economy. Few if any come free of distributional consequences altogether.

Short-term experiences of gain and loss will shape the incentive structures of individuals and institutions tasked with responding to infrastructural change. This in turn will shape the climate within which infrastructures struggle to emerge: broadly receptive, with allies adding support and innovation to extend the reach, quality, and fit of infrastructure? Or openly or covertly hostile, with important user groups and audiences dragging their heels, undermining change, putting forward counter-projects, or simply refusing to play

along? Failing to think proactively about the distributional consequences of infrastructure is not only bad politics, but bad business.

Thinking in this way calls attention to what we might term, borrowing from British geographer Doreen Massey (1993), the distinctive “power geometries” of infrastructure. In her study of the uneven distribution of benefits, burdens, and mobilities afforded by advanced transport and communication networks, Massey notes that

different social groups and different individuals are placed in very distinct ways in relation to these flows and interconnections. This point concerns not merely the issue of who moves and who doesn't, although that is an important element of it; it is also about power in relation to the flows and the movement. Different social groups have distinct relationships to this anyway-differentiated mobility: some are more in charge of it than others; some initiate flows and movement, others don't; some are more on the receiving end of it than others; some are effectively imprisoned by it.

In Leigh Star's terms, these are the ever-present “orphans” of infrastructure – the individuals, groups, and forms of social or professional practice that fit uneasily or not at all within the emerging infrastructural paradigm.

While being marginal vis-à-vis infrastructure can afford a measure of freedom, autonomy, and a potential for creative action, being marginal vis-à-vis consequential infrastructures will just as often come at the expense of heightened costs or barriers to action. Being “orphaned” by infrastructure is therefore to be on the receiving end of a significant exercise of exclusionary power, which post-hoc accommodations may never overcome. (Witness the retro-fitting of old buildings with ramps and elevators to accommodate people with disabilities; these rarely fully “correct” what are basically disability-unfriendly designs.) The frequent layering of inequalities, so that persons disadvantaged vis-à-vis one infrastructure (e.g. architectural design) may also suffer in relation to others (e.g. the educational system) only exacerbates this effect.

Banking the Unbanked

Banking and credit infrastructure plays a major role in the financial transactions that underpin daily economic life in the U.S., not least in the sphere of consumer and personal finance. In many ways, to be an economically competent member of society requires membership in this infrastructure: at minimum, access to a bank account, and more prosaically, reasonable physical access to banks themselves (or ATMs).

Such access is widely, but not universally, distributed. Some people have lost access through past defaults; some (e.g. undocumented aliens) lack the requisite paper trail; some live in neighborhoods un- or under-served by the established banking system. In such circumstances, infrastructural alternatives have arisen (family credit systems, private check-cashing establishments, etc.), but these often offer disadvantageous terms, increasing transaction costs for people who are already poor. Perhaps more importantly, they do little to connect people with other infrastructures (e.g. employment opportunities, credit markets) which depend on access to banking. Under these circumstances, seemingly mundane infrastructural moves, for example the extension or withdrawal of ATM services from poor neighborhoods, may have significant impacts on those people's life choices and life chances. Upstream exclusions may have important and expanding downstream consequences.

A different kind of tension emerges from the other end of the scale, namely the existence of established actors whose status and power owe much to their position vis-à-vis current infrastructural arrangements. If the earlier groups are orphaned by infrastructure, these groups are entitled by it – and will have a stake in furthering infrastructural development along lines that support and extend their strengths. The historical constitution of powerful classes of infrastructural users, both individual and collective, may constitute a powerful conservative force confronting and constraining new infrastructural development. Patterns of momentum and path dependency noted above are sometimes driven by this dynamic. Infrastructure, like regulation, may be subject to “capture,” in which the interests of powerful established constituencies come to overwhelm and crowd out potential innovations. Infrastructural incumbents may exploit their historically-accrued strengths to effectively hold infrastructure in place, stacking the deck against new, less organized, or less favorably placed actors, thereby limiting the scope and vision of new infrastructural possibilities.

Infrastructural Entitlement: The California Water System

Techniques of water management stand as a central contributor to California's explosive economic growth in the twentieth century. From the build-out of the San Francisco and Los Angeles urban systems to the later Central Valley and State Water Projects, the managed water system of California constitutes one of the largest sustained investments in public infrastructure in U.S. history – and by some measures, one of the most successful. Underlying this continuity, however, is a politically ambiguous logic of entitlement. The ongoing “mastery” of water via successive waves of infrastructural development has helped to produce not only an agricultural economy but an agro-industrial class which has emerged in recent decades as an extremely powerful, and some would argue fundamentally conservative, voice within California water policy.

Cases like this reveal the extent to which infrastructure can help to produce new kinds of stakeholders, whose influence may then work to fix or constrain future policy. Just as agriculture gave rise to agriculturalists in California, the ongoing push to develop cyberinfrastructure will carry with it the production of “cyberinfrastructuralists.” This is a necessary and in many regards positive part of any infrastructural development process. Left unmanaged, however, such entitlements may shape, fix, and limit the ongoing development of cyberinfrastructure.

The uneven distributional consequences of infrastructural change are matched by deep discrepancies in fundamental experiences and visions of infrastructural change. There is a tendency, already questioned in this report, to speak of “building” infrastructure (including cyberinfrastructure), as if the only perspective that matters is that of a putatively omniscient (and ideally well-intentioned!) system-builder. This architectural conceit echoes a long-standing tendency towards “great man” theories of scientific advance, in which heroic individuals acting more or less alone have shifted the terrain of human knowledge. While individual contributions *can* matter enormously, recent scholarship in the history and sociology of science should lead to caution in this regard. As this work shows, infrastructure is a deeply distributed phenomenon, involving actors of many types and levels. The variety of positions vis-à-vis infrastructure can lead to widely variant experiences and responses to infrastructure – many or all of which will need to be taken into account if the process of infrastructural development is to move forward effectively.

Infrastructure, Lived and Conceived

In The Practice of Everyday Life, French sociologist and historian Michel de Certeau offers the following striking (and now poignant) reflection on the nature and experience of urban space. Imagine, de Certeau asks us, two different observers of New York: one perched atop the World Trade Center and looking out over the city; the other encountering New York at street level, engaging the city through the practice of walking its streets. To the observer atop the tower, seeing the grid of city blocks and orderly rows of skyscrapers, the city will appear as an object of design: neat, orderly, geometric, amenable to rational planning and control. To the observer at ground level, enmeshed in more pedestrian, mundane, and everyday encounters, an entirely different city appears: disorderly, accidental, prone to disruption and surprise, and requiring of its users a good deal of day-to-day work.

The metaphor carries lessons for scholars and practitioners of infrastructure. Like cities, infrastructures are built phenomena, but also like cities, the experience of infrastructure is never fully determined or exhausted in design. Just as the design-eye view of the observer perched above the city can never fully capture or predict the complexity of the city-in-use below, so the perspectives of would-be infrastructure designers can never offer more than a first and partial approximation on infrastructure at work in local contexts.

Finding ways to translate between such design-level perspectives and the more “pedestrian” experience of infrastructure — and to continually incorporate lessons learned “below” into the next round of design from “above” — is among the central challenges in realizing the NSF vision for cyberinfrastructure.

Historical and comparative studies of infrastructure reveal frequent disconnects between such “design-centric” and “user-centric” visions of infrastructural development. Not surprisingly then, relations between designers and users emerge as repeated sites of tension in the real-world practice of infrastructure development. In some cases, discrepancies between the assumptions of designers and expectations of users have caused infrastructures to be questioned, opened up, and subjected to “user revolts” that have challenged, undermined, or in some cases improved upon, what had previously been regarded as elegant technical solutions. In others, perhaps the majority, the design-use disconnect is most eloquently expressed through neglect, as ambitious and well-intentioned systems languish on the shelf or desktops of users opting for alternative (perhaps local, perhaps kluged) solutions.

The careful nurturance of infrastructural change, and attending to the tensions that emerge from it, is a managerial and political skill of the highest order. It is also true that management often fails, and the quiet politics of infrastructure emerge as politics of a more recognizable and sometimes uncomfortable type. *Such instances of tension and resistance may constitute important sites of infrastructural learning and improvement, provided we can produce mechanisms that reliably surface and honestly report on difficulty, limitation, and failure (not a simple prescription, given the incentive structures*

prevailing among funders, sponsors, and builders of infrastructure). Tensions are best thought of as both barriers and resources to infrastructural development, and should be engaged constructively; in particular, they should be leveraged for their contributions to long-term properties of infrastructural fit, equity, and sustainability. Approaching tension from this perspective represents one way out of what we might term the *edifice complex* – the tendency to build first and ask questions later, or to treat the technical “code-and-wires” core as the realest or most essential thing about infrastructure, and the rest a social add-on – that has too frequently defined and limited the work of infrastructural development.

Ownership and investment models

A second class of tensions can be identified in instances where changing infrastructures bump up against the constraints of political economy: intellectual property regimes, public/private investment models, ancillary policy objectives, etc. Historically, the pervasive and foundational character of infrastructural systems such as road, rail, water, energy and telecommunication networks often led them to be assigned a public good or quasi-public good status,⁹ and therefore taken on as a primary venue of public investment. Many of the infrastructures transecting our everyday lives, including the systems of higher education and advanced research that many of us live and work in, have grown up on a diet rich in public funding. In recent decades, perhaps most notably in the U.S., long-standing public investment models have come under attack, and there is increasing pressure to constrain spending and/or partner with industry in ways argued to promote efficiency, innovation, and the transfer of ownership and function to private sector entities. At the same time, new and highly distributed development models (such as the open source movement) have appeared, offering what appear to be attractive alternatives to centralized and top-down development models. While such innovations should be welcomed and explored for their potential contributions, the historical analysis of past infrastructures should give us pause before wholeheartedly embracing a radically decentralized vision and strategy for infrastructure development.

Despite the difficulties of centralized development models, many now-mature infrastructures in the U.S. and elsewhere were built substantially through collective investments oriented to a public good logic. Sometimes this was achieved through strategic pairings of private ownership and regulatory oversight (witness the considerable success of AT&T’s regulated monopoly in extending telephony across the U.S. in the early to mid-20th century). In other cases, large-scale infrastructure was funded, shaped, and driven directly by the state, often in response to the demands of national security and/or economic competitiveness. The Internet, now widely seen as a model of decentralized, private sector-led development, depended through its formative years on an almost exclusive diet of DARPA and later NSF money.

⁹ Formal ownership is not necessarily the best marker here. While a distinction is sometimes drawn between the publicly-owned entities of Western and Northern Europe and the nominally private infrastructural history of the U.S. (e.g. AT&T, the 19th-century rail oligopoly, etc.), it should be noted that even the latter was built on an unusually intense coupling of private and public interest that recreated, in substance if not in form, many of the characteristics of a public infrastructure.

Some of the trends and tensions described above call into question this traditional road to large-scale infrastructural development. To the extent that targeted public monies remain an important spark and catalyst for infrastructural development, a key long-term challenge for American CI proponents will be to articulate a compelling and forward thinking *public investment rationale* for cyberinfrastructure. The effectiveness of such arguments will depend on the soundness of vision, demonstrable successes, and broad public contributions of the cyberinfrastructure program. Here again, targeted research on the political economy of infrastructure may help present-day infrastructure builders re-imagine and re-articulate the appropriate form, focus, and rationale for collective investment in infrastructure.

At the project level, practices of data handling, sharing, and the extended collaborative forms pursued under the NSF cyberinfrastructure vision may pose new challenges to existing regimes of intellectual property. Beyond tensions tied to the “internal” cultures and career structures of science (priority, publication, etc.), sorting out formal questions of ownership in vastly distributed projects may be a source of acute tension. This will be true most immediately in fields where the commercialization of research results is common place. For example:

- Who “owns” the results of deeply collaborative work?
- By what mechanisms can (or should?) downstream revenues from such work be distributed?
- How far does (or should?) property in “raw” data extend, when the reworking of community repositories leads to new results?
- Where researchers work against a mixed background of publicly owned and privately held intellectual property, how are rights to use and compensation to be balanced?

Such concerns are likely to multiply with the advent of increasingly networked and collaborative forms of research (as seen, for example, in NSF Cyberinfrastructure Council discussions of “virtual organizations”). They may arise in heightened form wherever academic research engages in direct and sustained contact with industrially-based research. Without clear answers to these questions, fears and concerns around the distribution of intellectual property flowing from deeply networked research may limit the scope of collaborations envisioned under cyberinfrastructure.

Similar tensions greet the relationship between national policy objectives and the transnational pull of science. Put simply, where broad-scale policy interests (in national economic competitiveness, security interests, global scientific leadership, etc.) stop at the borders of the nation-state, the practice of science spills into the world at large, connecting researchers and communities from multiple institutional and political locales. This has long posed a tension in science and education policy, showing up in practical terms in the complications of co-funding arrangements across multiple national agencies; restrictive structures greeting would-be foreign graduate students and researchers at U.S. universities; negotiating time allocations on transnational scientific resources such as CERN, etc.

Such national / transnational tensions have long shaped the development of infrastructure.¹⁰ At the extreme, the rail networks of Europe show the continuing effects of strategic decoupling, as infrastructure builders of the nineteenth and early twentieth century sought to restrict the feared advance of foreign armies. In other cases, strategic decoupling has been pursued for reasons of national economic advantage or protection; the enduring division between North American (NTSC), pan-European (PAL), and French (SECAM) color television standards is sometimes offered as an example of this. To the extent that cyberinfrastructure supports research collaborations across national borders, such national / transnational tensions may get picked up and replayed at the project level.

Data cultures, data tensions

In the daily working world of science, infrastructural tensions and conflicts are very often played out and resolved (or not) at the level of data. Data, and the anxieties and tensions it occasions, represents the front line of cyberinfrastructure development: its main site of operation, its most tangible output, and, in some ways (as the NSF's Cyberinfrastructure Vision document lays out) the target of its highest ambitions. From one view at least, cyberinfrastructure is principally *about* data: how to get it, how to share it, how to store it, and how to leverage it into the major downstream products (knowledge, discoveries, learning, applications, etc.) we want our sciences to produce.

It should come as little surprise then that the flat-sounding word “data” stands at the center of some of the most vexing tensions confronting the development of cyberinfrastructure. To begin, the word condenses a vast range of potential meanings. What counts as data varies profoundly across the fields of science addressed by the NSF. For some, data is first and foremost a question of *things*: samples, specimens, collections. For others, data is what comes out of a model – or perhaps the model itself. Data may be tactile, visual, textual, numeric, tabular, classificatory, statistical. Data may be an intermediate outcome, a step on the road to higher-order products of science (publications, patents, etc.). Or data may be the product itself. Where a discipline or research project fits within this spectrum will have enormous consequences for its positioning vis-à-vis cyberinfrastructure. This specificity alone guarantees that cyberinfrastructure should and assuredly never will be a singular or unified thing.

An important set of data challenges confronting and driving cyberinfrastructure development concerns the problems of storage, preservation, and effective curation. In some sciences, the sheer volume of data created on an ongoing basis makes effective data retention and back-up a challenge of the highest order. This raises important questions of form and granularity. How much data, and in what form, must one reasonably preserve to accomplish the task at hand? Is every data point required, or will second-order sets or inscriptions suffice? The answers to such questions are tied in turn to questions of short- and long-term audience and purpose. Is the data meant primarily to support the work-in-progress of a distinct team of researchers (what the NSF's

¹⁰ See, for example, the collective work of the “Tensions of Europe” project (www.histech.nl/Tensions/AboutUs/TekstAboutUs.htm) and its successor, “Transnational Infrastructures and the Rise of Contemporary Europe” (www.tie-project.nl/).

Cyberinfrastructure Vision defines as a “research collection”)? Is the data to be targeted and tailored to fit a larger, perhaps domain-level community (a “resource collection”)? Should the scope be wider still, aimed at supporting vast and multi-disciplinary teams over long spans of time (a “reference collection”)? Under each of these scenarios (which rarely separate so neatly in practice) fundamentally different data practices may be indicated, touching on such basic questions as the optimal form, format, and granularity of data. These will also often be in tension: for example, in trade-offs between the goals of optimizing to present and local purposes vs. the longer-term flexibility and wider usability of data.

Questions of preservation become still more complicated when prospects of reuse outside the immediate context of data production are considered. Here the thorny problem of metadata emerges: how much “data about data” do we need to support what may be as-yet dimly anticipated future, alternative, or comparative use? Historical solutions to this problem have been distinctly human: beyond the thin accounting of journal reports, scientists come to nuanced assessments of the techniques and findings of their colleagues by correspondence (now, typically, email), by hallway or dinner-time conversations during site visits or academic conferences, and by trading grad students or post-docs. For years now, the NSF and other funders have exhorted their grantees to collect and preserve metadata – a prescription that has, for the same number of years, been routinely ignored or under-performed. The metadata conundrum represents a classic mismatch of incentives: while of clear value to the larger community, metadata offers little to nothing to those tasked with producing it and may prove costly and time-intensive to boot.

Robust and widespread practices of data sharing within and across disciplines represent an equally important goal of the cyberinfrastructure program, and one that has been similarly challenging to accomplish. An important class of barriers lies in the sheer diversity cited above – how does one design tools with the range and ability to accommodate and translate between the distinctly different data needs of the various domain communities? Even if “technical” solutions can be put in place, how can participants from one disciplinary community make sense of data produced under the very different procedures and understandings of another? As pointed out in the workshop, data are the product of “working epistemologies” that are very often particular to disciplinary, geographic, or institutional locations. Data in the atmospheric sciences might not be understood as such by oceanographers, *in the absence of a good deal of articulation work*. To the extent that this articulation work has long been wrapped up in very human assessments of quality and fit, dreams of a seamless and self-sufficient web of data may be misguided.

It is also possible that a tech-centered approach to the challenge of data sharing inclines us toward failure from the beginning, because it leaves untouched underlying questions of incentives, organization, and culture that have in fact *always* structured the nature and viability of distributed scientific work. Questions of trust loom large here, and run both ways. Can I trust those I agree to share my data with to make reasonable and appropriate use of it, and on a timeline which doesn’t impede my own requirements re: publication, credit, and priority? On the receiving end, can I trust the data I’m getting, particularly as collaborative webs lengthen and my first-hand knowledge of the data and its producers recedes? Here again, there is considerable local and disciplinary variation in the way such norms and routines have been structured. In domain fields with long and robust histories of collaborative research, norms of sharing may be well advanced,

widespread, and highly structured. In others, the collaborative terrain may be more uneven, and norms and procedures for sharing relatively ill-defined. Where uncertainty exists, and where two data cultures collide, well-intentioned efforts to promote sharing via technical or organizational fixes are unlikely to succeed. This aspect of scientific work, absolutely central to the daily practice of scientists, is arguably the one least well understood within the current cyberinfrastructure development process.

IV. Design

When you are designing a cyberinfrastructure, there are several things going on simultaneously. You are trying to deploy the latest computing infrastructure to:

- Permit distributed collaborative work;
- Engineer changes in the organization of scientific work (e.g. altering reward structures for database work, or encouraging early sharing of results);
- Enable interdisciplinarity in a way which will get scientists from disparate communities working together.

This work is principally social and organizational. Yet in general designers are not trained to recognize these dimensions of their work practice – hence the horrific image of throwing new products over the wall to their designated community. One clear finding of the *History & Theory of Infrastructure* workshop was that based on the previous history of infrastructure development, cyberinfrastructure will not be built from the center with a single design philosophy. Instead, it will be built from the ground up, and in modular units. As one participant stated: “It’s simply unrealistic to talk about designing cyberinfrastructure. Rather, each project produces a set of modules which ideally interoperate to create a larger whole.”

Navigating. Michel Serres’ wonderful metaphor of the Northwest Passage is evocative (Serres 1980). Serres uses it to talk about how we “get between” the social sciences and the humanities and the natural sciences, but we will use it for a subset of this issue, namely how to get between social organization and technical infrastructure. The point, he says, is that the Northwest Passage is ever changing: shifting ice floes mean that last year’s route will never be the same as the current one. What we need to teach, then, is not a rigid road map but *principles of navigation*. There is no one way to design cyberinfrastructure, but there are tools we can teach the designers to help them appreciate the true size of the solution space – which is often much larger than they may think, if they are tied into technical fixes for all problems.

Indeed, one of the most interesting and challenging questions for the designer is to know when something needs an organizational and when a technical fix. For example, you can get two databases talking to each other either by producing a single, shared ontology, or through a translation process between two native ontologies. When computer scientists do most of this work, we can characterize this as a technical fix; the communities involved need not change their work practices. The process of infrastructure building can remain mostly invisible to them. However, you can also get the communities involved to sit around a table and hammer out a set of common standards, to be implemented subsequently in technical form. This is often a long and

difficult task, as witnessed by the fights between plant and animal virologists in the 1960s, when they discovered that plants and animals could share the same virus (Bowker and Star 1999): neither group wanted to give up its own classificatory practices, but a compromise was clearly necessary. This is organizational and community work. With the technical fix you don't have to go through the process of consensus-building – you just produce a two-way mapping based on regular elicitation processes.

**Notes from the Workshop:
How does one plug in a computer at an airport?**

To use the same device in multiple power networks as we travel the world, plugs and sockets must be adapted to each other. This is an instance of the gateway problem. Ordinarily we solve it with a plug adapter, i.e. a special-purpose gateway that connects one kind of plug with a different kind of socket. Other solutions could exist, however, and all of them are simultaneously technical, political and economic. Other adaptations come with other locations of design, cost, and responsibility.

In this example, adaptation could also occur on the device plug (a plug with multiple pin configurations to fit multiple sockets). This would shift design, cost, and responsibility to the device designer. Or it could happen on the socket (a reconfigurable socket to fit multiple plugs). This solution would shift cost and responsibility to builders, building codes, socket manufacturers, and potentially to electric power providers.

Either of these solutions would place less burden on travelers than the one we actually use!

A “live” version of this problem has cropped up around the new Airbus A380 superjumbo jet. Both runways and airport terminal gates, standardized to the current generation of large aircraft, must be widened to accommodate the A380's enormous wingspan. New jetways and service vehicles high enough to reach the A380's upper deck must be procured. Is the cost of these modifications the responsibility of the airports where the planes will land? Or of the airlines that will fly the jets?

Boundary work. A number of scholars, who might be characterized as “boundary workers,” have been seeking answers to the question of how to bridge the “great divide” between system building and social analysis (Bowker, Turner et al. 1997). Some good bridges have been constructed. One is the computer supported cooperative work (CSCW) community, which builds in part on the Scandinavian participatory design movement. The latter is based on the legal requirement in Scandinavian countries that when new information technology is brought into the workplace, it should be co-designed with the user community. A group of ethnographers has been trained to act as honest brokers between designers and users, explaining the contingencies of each to the other and suggesting ways forward. This is particularly useful since users don't often know quite how to express their needs to designers, and designers don't know quite how to express their limitations and possibilities to users. In another example, social network

analysis is being used increasingly to chart patterns of communication in distributed organizations and thus to provide a window for designers into the social world of their users and guide their strategy.

Three immediate lessons were drawn during the workshop. First, we need to move away from the model of seeing the social and organizational as sitting on top of the technical. Take the OSI protocol layers (Figure 5):




Figure 5. Open Systems Interconnection (OSI) reference model.

Since all of these — except perhaps the application layer — lie “below” computer users and their communities, we might naïvely conclude that the latter are irrelevant to protocol design. Instead, the social and organizational are intimately mixed with the technical at each of these layers – and in different configurations.

In the wonderful words of Jacques Revel: “the change in the scale of observation revealed not just familiar objects in miniature but different configurations of the social” (206). Second, we need to train up not only computer systems designers. We need to prioritize training up informatics workers. Typically, information managers in cyberinfrastructure projects come from a wide range of backgrounds, having migrated from computer science into the interstices of domain science or vice versa. Each migration path has its problems. It is through the information managers that much of the

eventual design work is done. As Steward Brand argued so well about buildings, the design process has only just begun when the building is completed (Brand 1994). This observation is consonant with studies of innovation carried out over the last thirty years at the Center for the Sociology of Innovation at the Paris School of Mines. Third, as a corollary, we need to recognize that cyberinfrastructure design is about recognizing emergent possibilities – being flexible and nimble.

Standards and Flexibility

This last observation leads to the fundamental difficulty of cyberinfrastructure design: path dependence – what you do now affects what you can build in the future. We have all suffered the deleterious consequences of path dependence as we have struggled over the years with Windows operating systems, or with QWERTY keyboards (the only technical artifact in the world to favor left-handers). Path dependence recognizes that past decisions limit future options. These limits are all the more obdurate when built into large, complex infrastructures. Is there a way to include flexibility in infrastructure design in order to facilitate change, potentially making infrastructure more responsive to evolving conditions?

This is a challenging problem, but it can be tackled in several different ways. For example, one can emphasize modular design of technical components. Another possibility is to focus on standards as a means of creating flexibility. Standards are knots in the web of infrastructure technologies and concurrent socio-institutional provisions. They are key focal points for actors' negotiation of their differences, interests, and opinions.

Standards are often associated with freezing development, in the light of enhancing system flexibility. It requires a different mindset. This mindset or conceptual framework is being developed and applied across infrastructures (i.e. transportation, information, batch processing industry, energy) to explore the conditions of and restrictions to our flexibility claim. Key concepts in the framework are: the gateway (dedicated and generic), compatibility between the political, operational and technical domains; flexibility objectives and flexibility characteristics of standards. Important work on these objectives is now being carried out by the Next Generation Infrastructures project based at Delft, the Netherlands (www.nginfra.nl/).

To design for flexibility, at each stage, with each dynamic, and in every type of intervention or evolution, possibilities for seeking alternative pathways must always be kept open. However, it is crucial not to underestimate the economic context within which such flexibility is permissible, and the resources it demands in both human and financial capital. Research communities should think of “resource pooling” or “collaboration” by first auditing what is being shared, comparing and contrasting across projects, in different countries, different parts of the world, in different experiences.

V. Conclusions and Recommendations

This report has drawn out some broad-scale findings from the comparative history, social study, and practice of infrastructure, with an eventual eye to policy- and program-level recommendations of relevance to CI and CI-related activities at the National Science Foundation. As we have suggested, while each infrastructure grows and evolves according to circumstances and dynamics all its own, there are also higher order patterns and processes that mark with some consistency the development of infrastructure. This suggests a modest but real advisory role for history and comparative social analysis: knowing about railroads, water systems, power grids, the Internet, etc. in their periods of formation can — and we believe *should* — inform the way we approach the current work of cyberinfrastructure development.

The *Dynamics* section shows that understanding when, where, and how past infrastructures made the jump from isolated systems to widely shared, highly accessible networks; how they came to achieve scale and scope; how they responded to and incorporated technological, social, legal, and organizational change; and how they dealt with the series of tensions and challenges generated by their own development: all of these can help us to think strategically about the present challenges facing the development of cyberinfrastructure. The lessons of history also include a rich set of instructive failures — the infrastructures that weren't — which may be mined for cautionary tales for today's cyberinfrastructure practitioners. Such lessons will rarely take the form of simple plug-and-play solutions or definitive “no-go” recommendations. Indeed, extracting “lessons” will require a good deal of translation work, which historians themselves have often been unwilling or uncomfortable to assume. The patterns and concepts described in the *Dynamics* section represent an early attempt at such translation; we believe that further work along these lines could provide both language and landmarks — if rarely anything so structured as a map — to the ongoing process of cyberinfrastructure development.

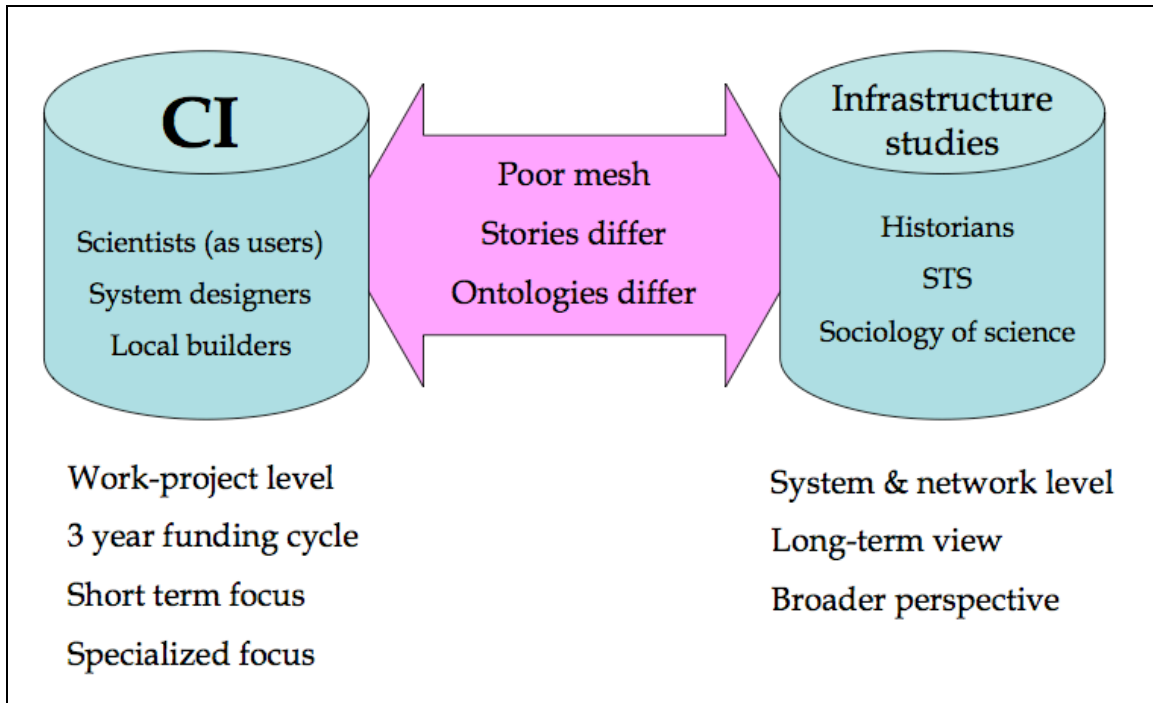


Figure 6. Perspectives and time scales: how social studies of infrastructure differ from cyberinfrastructure projects (CI). (Graphic prepared by Trond Jacobsen.)

As the *Tensions* section explains, infrastructural development is always a contested process, tied as it is to questions around access, power, and the life chances of groups and individuals. Would-be developers of infrastructure work within pre-constituted fields, and regularly encounter actors, both entrenched and emergent, who will see in the development of infrastructure opportunities for both gain and loss — and gauge their responses accordingly. This aspect of infrastructure has obvious implications for equity, participation, and a range of other broad social goals (including those expressed in *NSF’s Cyberinfrastructure Vision for the 21st Century*). It is also an issue of strategic management, fit, and long-term sustainability. Systems that fail to acknowledge and accommodate the tensions they inherit or provoke will have little chance to attract and sustain a broad scale base of users over time – and therefore little chance of rising (or sinking) to the level of infrastructure. In the world of cyberinfrastructure to date, we have often seen such tensions play out over the production, curation, and sharing of data — though just as often such data tensions serve as proxies for conflicts of a disciplinary, institutional, biographical, or broadly “cultural” sort. Here again, our workshop findings point to the need for more and better research into such dynamics.

Finally, our discussions around *Design* suggest that the language of “building” (cyber)infrastructure (in the sense of creating it either from scratch or according to an orderly progression from plan) may be misguided, and seriously overstate the capacities for action and control available to central system-builders.

Nevertheless, action can be undertaken now to nurture, foster, and support the growth of an advanced computational infrastructure for scientific work that might realize at least some of the goals and ideals laid out in the NSF cyberinfrastructure vision. Under such circumstances, the patient art of “growing” infrastructures will depend less on the

Herculean figure of the master engineer, and more on a series of pragmatic, modest, and strategically-informed interventions undertaken on the basis of imperfect knowledge and limited control. Our report indicates first, modest findings in that direction. Further research into promising (and otherwise) design strategies is once again needed.

Recommendations

In addition to the themes referenced above, workshop participants were asked to distill a more bounded set of recommendations targeted to funders and program managers associated with the NSF cyberinfrastructure program (and related initiatives). These recommendations, necessarily preliminary in nature, are intended to bring the emerging field of historical and comparative infrastructure studies to bear on the immediate needs and concerns of the cyberinfrastructure program. All are intended to enhance the quality, effectiveness, and long-term sustainability of CI investments.

Learning from cyberinfrastructure

How we can learn more about “growing” infrastructures by studying current cyberinfrastructure projects, in an iterative and informative cycle potentially beneficial to those projects and future ones?

Comparative social scientific analysis of cyberinfrastructure projects. A large number of cyberinfrastructure projects are presently proceeding in parallel, with little overlap or cross-communication. Which ones are succeeding? How, and why? Comparative study of these projects and performing high-level requirements analysis,¹¹ with particular attention to their handling of social, cultural, and organizational issues, may aid in developing programmatic principles for future projects. This would extend, deepen, and generalize from the limited social scientific studies of cyberinfrastructure that have been conducted to date.¹²

Cross-agency and cross-national comparisons. Beyond the comparative analysis of NSF-funded projects indicated above, broader comparisons with related initiatives undertaken by other federal funders (NIH, DOE, etc.) and other national funding agencies federal agencies (e.g. UK and EU “e-Science” initiatives) should be supported. Such analyses would extend the effective case set of cyberinfrastructure studies. Stepping beyond NSF and other institutions of U.S. science funding would allow comparative insights at a broader scale to emerge. Such analyses would ideally be undertaken by multinational teams of researchers using common research questions, methods, and protocols.

¹¹ Bergman, M., King, J. L., and Lyytinen, K. (2002). Large-scale requirements analysis as heterogeneous engineering. In *Social Thinking: Software Practice*, Y. Dittrich, C. Floyd, and R. Klischewski, Eds. MIT Press, Cambridge, MA, 357-386.

¹² E.g. the Comparative Interoperability Project and the Science of Collaboratories Project.

Accurate and realistic reporting mechanisms. Anecdotal evidence from many of the workshop participants suggests that standard forms of project reporting, given the incentives of both funder and grantee, will tend to over-report experiences of success and under-report those of difficulty or failure. Efforts to accommodate and encourage the honest reporting of failure could go a long way to supporting long-term and comparative learning across the varieties of cyberinfrastructural experience. As science itself has proceeded through the disciplined and even-handed study of failure, funders and proponents of cyberinfrastructure must learn to stop hiding the bodies.

Aligning extant or creating new incentive structures. Many workshop attendees pointed out that the current reward system in the academy does not provide incentives to scientists, researchers, students, or administrators to contribute to cyberinfrastructure-based activities. Creating metadata to share effectively, engaging and accommodating the local routines and contexts of distant collaborators, and archiving project data for wide distribution and re-use are critical to the development of cyberinfrastructure, but tend to be time and effort intensive. Further, these activities go unrecognized when appraising the value of research activities, and are considered articulation work when performed.

The NSF can consider two routes to provide alternate incentives. First, through mandate, providing specific provisions for the preparation, sharing, and distribution of data resulting from NSF-sponsored research. Alone, this strategy is likely to be resented and unsuccessful. A second method is to work with institutions to legitimate CI-building activities as rewardable academic work, gently guiding the culture to provide personal incentives to researchers. The combination of these two strategies, providing a framework through requirements and creating or encouraging reward structures for meeting these requirements, may slowly shift academic culture toward natural and local support for CI.

Instrumenting cyberinfrastructure. In addition to improved and incentive-sensitive reporting mechanisms, learning from CI development efforts to date could be enhanced by the proactive “instrumenting” of CI projects for social scientific analysis. This could include provisions or protocols for data collection and preservation built into the terms of the grants themselves, as well as into the mechanisms of service delivery. Examples of these could include provisions for retention of project-related communications (emails, meeting notes, etc.); and “sampling” interviews across the range and duration of specific projects (to capture subtle shifts in project composition, participants, and evaluation criteria over time). Such efforts at instrumentation could be coordinated with one or more of the larger comparative research programs outlined above. Individual CI projects may also benefit from the early (rather than the post-hoc, primarily evaluative) inclusion of social scientists on project teams.

Improving cyberinfrastructural practice

How can we enhance training, preservation, and other measures to improve the quality of cyberinfrastructure projects and personnel?

Infrastructural diagnostics. In addition to the evaluative goals expressed in the first set of recommendations, a capacity for in-process analyses could become a regular part of the working repertoire of cyberinfrastructure practitioners. Sensitization to the dynamics and tensions of infrastructure, and an analytic capacity to see these early on, could support in-course corrections contributing to local project efficacy, relevance, and sustainability. Improved diagnostic capacities would need to be matched with sufficient funding flexibility at the project or broader program level to support necessary changes in course.

Training for information managers. This might take the form of an ongoing center, summer workshop, or eventually a formal graduate program, perhaps most logically based in one (or more) of the information schools.¹³ As work in the Long Term Ecological Research (LTER), ocean science, and other research communities suggests, information managers occupy a strategic, under-recognized, and hard to fill niche in the contemporary practice of science. IMs are often the clearest point of human articulation between domain-specific needs and demands and the emerging opportunities and constraints of cyberinfrastructure. IMs are often also the principal keepers, curators, and coordinators of data. At present, the training of IMs remains largely haphazard.

Cross-disciplinary symposia for graduate students, post-docs, and early-career faculty. Cross-disciplinary meetings of researchers engaged (or soon to engage) in cyberinfrastructure-related work may produce a common focus, above domain-level specifics, on shared experiences, problems, skill sets, and application potentials for cyberinfrastructure. Reaching researchers at such critical career junctures would also help to cultivate an important and strategically placed future user class, producing important long-term benefits in the understanding and practice of cyberinfrastructure. Attendance at such meetings could also be structured to support the goals of inclusion and ecological diversity noted below.

Fitting funding to time scales. Workshop participants noted an often serious mismatch between the decadal time scales at which large-scale infrastructures have historically

¹³ There are now 19 schools of information in the USA and Canada, engaged in graduate and undergraduate education of information professionals. They are loosely defined in the I-Schools Charter as "schools interested in the relationship between information, technology, and people. This is characterized by a commitment to learning and understanding of the role of information in human endeavors. The I-Schools take it as given that expertise in all forms of information is required for progress in science, business, education, and culture. This expertise must include understanding of the uses and users of information, as well as information technologies and their applications" (www.ischools.org/oc/charter.html). The second annual conference of I-Schools was held at the University of Michigan in October 2006 (iconference.si.umich.edu/index.htm); the third is scheduled for UCLA in October 2007.

developed and the typically much shorter, 1-5 year time horizons of NSF funding at the project and even program levels.

Greater experimentation with multiple temporalities of funding, and in particular the extension of at least some categories of support on a longer-term basis, may help to align and accommodate the real time scales on which infrastructures grow. Insofar as such extensions relax or suspend expectations around short-term turnaround of results, they may also indirectly support the goals of infrastructural learning outlined above.

Enhancing resiliency, sustainability, and reach

How can we accommodate change, enhance flexibility, and extend the reach of cyberinfrastructure and the cyberinfrastructure program?

Designing for flexibility and change. Given its relative immaturity and the rapidly changing technological backdrop against which cyberinfrastructure is unfolding, efforts not to prematurely “sink” or “fix” the form and vision of cyberinfrastructure (or distinct cyberinfrastructure projects) should be supported. As the ecosystem sciences teach, resiliency in the face of change is very often a feature of diverse local ecologies. Accordingly, efforts to support diverse cyberinfrastructural ecologies should be undertaken. Dimensions along which such diversity might be assessed would include, at minimum: disciplinary representation; actor type and scale; and form, mechanism, and timescale of investment. Research into standard-setting processes that allow for flexibility, as mentioned in the Design section, should be a priority.

Continued and expanded mechanisms to incorporate under-represented groups and institutions. Given the distributional and exclusionary concerns cited in the Tensions section above, ongoing efforts to extend participation in cyberinfrastructure to under-represented groups and institutions are needed. Early efforts in this direction (e.g. the Advanced Networking with Minority-Serving Institutions Project, www.msinetwork.org/AboutUs.htm) should be sustained and extended.

Form alliances with experienced niche organizations to eliminate redundant expertise. Expertise that is not core to scientific research, but which such research requires on an ongoing basis, may already exist outside the usual academic and governmental channels. For example, the management of intellectual property and licensing is a morass within one institution, let alone a multi-institution collaboration, and tensions between technology transfer offices is not uncommon. The Creative Commons (creativecommons.org), Science Commons (sciencecommons.org), Scholar’s Copyright Project, and Biological Materials Transfer Project are examples of Open IP schemes that may partner with the NSF to standardize aspects of information and data sharing for CI projects. Since such resources already exist, and have been performing this work for several years, it is easier and more efficient to capitalize on this experience through partnership.

VI. Bibliography

- Abbate, Janet (1999). *Inventing the Internet*. Cambridge, MA: MIT Press.
- Alexander, C. (1979). *A Timeless Way of Building*. New York, NY, Oxford University Press.
- Axelrod, R. and M. D. Cohen (2001). *Harnessing Complexity: Organizational Implications of a Scientific Frontier*. New York, NY, Basic Books.
- Baker, Karen, Steven Jackson, and Jerry Wanetick (2005). "Strategies Supporting Heterogeneous Data and Interdisciplinary Collaboration: Towards an Ocean Informatics Environment," in *HICSS38: Proceedings of the Hawaii International Conference on System Sciences* (Honolulu).
- Beniger, J. R. (1986). *The Control Revolution: Technological and Economic Origins of the Information Society*. Cambridge, MA: Harvard University Press.
- Bijker, Wiebe, Thomas P. Hughes, and Trevor Pinch, eds. (1987). *The Social Construction of Technological Systems*, Cambridge, MA: MIT Press.
- Bijker, Wiebe E. and John Law, eds. (1992). *Shaping Technology/Building Society: Studies in Sociotechnical Change*. Cambridge, MA: MIT Press.
- Borgman, Christine L (2000). *From Gutenberg to the Global Information Infrastructure: Access to Information in the Networked World*. Cambridge, MA, MIT Press.
- Borgman, Christine L (2007). *Scholarship in the digital age : information, infrastructure, and the Internet*. Cambridge, MA: MIT Press.
- Bowker, G. C. and S. L. Star (1999). *Sorting Things Out: Classification and its Consequences*. Cambridge, MA: MIT Press.
- Bowker, G. C., W. Turner, et al., Eds. (1997). *Social Science, Technical Systems and Cooperative Work: Beyond the Great Divide*. Mahwah, NJ: Erlbaum.
- Brand, S. (1994). *How Buildings Learn: What Happens after they're Built*. New York: Viking Press.
- Brand, S. (1999). *The Clock of the Long Now: Time and Responsibility*. New York, NY: Basic Books.
- Braun, Ingo (1994). "Geflügelte Saurier: zur intersystemische vernetzung grosser technische Netze," in Braun, Ingo and Bernward Joerges, eds., *Technik ohne Grenzen*. Frankfurt am Main: Suhrkamp, 446-500.
- Braun, Ingo and Bernward Joerges (1994). *Technik ohne Grenzen*. Frankfurt am Main: Suhrkamp.
- Brown, John Seely and Paul Duguid (2000). *The Social Life of Information*. Cambridge, MA: Harvard Business School Press.
- Buckland, Michael (2006). *Emanuel Goldberg and His Knowledge Machine: Information, Invention, and Political Forces*. Westport, CO: Libraries Unlimited.
- Castells, M. (1996). *The Rise of the Network Society*. Cambridge, MA: Blackwell Publishers.
- Chandler, A. D., Jr. and J. W. Cortada (2003). *A Nation Transformed by Information: How Information has Shaped the United States from Colonial Times to the Present*. New York, NY: Oxford University Press.

- Coutard, Olivier (1999). *The Governance of Large Technical Systems*. New York: Routledge.
- Coutard, Olivier, Richard Hanley, and Rae Zimmerman (2004). *Sustaining Urban Networks: the social diffusion of large technical systems*. New York, NY: Routledge.
- David, Paul (1985). "Clio and the Economics of QWERTY," *American Economic Review* 75, 332-37.
- David, Paul A. and Julie Ann Bunn (1988). "The Economics of Gateway Technologies and Network Evolution: Lessons from Electricity Supply History," *Information Economics and Policy* 3, 165-202.
- De Certeau, Michel (1984). *The Practice of Everyday Life* (tr. Steven Rendall). Los Angeles: University of California Press.
- Dennett, D. (1996). *Darwin's Dangerous Idea: Evolution and the Meanings of Life*. New York, NY, Simon & Schuster Publishers.
- Downey, Greg (2001). "Virtual Webs, Physical Technologies, and Hidden Workers: The Spaces of Labor in Information Internetworks," *Technology and Culture* 42, 209-35.
- Duguid, Paul (2005). "'The Art of Knowing': Social and Tacit Dimensions of Knowledge and the Limits of the Community of Practice," *The Information Society* 21 (2), 109-118.
- Duguid, Paul (2005). "Networks and Knowledge: The Beginning and End of the Port Commodity Chain, 1703-1860," *Business History Review* 79 (3), 493-526.
- Dutton, W.H., B. Kahin, R. O'Callaghan, and Wyckoff, A., eds. (2005). *Transforming Enterprise*. Cambridge, MA: MIT Press.
- Dutton, W.H. (2004). *Social Transformation in the Information Society*. Paris, UNESCO Publications for the World Summit on the Information Society.
- Edwards, Paul N. (1996). *The Closed World: Computers and the Politics of Discourse in Cold War America*. Cambridge, MA: MIT Press.
- (1998a). "Y2K: Millennial Reflections on Computers as Infrastructure," *History and Technology* 15, 7-29.
- (1998b). "Virtual Machines, Virtual Infrastructures: The New Historiography of Information Technology," *Isis* 89 (1), 93-99.
- (2006). "Meteorology as Infrastructural Globalism," *Osiris* 21 (special issue on *Science, Technology, and International Affairs: Historical Perspectives*, guest eds. John Krige and Kai-Henrik Barth), 229-50.
- Egyedi, Tineke (2001). "Infrastructure Flexibility Created by Standardized Gateways: The Cases of XML and the ISO Container," *Knowledge, Technology & Policy* 14 (3), 41-54.
- Egyedi, Tineke (2005). "A Standard's Integrity: Can It Be Safeguarded?" *IEEE Communications Magazine* 43 (2), 151-155.
- Egyedi, Tineke (2006). "Standards and Sustainable Infrastructures: Matching Compatibility Strategies With System Flexibility Objectives," in Sh. Bolin, ed. *The Standards Edge: Unifier or Divider*. Menlo Park, CA: Bolin Communications.
- Fano, Robert M. (1992 [1967]). "The Computer Utility and the Community," reprinted in *IEEE Annals of the History of Computing* 14 (2), 39-41.
- Fischer, Claude S. (1992). *America Calling: A Social History of the Telephone to 1940*. Berkeley: University of California Press.
- Finholt, T.A. (2003). "Collaboratories as a New Form of Scientific Organization," *Economics of Innovation and New Technology* 12, 5-25.

- Finholt, T.A. (2002). "Collaboratories," in B. Cronin, ed. *Annual Review of Information Science and Technology* 36. Washington D.C., American Society for Information Science.
- Foray, Dominique (1997). "The Dynamic Implications of Increasing Returns: Technological Change and Path Dependent Inefficiency," *International Journal of Industrial Organization* 15, 733-52.
- Forster, Paul W. and John Leslie King (1995). "Information Infrastructure Standards in Heterogeneous Sectors: Lessons from the Worldwide Air Cargo Community," in Kahin, Brian and Janet Abbate, eds., *Standards Policy for Information Infrastructure*. Cambridge, MA: MIT Press, 148-77.
- Foster, Ian (2006). "A Two-Way Street to Science's Future," *Nature* 440 (23 March), 419.
- Foster, Ian (2005). "Service-Oriented Science," *Science* 308 (6 May), 814-817.
- Fountain, Jane (2001). *Building the Virtual State: Information Technology and Institutional Change*. Washington, D.C., Brookings Institution Press.
- Fountain, Jane (2006). "Central Issues in the Political Development of the Virtual State," in Manuel Castells and Gustavo Cardoso, eds. *The Network Society: From Knowledge to Policy*. Washington, D.C., Brookings Institution Press.
- Friedlander, Amy (1995a). *Emerging Infrastructure: The Growth of Railroads*. Reston, VA: Corporation for National Research Initiatives.
- (1995b). *Natural Monopoly and Universal Service: Telephones and Telegraphs in the U.S. Communications Infrastructure, 1837-194*. Reston, VA: Corporation for National Research Initiatives.
- (1996). *Power and Light: Electricity in the U.S. Energy Infrastructure, 1870- 1940*. Reston, VA: Corporation for National Research Initiatives.
- Graham, Stephen and Simon Marvin (2001). *Splintering urbanism: networked infrastructures, technological mobilities and the urban condition*. New York: Routledge.
- Grübler, Arnulf and Nebojsa Nakićenoviâc (1991). *Long waves, technology diffusion, and substitution* (RR-91-17) Laxenburg, Austria: International Institute for Applied Systems Analysis, 313-342.
- Grudin, Jonathan (1989). "Why Groupware Applications Fail: Problems in Design and Evaluation," *Office: Technology and People* 4 (3), 245-64.
- Hafner, Katie (1996). *Where Wizards Stay Up Late: the Origins of the Internet*. New York: Simon and Schuster.
- Hannemayr, Gisle (2003). "The Internet As Hyperbole: A Critical Examination of Adoption Rates," *The Information Society* 19 (2), 111-21.
- Hanseth, Ole and Eric Monteiro (1998). *Understanding Information Infrastructure* (unfinished book manuscript), www.ifi.uio.no/~oleha/Publications/bok.html
- Hartmanis, Juris and Herbert Lin (1992). *Computing the future: a broader agenda for computer science and engineering*. Washington, D.C: National Academy Press.
- Hauben, Michael and Ronda Hauben (1997). *Netizens: On the History and Impact of Usenet and the Internet*. Los Alamitos, CA: IEEE Computer Society Press.
- Held, David, et al. (1999). *Global Transformations: Politics, Economics and Culture*. Stanford: Stanford University Press.
- Hughes, Thomas P. (1983). *Networks of Power: Electrification in Western Society, 1880-1930*. Baltimore: Johns Hopkins University Press.

- Jackson, Steven and Karen Baker (2004). "Ecological Design, Collaborative Care, and the Ocean Informatics Project," in *Artful Integrations: Proceedings of the 2004 Participatory Design Conference* (Toronto).
- Jackson, Steven (2005). *Building the Virtual River: Numbers, Models, and the Politics of Water in California*. Unpublished Ph.D. dissertation, Department of Communication, University of California, San Diego.
- Jackson, Steven (2006). "Water Models and Water Politics: Deliberative Design and Virtual Accountability," in *DG.O2006: Proceedings of the 2006 Digital Government Conference* (San Diego).
- John, Richard (1995). *Spreading the News: The American Postal System from Franklin to Morse*. Cambridge, MA. Harvard University Press.
- Kahin, Brian and Janet Abbate (1995). *Standards policy for information infrastructure*. Cambridge, Ma: MIT Press.
- Kahin, Brian and Dominique Foray, eds. (2006). *Advancing Knowledge and the Knowledge Economy*. Cambridge, MA, MIT Press.
- Kaijser, Arne, and Marika Hedin, eds. (1995). *Nordic energy systems: historical perspectives and current issues*. Canton, MA: Science History Publications/USA.
- Kaijser, Arne (2003). "Redirecting infrasystems towards sustainability," in Biel, Anders, Bengt Hansson, and Mona Martensson, eds., *Individual and Structural Determinants of Environmental Practice*. Williston, VT: Ashgate, 152-79.
- King, John (2006). "Modern Information Infrastructure in the Support of Distributed Collective Practice in Transport," *Computer Supported Cooperative Work* 15 (2/3), 111-121.
- La Porte, Todd R., ed. (1991). *Social Responses to Large Technical Systems: Control or Adaptation*, NATO ASI Series D, Behavioural and social sciences, vol. 58. Boston: Kluwer Academic Publishers.
- Lee, Christopher (2005). *Defining Digital Preservation Work: A Case Study of the Development of the Reference Model for an Open Archival Information System*. Unpublished Ph.D. dissertation, School of Information, University of Michigan.
- Libicki, Martin (1995). "Standards: The Rough Road to the Common Byte," in Kahin, Brian and Janet Abbate, eds., *Standards Policy for Information Infrastructure*. Cambridge, MA: MIT Press, 35-78.
- Massey, Doreen (1993). "Power-Geometry and a Progressive Sense of Place," in J. Bird, B. Curtis, T. Putnam, G. Robertson, and L. Tickner, eds., *Mapping the Futures: Local Cultures, Global Change*. New York, Routledge.
- Mayntz, Renate and Thomas P. Hughes, eds. (1988). *The Development of Large Technical Systems*. Boulder, CO: Westview Press.
- McKenney, James L., Duncan G. Copeland, and Richard O. Mason (1995). *Waves of Change: Business Evolution through Information Technology*. Boston: Harvard Business School Press.
- Misa, Thomas, Philip Brey and Andrew Feenberg, eds. (2003). *Modernity and Technology*. Cambridge, MA, MIT Press.
- Misa, Thomas (2004). *Leonardo to the Internet: Technology and Culture from the Renaissance to the Present*. Baltimore, MD, Johns Hopkins University Press.
- Misa, Thomas (1998). *A Nation of Steel: The Making of Modern America 1865-1925*. Baltimore: Johns Hopkins University Press.

- Olson, Gary M. and Judith S. Olson (2000). "Distance Matters," *Human-Computer Interaction* 15, 139-178.
- Orlikowski, W.J. (1996). "Improvising Organization Transformation Over Time: A Situated Change Perspective," *Information Systems Research* 7 (1), 63-92.
- Porter, T. M. (1986). *The Rise of Statistical Thinking, 1820-1900*. Princeton, NJ: Princeton University Press.
- Russell, A.L. (2006). "'Rough Consensus and Running Code' and the Internet-OSI Standards War," *IEEE Annals of the History of Computing* 28 (3), 48-61.
- Schmidt, Kjeld and Ina Wagner (2004). "Ordering Systems: Coordinative Practices and Artifacts in Architectural Design and Planning," *Computer Supported Cooperative Work* 13 (5-6), 349-408.
- Schot, Johan, Thomas J. Misa, and Ruth Oldenziel, eds. (2006). "Tensions of Europe: The Role of Technology in the Making of Europe" (special issue), *History and Technology*, 21(1).
- Serres, M. (1980). *Le Passage du Nord-Ouest*. Paris: Editions de Minuit.
- Serres, M. (1990). *Le contrat naturel*. Paris : Editions François Bourin.
- Star, S. L. and K. Ruhleder (1996). "Steps toward an ecology of infrastructure: Design and access for large information spaces," *Information Systems Research* 7(1): 111-134.
- Suchman, Lucy (2006). *Human-Machine Reconfigurations: Plans and Situated Actions, 2nd Edition*. Cambridge, Cambridge University Press.
- Summerton, Jane, ed. (1994). *Changing Large Technical Systems*. Boulder, CO: Westview Press.
- Summerton, Jane, and Boel Berner, eds. (2002). *Constructing Risk and Safety in Technological Practice*. New York: Routledge.
- Turner, Fred (2006). *From Counterculture to Cyberculture: Stewart Brand, the Whole Earth Network, and the Rise of Digital Utopianism*. Chicago: University of Chicago Press.
- Vandenbosch, B. and M.J. Ginzberg (1996). "Lotus Notes® and Collaboration: le plus ça change.," *Proceedings of the Twenty-Ninth Hawaii International Conference on System Sciences*, 61-71.
- van der Vleuten, Erik, and Arne Kaijser (2006). *Networking Europe: transnational infrastructures and the shaping of Europe, 1850-2000*, Sagamore Beach, MA: Science History Publications/USA.
- Wiener, N. (1951). *Cybernetics and Society*. New York: Executive Techniques.
- Yates, JoAnne (1989). *Control through Communication: The Rise of System in American Management*. Baltimore: Johns Hopkins University Press.
- Yates, JoAnne (2005). *Structuring the Information Age: Life Insurance and Technology in the Twentieth Century*. Baltimore: Johns Hopkins University Press.

Appendix A. Conference Participants

Mark Ackerman
Associate Professor, School of Information and Department of Computer Engineering (EECS)
University of Michigan

Paul Avery
Professor, Department of Physics
University of Florida

Archer Batcheller
Doctoral Student, School of Information
University of Michigan

Fran Berman
Professor, Department of Computer Science and Engineering
University of California – San Diego

Christine Borgman
Professor, Department of Information Studies
University of California – Los Angeles

Geoffrey Bowker (co-PI)
Professor and Executive Director, Center for Science, Technology, and Society
Santa Clara University

Michael Buckland
Professor Emeritus, School of Information
University of California - Berkeley

Cecilia Deluca
Team Lead, Earth System Modeling Infrastructure Group
National Center for Atmospheric Research

Paul Duguid
Professorial Research Fellow
Queen Mary, University of London

Bill Dutton
Director, Oxford Internet Institute
Oxford University

Paul Edwards (PI)
Associate Professor, School of Information
University of Michigan

Tineke Egyedi
Senior Researcher, ICT Department, Delft University of Technology
President, European Academy of Standardisation

Tom Finholt
Associate Dean for Research and Innovation, School of Information
University of Michigan

Ian Foster
Professor and Director, Computation Institute
University of Chicago and Argonne National Laboratories

Jane Fountain
Professor and Director, National Center for Digital Government
University of Massachusetts - Amherst

Suzi Iacono
Acting Division Director for Information and Intelligent Systems
Directorate for Computer and Information Science and Engineering (CISE)
National Science Foundation

Steven Jackson (co-PI)
Assistant Professor, School of Information
University of Michigan

Robert Kahn
Chairman, CEO, and President
Corporation for National Research Initiatives (CNRI)

John King
Vice Provost for Academic Information
University of Michigan

Cory Knobel
Doctoral Student, School of Information
University of Michigan

Terry Langendoen
Coordinator for Cyberinfrastructure
Division of Behavioral and Cognitive Sciences
National Science Foundation

Christopher A. Lee
Assistant Professor, School of Information and Library Science
University of North Carolina – Chapel Hill

Kalle Lyytinen
Professor, Weatherhead School of Management
Case Western Reserve University

Tom Misa
Director, Charles Babbage Institute
University of Minnesota

Gary Olson
Professor, School of Information
University of Michigan

David Ribes
Post-Doctoral Fellow, School of Information
University of Michigan

Richard Rood
Professor, Department of Atmospheric, Oceanic, and Space Sciences
University of Michigan

Kjeld Schmidt
Associate Professor, Design of Organizational IT (DOIT) Group
IT University of Copenhagen

Johan Schot
Professor, Social History of Technology
Eindhoven University of Technology

Leigh Star
Professor, Center for Science, Technology, and Society
Santa Clara University

Jane Summerton
Professor, Department of Technology and Social Change
Linköping University

Fred Turner
Assistant Professor, Department of Communication
Stanford University

Erik van der Vleuten
Assistant Professor, History of Technology Department
Eindhoven University of Technology

Kevin Walsh
Technical Staff and PhD Student
San Diego Supercomputing Center and University of California – San Diego

Mark Weiss
Senior Science Advisor
Directorate for Social, Behavioral, and Economic Sciences
National Science Foundation

JoAnne Yates
Professor, Sloan School of Management
Massachusetts Institute of Technology