Information and Institutional Change: The Case of Digital Libraries

Philip E. Agre
Department of Information Studies
University of California, Los Angeles
Los Angeles, California 90095-1520
USA

pagre@ucla.edu
http://polaris.gseis.ucla.edu/pagre/

This is a chapter in Ann P. Bishop, Barbara P. Buttenfield, and Nancy Van House, eds, *Digital Library Use: Social Practice in Design and Evaluation*, MIT Press, 2003.

7000 words.

1 Introduction

Is a digital library a machine or an institution? The term has been used both ways (Borgman 1999). Computer scientists are willing to build a database and simply call it a "library". But whatever they are called, society will evaluate digital libraries in terms of the ways that they fit, or fail to fit, into the institutional world around them (Kling and Elliott 1994). Institutions, for present purposes, are the enduring categories of society: social roles, legal systems, linguistic forms, technical standards, and all of the other components of the playing field upon which human relationships are conducted (Commons 1924, Goodin 1996, March and Olsen 1989, North 1990, Powell and DiMaggio 1991). An institutional field is a particular relatively enduring ensemble of institutional categories; examples include particular historical forms of the market, the political system, or the university.

In this chapter, I propose to explore certain ways in which the "fit" between technology and institutions might be conceptualized and evaluated. I cannot survey all of the technical issues, or all of the institutional issues. Instead, I will focus on

the boundary between technology and institutions. Investigation of this boundary requires considerable preparatory analysis, and it will be necessary to identify and transcend several intellectual traps. In particular, evaluation of the interaction between technology and institutions must commence with substantive ideas on both sides: if we derive our institutional ideas from the metaphors already embedded in the technology, then the process may be tautological from the start. Every technology is embedded in the social world in complicated ways, and this is particularly true for digital libraries, which are intertwined with the cognitive processes of a complex society. Unless our conceptualization of society stands on an equal footing with our conceptualization of the technology it uses, our analysis will inevitably be overwhelmed by myths.

Digital libraries are distinctive in another way. The library world, like any institutional field, maintains a distinct identity. But the library world also articulates with other institutional fields, that is, it interacts with them in relatively stable and structured ways. In fact, the library world (along with some others, such as the legal system, higher education, the legislature, and journalism) articulates with virtually every other institution in society. This is the central institutional tension of the library: its need to maintain relatively uniform practices despite the great diversity of the social worlds whose members it serves (Agre 1995). Although it has never been easy, librarians have historically be able to maintain this uniformity of practices because of the limited variety of physical media (books, films, records, magazines) that various social institutions have produced. Networked computing, however, permits libraries to be articulated much more elaborately with the institutions they serve. Will the "common coin" of digital representation enable libraries to maintain a manageable uniformity in their practices despite the diversity of their articulations with other institutions? Or will be immense flexibility of digital computation unleash unmanageable pressures for heterogeneity?

My approach in this chapter is analytical, not empirical, and I will proceed as follows.

Section 2 will argue for the "priority of analysis": the idea that sociological conceptualization of user communities and institutions is logically prior to the design and evaluation of technical systems. When the priority of analysis is not respected, an intellectual vacuum opens up and various patterned cultural

myths flow in to colonize our thinking.

Section 3 considers the case of scholarship, with particular attention to the institutional conditions for the construction of healthy scholarly communities. In both cases, conventional theories ill-serve us by dividing the world into extremes when the reality, both descriptively and normatively, falls in the middle.

Section 4 considers the case of the public sphere. The health of a democratic society is founded, at least partly, in the pervasive processes of collective cognition that tie innumerable overlapping subcommunities. The conditions of these processes can be usefully compared and contrasted with those already analyzed in the world of scholarship.

Section 5 builds on these cases by considering more abstractly the embedding of digital libraries in their institutional environments. I will sketch several potentially useful themes, including convergence, specialization, standards dynamics, organizational boundaries, and genres.

Section 6 concludes by summarizing some of the positive contributions by which digital libraries might be evaluated.

2 Priority of analysis

The design of computer systems begins with concepts: concepts that describe the people, places, and things that the computer is supposed to represent, the attributes they can possess, the actions they can take, and the actions that can be taken upon them. The concepts that become embodied in computers are part of intellectual history: they come from somewhere, and indeed the usefulness of the computer will consist largely in the accuracy with which the users' concepts can be used to explain what the computer does. When designing a computer to predict the weather, for example, most of the relevant concepts are derived from meteorology. Because the concepts of meteorology are already stable and codified, the design process has a clear starting-point.

When designing a digital library, on the other hand, one must comprehend social phenomena of great complexity. No single discipline will provide all of the necessary concepts. Instead, it is necessary to employ concepts on several distinct levels of analysis. One level of analysis pertains to the physical and

cognitive mechanics of work; on this level the necessary concepts derive from ergonomics and human-computer interaction. Another level of analysis pertains to the principled organization of information and the search habits of individual library users, and on this level the necessary concepts can be obtained in reasonably stable and codified form from the tradition of library and information science (Borgman in press).

In this chapter, I will be concerned principally with an even higher level of analysis -- the embedding of a digital library in the larger social world -- for which the necessary concepts derive from social theory. The design of technical systems and institutions has not usually been informed by concepts from social theory, however, and so in this section I will consider the role of social theory in design.

A central challenge for social theory is the great complexity of social phenomena. No single concept will explain everything. The social theorist is therefore necessarily engaged in traffic control: working consciously with the relationships among a large number of concepts. For example, useful concepts are found on several different scales, or levels, and I have already informally sketched some of the levels that are relevant to the analysis of digital libraries. Each level of analysis is equally important, and analyses on the different levels will regularly inform one another. The concepts themselves differ from the concepts of science: they cannot be defined in mathematical terms. Their purpose, rather, is to help describe particular examples of social practice. These descriptions are necessarily intricate, and one does not expect to derive simple generalizations from them. The purpose of theoretical work, for example in this chapter, is to clarify concepts and their relationships. Whether the concepts are useful in the analysis of particular empirical situations is a different question, and each project -- the theoretical and the empirical -- provides important guidance to the other.

Concepts play at least three roles in design:

(1) They are employed in studying the task and the context of use.

(2) They are inscribed directly into the software, and into the categories and policies of the institution.

(3) They define the criteria by which the technical and

institutional systems are evaluated.

It follows that the analysis of concepts should precede design, or at least that conceptual analysis is necessary for design to make any progress. This is the *priority of analysis*. Although the point may seem like common sense, few design projects make any explicit provision for this kind of conceptual work before making irreversible design commitments. The traditional methods of systems design do employ the word "analysis", but in a narrower sense, in which the concepts are assumed to be given in advance. This is a dangerous assumption: without analysis, designers must necessarily employ whatever concepts they find lying around. These concepts might be incomplete or incoherent; they might distort the practices or omit large parts of them. Concepts that derive from the millenarian ideologies of computerization movements (Kling and Iacono 1988) are likely to be misleading as well. A design process that does not analyze its concepts forecloses much of the design space before it even begins, and it risks catastrophe if its concepts are broken. And because the same potentially problematic concepts are used in evaluating the system that is designed, fundamental design flaws will not necessarily be detected.

The design of digital libraries requires conceptual analysis because of the great complexity of a library's relationship to its institutional context. Advanced computing and broadband networking will enable digital libraries to become highly integrated with the institutions in which they are used, but little is known about the forms that this integration might take. A digital library does not require its users to extract themselves from their ongoing patterns of activity. To the contrary, the library can conform itself to those patterns of activity in numerous ways. Thoughtful design will require substantive ideas about those patterns of activity, and about what it might mean for a digital library -- or anything else -- to "fit" within them.

Conceptual analysis faces other challenges as well. Information technology often requires designers to revisit and clarify old concepts so that design thinking does not fall into simple dichotomies. What, for example, is a library? A concept of "library" that is too fully rooted in past historical forms will make innovation impossible, but a superficial concept of "library" that draws out only a few aspects of those past historical forms -- for example, the library as a big container of documents -- will pass over phenomena whose absence in newly designed system

may be fatal.

The middle ground between the maximal and simplistic conceptions of "library" is enormous, and it is not easily mapped. In mapping that ground, it helps to have two kinds of concepts. *Bridging concepts* are concepts that enable designers to move back and forth between the technical and institutional sides of their work. An example of a bridging concept is "inscription": the process by which social discourses are translated into the workings of software (Agre 1998a). And *meso-level concepts* are concepts that describe medium-sized social phenomena, for example "institutions" and "social networks", thus avoiding the sterile opposition between macro and micro that frustrates many applications of social theory (e.g., O'Neil 1998: 10).

Much of the skill of conceptual analysis consists in watching out for common traps that can confine a project's concepts within the bounds of unnecessary assumptions. Here, for example, are several conceptual traps that may afflict the unwary designer of digital libraries:

(1) The trap of presupposing standardization. Fantasies about computers in popular culture often assume an implausibly high level of interoperability among systems that have arisen independently of one another. This is certain to be a substantial issue as digital libraries are integrated with the systems of their diverse users, and the effort that goes into technical implementation in a narrow sense may be slight in comparison to the effort of consensus-building around standards.

(2) The trap of deriving political consequences straight from the technology. Authors such as Gilder (1992: 48-50, 126) have predicted that the decentralized nature of networked information technology would lead to a decentralization of power in society. But this consequence hardly follows. Computer networks are just as capable of projecting the instruments of control into far-flung locations. Likewise, librarians know well that uniform technical standards for access to digital libraries do not imply equal access in any effective social sense.

(3) The trap of automation. The word "automation" often slips back and forth between two distinct senses. In one sense, the word simply refers to any use of technology. But in another sense, it refers to a particular way of designing and using

technology, whereby the workings of a machine are modeled on the activities involved in a particular job, and the purpose of the machine is to replace the human effort that the job involves. When the word does slip silently between these two meanings, the design process can be led to presuppose the narrower sense of the term rather than consciously choosing it. It is sometimes both practical and beneficial to replace a human job with a machine on a one-for-one basis, but the possibilities of technology are vastly larger. In most cases, a new technology will lead to a renegotiation of the roles of people and machines, and this renegotiation should be part of the design process.

(4) The trap of assuming rapid change. The capacities of computer chips and fiber optic cables are growing rapidly, but it does not follow that social institutions will change as fast, or that they *can* change that fast, or that they should. Institutions become intertwined in with information technologies in many ways (Kling and Iacono 1989). Technical standards, once entrenched in the installed base and practices of an institution, are exceedingly difficult to change. The institutions themselves, as carriers of collective memory and skill, are usually slow to change as well, and for good reason. A design process that assumes rapid change will become preoccupied with "keeping up", and with "not being left behind", and will therefore not perceive the need for sober analysis of concepts -- not least one's concepts about change itself.

(5) The trap of all-or-nothing change. Many highly developed discourses presuppose that computing will give rise to total, discontinuous changes, either in society generally or in particular institutions. (For an example in higher education see Dolence and Norris (1995).) The world is thus divided, in Manichean fashion, into revolutionaries and reactionaries: those who embrace change and those who resist it. This kind of opposition is understandable in the absence of analysis: exactly because the changes are new, society has not needed concepts to describe the relationship between what is changing and what is not. Describing that relationship is a central role of conceptual analysis in the design process. Visions of discontinuous change can also arise from an over-simple understanding of the relationship between technology and sociology. Designers often associate technology with the future and sociology with the present and past. Because technology is supposed to change things, conceptually sophisticated investigation of the social world can seem irrelevant to design. But this understanding of

design is simplistic and even dangerous: it amounts to a willful blindness to the context in which the designed systems will be used.

(6) The trap of command-and-control computing. The main tradition of computer system design arose in military and industrial contexts in which the designers were closely allied with authorities who possess great power to direct the activities of the users. As a result, computers have long been associated with rationalization and hierarchy. The rise of the Internet, with its decentralized control structures, has shaken this association somewhat, but many legacies of the command-and-control era remain. For example, database design still assumes that individuals will be assigned unique identifiers, despite the serious privacy problems that this practice can raise. A new generation of privacy-enhancing technologies (Agre and Rotenberg 1997) has not been integrated into day-to-day design practice.

(7) The trap of inventing a new world. Several of these traps can combine to persuade designers that they can use technology to impose entirely new patterns of activity on their users. After all, the purpose of design is innovation, and true technical innovation is impossible unless the users change their habits. While it does make sense to speak of institutional design (Goodin 1996), new institutional structures usually cannot be imposed through technology. Existing patterns of activity are usually shaped by many factors beyond technology. Designers can consciously choose to amplify an existing force in society, but they probably cannot create new forces.

(8) The trap of blaming "resistance". Some technologies are rapidly adopted, and others are not adopted at all. The difficulty of predicting adoption can frustrate designers, and the language of "resistance" provides a simple explanation of the problem. But a responsible designer will try to distinguish between resistance that is irrational and resistance that arises from a poorly designed system.

(9) The trap of assuming away intermediaries. Networked information technology is frequently held to eliminate the need for intermediaries: those individuals and organizations that facilitate connections between buyers and sellers, citizens and government, people and information, and so on. After all, if the network can connect the parties directly, what is the purpose of

the intermediary? This argument depends on an ambiguity in words like "connect". A computer network can transport data between point A and point B, and it can make information available in a standard format at many points A so that computers located at various points B can search it. But intermediaries can serve many other purposes, and most of the successful new businesses on the Internet are in fact intermediaries (Sarkar, Butler, and Steinfield 1995; Shapiro 1999). A more suitable term -- and a less constraining concept -- is reintermediation (Halper 1998, Negroponte 1997).

(10) The trap of technology- and economics-driven scenarios. Institutions must be described using vocabularies from several disciplines, and great havoc can result when one discipline's language is employed to the exclusion of all others. As this list should already have made clear, design is too often technology-driven: given a hammer, one looks at things as if they were nails. Nonetheless, economics-driven design is equally hazardous. Economics is a powerful mode of analysis. But economic theories simplify and idealize the world. The mainstream neoclassical theories, for example, almost entirely ignore inormation and institutions (Casson 1997, Hodgson 1988). Economic theories have also tended to homogenize things by treating them as a uniform array of resources to be allocated, and they have likewise tended to oversimplify the web of human relationships within which economic exchange takes place (Granovetter 1992). When economic analyses are turned into institutional prescriptions, these simplifications can become serious blind spots. Economic analysis is an increasingly important component of the design process, but it is only useful as part of a dialogue.

(11) The trap of designing for a limited range of cases. Much of the design process necessarily takes place far from the places where the resulting systems will be used. As a result, designers must depend on their own imaginations. Designers whose imaginations are shaped by experience with one setting, or one type of user, risk designing systems that discriminate against other settings or users (cf. Friedman and Nissenbaum 1996).

(12) The trap of presupposing transparency. Experts usually forget what it was like to be a beginner, and designers usually cannot imagine what it is like to confront their systems anew. A generation of user-interface design has developed great solicitude for the situation of the beginning user (Shneiderman

1998), but this work has been primarily ergonomic and cognitive in orientation. Digital libraries exemplify a new generation of systems that cross institutional boundaries (Friedman 1989), and little is known about the challenges that such systems present to the beginner. Designers can too easily assume that the user possesses the whole tacit worldview of the designers' own community, and analysis will be required to understand the ways in which differences in worldview can affect the assumptions that users bring to a system.

These, then, are twelve challenges for conceptual analysis as part of the design process. A good design will seem deceptively simple, precisely because these issues do not arise. A good set of analytical concepts will avoid these traps as well, and it will be useful to keep them in mind when evaluating the substantive discussions of the remainder of the chapter.

3 Scholarly community

The scholarly community is deeply intertwined with the library. Much of the library's contents are produced by scholars, and the structure of these materials reflects the institutional structure of scholarship. Much of a scholar's professional persona lives on the library's shelves. The success of scholarship depends on the health of scholarly communities, and digital libraries will participate in changes that can affect the health of these communities for better or worse. A central theme, then, will be the internal workings of these communities, and their linkages to the rest of the world.

Scholars need a space apart from the world. What is this space? It is not physically localized, and indeed stretches around the earth in the "invisible colleges" into which scholars form themselves (Crane 1972). Scholars' space is understood in many other ways: time to think, a private work space, access to books and journals, freedom from political pressure, the opportunity to try ideas that might not succeed, and so on. In particular, scholars need a space for the self-organizing mechanisms of their community. While popular authors and journalists make their living selling their writings in a straightforward market system, scholars cannot use market mechanisms to govern their work because their task is to produce public goods: ideas and discoveries that are difficult to buy and sell (Hallgren and McAdams 1997). This is the purpose of peer review and the informal assignment of credit to innovators (Latour and Woolgar

1986). These mechanisms obligate scholars to monitor one another's careers, for example by reading journals and through professional meetings and rumor networks, and scholars thereby have a powerful incentive to adopt new information and communication technologies.

Scholars' space is also a container for conflict. In his astonishing sociological history of philosophy, Collins (1997) has demonstrated that the intellectual health of philosophy has depended crucially on robust debate between scholarly movements or schools. When the institutional conditions are present to support scholarly work, Collins argues, philosophical schools strategically split and merge so that only a few schools compete in the intellectual "attention space". The debates among these schools keep them honest, and the need to respond to opposing schools' arguments is the motive force that moves philosophical inquiry forward. When the institutional conditions of orderly debate fail, for example through economic collapse or political controls, philosophical inquiry becomes rigid or fragmented.

New technologies are quite capable of affecting the system of incentives that makes these mechanisms work. If scholars can advance in their careers by leading coherent intellectual movements, then technology can make it easier to organize such movements by maintaining communications among their members. The Advanced Research Projects Agency (ARPA), which invented the Internet, has also pioneered methods for using the Internet to operate research communities. Indeed, rather than leaving the creation of such communities to the career agendas of individual researchers, ARPA has largely internalized the process, using its own research-funding procedures to organize technical communities (e.g., the Image Understanding community) that are defined in relation to ARPA's own needs. Technology can also make it easier for scholars to move from one community to another, staying long enough to apply their skills to particular problems and then moving along to another community when a given line of work yields diminishing marginal returns. These factors may have contributed to the tremendous growth of interdisciplinary research during the 1980's and 1990's: not just the crossing of disciplinary boundaries, but the continual creation of new communities with new, permeable boundaries.

Assuming that Collins' argument applies to disciplines besides

philosophy, what does it say about the current situation? The increased fluidity of scholarly communities may be a sign of fragmentation. Or perhaps the attention space can now accommodate a larger number of schools. Collins' argument also points to the need for a loose coupling between individual research communities and the rest of the world. On one hand, a community must be sufficiently coherent to define a common language and a shared set of problems, methods, and goals. Otherwise it could not contain its own internal conflicts or develop any depth of learning. On the other hand, each community must be accountable to the arguments and objections of other communities. If this accountability becomes too rigid, then research will devolve into a purely political struggle. But if it becomes too loose, then misguided research communities will be able to reproduce themselves indefinitely. This is fundamentally a question of institutional design, but technology can either hinder or facilitate any institutional mechanism.

A loose analogy might be drawn between scholarly communities and ecosystems. Island ecologies permit evolution to head off in new directions without being constrained by the competitors in other ecologies. Likewise, the intellectual communities of different countries have historically been somewhat isolated from one another. Scholars have always corresponded and traveled, not least because they can advance professionally by bringing new ideas home with them. But communications and travel have always been laborious and expensive, and the interactions among national scholarly communities have always been limited as a result. Language has been a barrier as well. Nonetheless, new technologies decrease the costs of scholarly interaction, so that now it is necessary to determine the optimum level of coupling among different intellectual ecosystems. The dangers of excessive technologically-facilitated homogeneity can be seen in the computer industry, which has gone through a long series of manias during the 1990's: virtual reality, agents, network computers, and so on. Institutional changes will be required if these wasteful storms of intellectual fashion begin to disrupt scholarship.

Technology might contribute to constructive changes in the institutions of scholarship through the invention of new genres (Agre 1998b). As a thought experiment, consider the problem of professional mobility: the possibility of advancing in one's career, either by building stature within a particular research

community or by moving to another one. Either kind of move requires an individual to master a complex landscape of scholars and their work. Where can maps of this landscape be found? One map is available in traditional library catalogs, but that map does not reflect many of the most important features of the territory (Agre 1995). Other maps are available in the narratives by which authors of scholarly literature give credit to the authors who have gone before them. And yet others are available in survey articles. A digital library might make these narratives more systematically available. Graduate students who are defining an area of dissertation research could be obliged to produce an extensive survey, as a structured hyperlinked document conforming to a particular XML document type, of the literature in that area. The scholarly community could use the Internet to organize a peer-review and publication system for these documents, and these mechanisms would help students to develop their professional voices. Once deposited in the digital library, the world's entire collection of literature surveys could then be searched by anyone wishing for an introduction to a given literature. Reverse links, from works to the surveys that mention them, would make available several narratives of a given work's place in intellectual history. Whether this proposal is feasible, of course, depends on much more than technology. It would require a great deal of consensus-building, and it would require individual research programs to surrender some degree of control over the progress of their students' careers. A digital library can be designed to support these kinds of technical mechanisms, and it can be evaluated in terms of the support it offers to the laborious process of building consensus around them. Increased professional mobility may have disadvantages as well. Students today are strongly bound to their dissertation advisors both by cognitive limits (it is hard to learn any other intellectual system than that of one's teacher) and by the mechanisms of professional evaluation (only one's teacher is in a position to write the evaluations upon which one's advancement depends). But if these bonds are loosened then it might become impossible to build a stable intellectual community.

These few ideas hardly exhaust the range of institutional linkages between a research community and the rest of the world. Other linkages include those between research and teaching, between scholarship and government, and between theoretical work and applications. A more careful treatment would consider these linkages systematically, revealing how

each linkage works in the present day and inquiring how a digital library might facilitate or disrupt the existing dynamics (cf. Lamb 1995). These considerations will return in the next section.

4 Public sphere

In addition to their role in supporting research, digital libraries can also be evaluated in terms of the contribution they make to the health of democracy. This would seem obvious enough: democracy is supposed to be a matter of rational deliberation, and a digital library ought to support the activities of research, reflection, and communication that rational deliberation requires. But much depends on one's conception of democracy, and of the cognitive processes that support it. Liberal political theory, for example, locates the practices of democracy in the individual: individual people gathering information, debating one another, and expressing their choices through aggregating mechanisms like voting. To the extent that library science conceptualizes library patrons as individuals, it embodies a liberal theory of politics. Given the epistemic and cognitive limitations of isolated individuals, however, library or no, such a theory cannot explain how citizens can effectively deliberate on matters that involve far-flung facts and affect the community as a whole. At the opposite extreme, authors like Volosinov (1973) interpret both politics and cognition entirely in collective terms, leaving no analytical space for the individual. This theory is no better, and (strikingly) for the same reason: collective cognition, to be effective, requires a substantial division of labor (cf. Hutchins 1996, Weick and Roberts 1993). Somewhere between these positions, communitarian authors imagine individual cognition and action to be constituted to a large extent by the norms and language of the community, but do not imagine that the community completely determines the individual's choices (Etzioni 1995). This is progress, but it provides no real theory of interests and conflict. What is needed, therefore, is a substantive account of the cognitive basis of social movements and other social groups (Melucci 1996), and of the ways that technologies and institutions can either support or disrupt this collective cognition.

One starting point might be the rough analogy between social groups and scholarly communities. Each type of community needs an autonomous space, loosely coupled to the spaces of other communities. Each provides its individual members with

relatively safe opportunities to develop their public voices. But the analogy stops there. Scholars need to be accountable in material terms for the coherence and utility of their ideas, so that the institutions of scholarship can allocate their resources in a productive way. Otherwise scholars would be paid to talk nonsense. No outside regulation is required to prevent a social movement in a democracy from talking nonsense, since the movement's ideas must be coherent enough to organize effective action and appealing enough to form the ideological basis for coalitions with other groups. Incoherent ideas can be exposed by other movements that compete to recruit the same social groups to its own coalition. The principal question, therefore, concerns the conditions under which different groups are able to organize themselves cognitively: unequal access to the means of collective cognition can lead to material inequalities of other sorts. Herein lies one of the central political claims for the Internet: online discussion groups provide cognitive infrastructure for a vast range of constituencies (Agre 1998b), and digital libraries seem certain to do the same. Buchstein (1997: 251) observes that "viewed in terms of contemporary democratic theory, the positive qualities attributed to the Internet strikingly resemble the Habermasian unrestricted public sphere". The public sphere is not singular but multiple, and "[t]echnologies of communication ... make possible a highly differentiated network of public spheres. ... The boundaries are porous; each public sphere is open to the other public spheres." (Habermas 1987: 359–360; cf. Fraser 1992). Several authors have even spoken of new communications technologies as providing the conditions for a collective intelligence, whether in organizations (Fisher and Fisher 1997, Smith 1994), on a societal level (Hayek 1948: 50–54) or globally (De Kerckhove 1997, Levy 1997, Rossman 1992, Wells 1938). But, just as obviously, technology does not provide all of the necessary conditions. Interest-group politics, notoriously, suffers from free-ridership (Olson 1965): group members who do not participate in developing an intelligent group consensus will nonetheless benefit from it, and technology does not automatically create the necessary incentives. Collective cognition requires a shared identity, social skills, and morale, each of which has conditions of its own.

What is more, technology has also raised the stakes by facilitating the explosion of "information-driven politics" (Greider 1992: 46) that has been accelerating since its origins in the open-government movements of the 1970's. The purpose of

"think tanks" is precisely to generate the steady stream of convenient facts, persuasive phrases, and finely-tuned ideologies that assemble coalitions around the agendas of their paying supporters. These organizations expose the great complexity buried beneath simple concepts such as the "marketplace of ideas" (e.g., Baker 1989, Ingber 1984). Ideas are public goods, and I have already mentioned the role of scholarly communities in alleviating the economic pathologies that public goods raise. But the "marketplace of ideas" is strange in another way: ideas in the public sphere are useful to me not because I "buy" them but because other people do. And that is the role of the think tank: selling one's ideas to others. Every group has an interest in influencing the thinking of every other group, for example through the public relations practice of providing "information subsidies" to the media (Gandy 1982), and so it can be extremely difficult for a social group to conduct its collective cognition autonomously (Habermas 1987). The problem is partially one of scale: a social group whose members are few in number but command great resources can organize its institutions of collective cognition more easily, other things being equal, than a group whose members are more numerous and less wealthy. Larger groups are easier to infiltrate and thus provide easier targets of surveillance, and several public relations firms now routinely monitor public Internet discussions, among other popular communications channels, on behalf of their clients. (See, for example, www.ewatch.com.) The problem of autonomy arises on the most basic level when provocateurs set about disrupting a community's cognitive institutions, and some online communities have developed sophisticated methods for maintaining their boundaries in the face of such attacks (Phillips 1996). New technologies can support the development of autonomous processes of collective cognition if they provide social groups with the tools to minimize these dangers, or at least to equalize them.

Finally, digital libraries bear on the relationship between the professions and the rest of society. This has been a crucial issue for democracy since the days of Lippmann (1922) and Dewey (1927). Although opinions differed on the extent of formal political power that should be invested in the experts, the elite consensus of that era was that nonetheless democracy should concede a great deal of cognitive authority to professions and their expertise (Schudson 1998: 211-219). Subsequent experience, however, has made clear that democracy requires an irreducible creative tension between professional and popular

voices. Digital libraries will presumably continue to facilitate the production and authorization of professional knowledge, but they also may also enable nonprofessionals to appropriate this knowledge in their own ways (cf. Blau 1999: 125–127). Once again, the conditions are largely institutional: now that it is technically possible to make professional knowledge accessible to the public, new incentives might be useful to encourage professionals to to make professional knowledge accessible in a fully effective sense. Digital libraries should also be evaluated for their capacity to support forms of collective cognition that differ from those of the traditional organized professions.

5 Institutional embedding

The previous sections have sketched a few of the ways in which a digital library might fit, or else fail to fit, into the institutional world around it. The discussion is necessarily schematic, and it will not be possible to offer any meaningful generalizations until digital libraries are being used on a large scale. Nonetheless, some general patterns can be anticipated. Most fundamentally, the design of digital libraries will require a dynamic approach: neither ignoring the institutional context nor trying to legislate it, but participating in the dialectical interaction between technology and institutions. Institutional processes shape technologies, and the technologies that result are then appropriated by the institutions' members. Experience with these appropriations helps to shape new generations of technology, which are appropriated in turn. These appropriations are famously unpredictable, but they can in fact be predicted to a certain degree: given an analysis of the existing forces in a given institutional field, one can safely say that those forces will shape the community's understandings of the technology and its potential uses.

The dialectical interaction between institutions and technology does not happen in isolation; quite the contrary, it is increasingly mediated by the global dynamics of technical standards (David and Shurmer 1996, Kahin and Abbate 1995). In the 1970's, much software was produced by organizations for their own use. But in the 1990's, the inherent economics of software has created tremendous forces away from bespoke applications and toward packaged software whose immense development costs can be spread across many different customers. As a result, few organizations determine their own fate. Even a whole institutional field, such as the libraries or the

educational system, can find itself hostage to global standards that emerge and develop a critical mass of users in other sectors. It is easy to speak of the design of digital libraries as if designers can freely choose their own directions, but in practice digital libraries emerge through negotiations in a tremendous variety of standards coalitions. Some of these coalitions are specific to libraries as an institutional field, but most are not. It follows that digital libraries can only be designed intelligently if their stakeholders join these negotiations (Oddy 1997: 83). Questions of institutional fit also arise in other design contexts, of course, and the substance of these standards negotiations will often pertain, explicitly or not, to fundamental ideas about institutions and the social relationships that they define. Despite their esoteric reputation, standards can very easily embody substantive commitments that shape and constrain people's activities (Reidenberg 1998), and they can bias a playing field toward some players and away from others (Mansell 1995).

Technology and institutions interact especially in regard to issues of centralization and decentralization. The library community has already gone a long way toward eliminating duplicate effort by pooling catalog records, and that experience can serve as a template for future issues of digital library governance. Centralization is also fundamental to the establishment of compatibility standards, inasmuch as standards require consensus that must usually be coordinated through some central body. In many cases centralized power is required to create the incentives for compatibility, but compatibility then creates the conditions for power to be decentralized. An institution field can easily become "stuck" with an overly centralized concentration of power, but it can just as easily become stuck at the opposite extreme when sufficient consensus cannot be established to adopt and implement new standards. These governance challenges are great enough when standards change slowly and in isolation from one another, but they become crucial when large number of standards are being developed and adopted simultaneously, as they are right now. In the worst case, the direction of digital library standards-setting could be captured by a single interest, for example a software vendor who can leverage a standard operating system, or else a coalition of intellectual property owners who can leverage their contractual control over digital library content. The potential for monopoly rent-extraction in that scenario is enormous, and so libraries will have to learn how to maintain their boundaries against such effects.

Digital libraries also face strong centrifugal forces. I have already mentioned the great diversity of institutional fields with which libraries interact, and each of these institutional fields is likely to have developed its standards and practices in relative isolation from the others. The technologies and policies of a digital library can be deeply integrated with any one of those neighboring fields, or with a few, but it will be hard to integrate with many of them. If the design of digital libraries is biased by the needs of a small number of powerful user groups (experts, for example, as opposed to lay persons), then they might discriminate against others. Or they may simply be pulled to pieces, with different digital libraries heading in different directions without being interoperable with one another. Managing these tensions will be a great institutional challenge. Digital libraries may also become the terrain upon which diverse institutions negotiate a common set of standards that facilitate activity in each area without artificial constraint.

6 Conclusion

This chapter has sketched some of the institutional problems with which the development of digital libraries must contend. It has also made clear that librarians, far from being automated into nonexistence by new technology, retain a considerable role in ensuring that libraries continue to encourage these values (cf. Nardi and O'Day 1996). This role is centrally one of design -- not the command-and-control style of design from which computers first emerged, but a participatory style in which the well-being of social institutions and their participants cannot be separated from the construction of technical systems. This new style of design thus involves leadership skills of a high order. But it also involves analytical skills, and I hope to have demonstrated the role of social theory in the practical work of designing digital libraries that can be truly useful in a complicated world. Among the contributions of social theory has been a clear sense in which a library, even when it is digital, is still a place: the place where a scholarly community or a social movement can conduct its collective cognition with a reasonable degree of autonomy. We still know little about the construction of such places, but perhaps we can renew our appreciation of the need for them.

* Acknowledgements

I appreciate helpful comments by Nancy Van House.

* References

Philip E. Agre, Institutional circuitry: Thinking about the forms and uses of information, *Information Technology and Libraries* 14(4), 1995, pages 225-230.

Philip E. Agre, The Internet and public discourse, *First Monday* 3(3), 1998a.

Philip E. Agre, Designing genres for new media, in Steve Jones, ed, *CyberSociety 2.0: Revisiting CMC and Community*, Sage, 1998b.

Philip E. Agre and Marc Rotenberg, eds, *Technology and Privacy: The New Landscape*, Cambridge: MIT Press, 1997.

C. Edwin Baker, *Human Liberty and Freedom of Speech*, New York: Oxford University Press, 1989.

Andrew Blau, Floods don't build bridges: Rich networks, poor citizens and the role of public libraries, in Sally Criddle, Lorcan Dempsey, and Richard Heseltine, eds, *Information Landscapes for a Learning Society*, London: Library Association, 1999.

Christine L. Borgman, What are digital libraries? Competing visions, *Information Processing and Management* 38(3), 1999, pages 227-243.

Christine L. Borgman, *From Gutenberg to the Global Information Infrastructure: Access to Information in the Networked World*, Cambridge: MIT Press, in press.

Hubertus Buchstein, Bytes that bite: The Internet and deliberative democracy, *Constellations* 4(2), 1997, pages 248-263.

Mark Casson, Institutional economics and business history: A way forward?, *Business History* 39(4), 1997, pages 151-171.

Randall Collins, *The Sociology of Philosophies: A Global Theory of Intellectual Change*, Cambridge: Harvard University Press, 1998.

John R. Commons, *Legal Foundations of Capitalism*, New York: Macmillan, 1924.

Diana Crane, *Invisible Colleges: Diffusion of Knowledge in Scientific*

Communities

Paul A. David and Mark Shurmer, Formal standards-setting for global telecommunications and information services, *Telecommunications Policy* 20(10), 1996, pages 789-815.

Derrick De Kerckhove, *Planetary Mind: Collective Intelligence in the Digital Age*, Hardwired, 1997.

John Dewey, *The Public and Its Problems*, New York: Holt, 1927.

Michael G. Dolence and Donald M. Norris, *Transforming Higher Education: A Vision for Learning in the 21st Century*, Ann Arbor: Society for College and University Planning, 1995.

Amitai Etzioni, *New Communitarian Thinking: Persons, Virtues, Institutions, and Communities*, Charlottesville: University Press of Virginia, 1995.

Kimball Fisher and Maureen Duncan Fisher, *The Distributed Mind: Achieving High Performance Through the Collective Intelligence of Knowledge Work Teams*, AMACOM, 1997.

Nancy Fraser, Rethinking the public sphere: A contribution to the critique of actually existing democracy, in Craig Calhoun, ed, *Habermas and the Public Sphere*, Cambridge: MIT Press, 1992.

Andrew L. Friedman, *Computer Systems Development: History, Organization and Implementation*, Chichester, UK: Wiley, 1989.

Batya Friedman and Helen Nissenbaum, Bias in computer systems, *ACM Transactions on Information Systems* 14(3), 1996, pages 330-347.

Oscar H. Gandy, Jr., *Beyond Agenda Setting: Information Subsidies and Public Policy*, Norwood, NJ: Ablex, 1982.

George Gilder, *Life After Television*, New York: Norton, 1992.

Robert E. Goodin, ed, *The Theory of Institutional Design*, Cambridge: Cambridge University Press, 1996.

Mark Granovetter, Economic action and social structure: The problem of embeddedness, in Mark Granovetter and Richard Swedberg, eds, *The Sociology of Economic Life*, Boulder: Westview, 1992.

Who Will Tell the People? The Betrayal of American Democracy, New York: Simon and Schuster, 1992.

Jurgen Habermas, *The Philosophical Discourse of Modernity: Twelve Lectures*, translated by Frederick Lawrence, Cambridge: MIT Press, 1987.

Martyne M. Hallgren and Alan K. McAdams, The economic efficiency of Internet public goods, in Lee W. McKnight and Joseph P. Bailey, eds, *Internet Economics*, Cambridge: MIT Press, 1997.

Mark Halper, Middlemania, *Business 2.0* 3(7), 1998, pages 45-60.

Friedrich A. Hayek, *Individualism and Economic Order*, Chicago: University of Chicago Press, 1948.

Geoffrey M. Hodgson, *Economics and Institutions: A Manifesto for a Modern Institutional Economics*, Cambridge, UK: Polity Press, 1988.

Edwin Hutchins, *Cognition in the Wild*, Cambridge: MIT Press, 1996.

Stanley Ingber, The marketplace of ideas: A legitimizing myth, *Duke Law Journal* 1984(1), 1984, pages 1-91.

Brian Kahin and Janet Abbate, eds, *Standards Policy for Information Infrastructure*, Cambridge: MIT Press, 1995.

Rob Kling and Margaret S. Elliott, Digital library design for usability, in John L. Schnase, John J. Leggett, Richard K. Futura, and Ted Metcalfe, eds, *Proceedings of Digital Libraries '94: The First Annual Conference on the Theory and Practice of Digital Libraries*, College Station, Texas, June 1994, pages 146-155.

Rob Kling and Suzanne Iacono, The mobilization of support for computerization: The role of computerization movements, *Social Problems* 35(3), 1988, pages 226-243.

Rob Kling and Suzanne Iacono, The institutional character of computerized information systems, *Office: Technology and People* 5(1), 1989, pages 7-28.

Roberta Lamb, Using online information resources: Reaching for the *.*'s, in Frank M. Shipman III, Richard K. Futura, and David

M. Levy, eds, *Proceedings of Digital Libraries '95: The Second Annual Conference on the Theory and Practice of Digital Libraries*, Austin, TX, June 1995, pages 137-146.

Bruno Latour and Steve Woolgar, *Laboratory Life: The Construction of Scientific Facts*, Princeton: Princeton University Press, 1986. Originally published in 1979.

Pierre Levy, *Collective Intelligence: Mankind's Emerging World in Cyberspace*, translated by Robert Bononno, Plenum Press, 1997.

Walter Lippmann, *Public Opinion*, New York: Macmillan, 1922.

Robin Mansell, Standards, industrial policy and innovation, in Richard Hawkins, Robin Mansell, and Jim Skea, eds, *Standards, Innovation and Competitiveness: The Politics and Economics of Standards in Natural and Technical Environments*, Aldershot, England: Edward Elgar, 1995.

James G. March and Johan P. Olsen, *Rediscovering Institutions: The Organizational Basis of Politics*, New York: Free Press, 1989.

Alberto Melucci, *Challenging Codes: Collective Action in the Information Age*, Cambridge University Press, 1996.

Bonnie Nardi and Vicki L. O'Day, Intelligent agents: What we learned at the library, *Libri* 46(2), 1996, pages 59-88.

Nicholas Negroponte, Reintermediated, *Wired* 5(9), 1997, page 208.

Douglass C. North, *Institutions, Institutional Change, and Economic Performance*, Cambridge: Cambridge University Press, 1990.

Pat Oddy, *Future Libraries, Future Catalogues*, London: Library Association, 1997.

Mancur Olson, Jr., *The Logic of Collective Action: Public Goods and the Theory of Groups*, Cambridge: Harvard University Press, 1965.

Patrick H. O'Neil, *Revolution from Within: The Hungarian Socialist Workers' Party and the Collapse of Communism*, Northampton, MA: Elgar, 1998.

Wanda J. Orlikowski, Learning from Notes: Organizational issues in groupware implementation, *The Information Society* 9(3), 1993,

pages 237–250.

David J. Phillips, Defending the boundaries: Identifying and countering threats in a Usenet newsgroup, *The Information Society* 12(1), 1996, pages 39–62.

Walter W. Powell and Paul J. DiMaggio, eds, *The New Institutionalism in Organizational Analysis*, Chicago: University of Chicago Press, 1991.

Joel Reidenberg, Lex Informatica: The formulation of information policy rules through technology, *Texas Law Review* 96, 1998, pages 553–593.

Parker Rossman, *The Emerging Worldwide Electronic University: Information Age Global Higher Education*, Westport, CT: Greenwood Press, 1992.

Mitra Barun Sarkar, Brian Butler, and Charles Steinfield, Intermediaries and cybermediaries: A continuing role for mediating players in the electronic marketplace, *Journal of Computer-Mediated Communication* 1(3), 1995.

Michael Schudson, *The Good Citizen: A History of American Civic Life*, New York: Free Press, 1998.

Andrew L. Shapiro, *The Control Revolution: How The Internet is Putting Individuals in Charge and Changing the World We Know*, Public Affairs, 1999.

Ben Shneiderman, *Designing the User Interface: Strategies for Effective Human-Computer Interaction*, third edition, Reading, MA: Addison Wesley, 1998.

John B. Smith, *Collective Intelligence in Computer-Based Collaboration*, Hillsdale, NJ: Erlbaum, 1994.

Stephen Toulmin, *Human Understanding: The Collective Use and Evolution of Concepts*, Princeton University Press, 1972.

V. N. Volosinov, *Marxism and the Philosophy of Language*, translated by Ladislav Matejka and I. R. Titunik, Cambridge: Harvard University Press, 1973.

Karl E. Weick and Karlene H. Roberts, Collective mind in organizations: Heedful interrelating on flight decks, *Administrative Science Quarterly* 38, 1993, pages 357–381.

H. G. Wells, *World Brain*, Garden City, NY: Doubleday, Doran, 1938.