
Going Digital: *A Look at Assumptions Underlying Digital Libraries*

David M. Levy and Catherine C. Marshal

What are digital libraries, how should they be designed, how will they be used, and what relationship will they bear to what we now call “libraries”? Although we cannot hope to answer all these crucial questions in this short article, we do hope to encourage, and in some small measure to shape, the dialog among computer scientists, librarians, and other interested parties out of which answers may arise. Our contribution here is to make explicit, and to question, certain assumptions that underlie current digital library efforts. We will argue that current efforts are limited by a largely unexamined



and unintended allegiance to an idealized view of what libraries have been, rather than what they actually are or could be. Since these limits come from current ways of thinking about the problem, rather than being inherent in the technology or in social practice, expanding our conception of digital libraries should serve to expand the scope and the utility of development efforts.

We will take a work-oriented perspective [5,7] on libraries—a perspective from which observations of people’s actual work is used to critique and guide technology development. We begin with a simple framework, derived from Yates [17] which highlights three crucial aspects of libraries (see Figure 1).

Documents: While there are many types of libraries—including national, corporate, community, and school libraries—all libraries seem to share at least one property: they house and provide access to collections of documents [13]. We use the term “document” in the broadest sense, to denote all enduring communicative records, including paper materials, electronic files, videotapes, and audiotapes.

Technology: Documents are communicative artifacts, and as *artifacts* (literally, “made with skill”), they inevitably require technology to be created and maintained. Handwritten grocery lists, printed books, and movies on videotape are all based on technologies that have required huge investments of human intelligence and labor to be realized. Traditionally, libraries have been principally associated with the technologies of the book—paper and print—but they also make use of many others as well, such as microform and audiotapes.

Work: Were we to talk only about collections of documents and the technologies through which they are realized, we would miss the most fundamental element of all—the actual work done by library users, as well as the work done by library personnel in support of them. We might call these two types of activities *research* and *service*, respectively. Without understanding the forms of research toward which a library is oriented, we could not understand how its collections are selected and organized; and without understanding its internal services, we could not understand how collections are maintained or how users’ research is supported [3].

This framework places work on a par with documents and technology as a legitimate domain of investigation and source of innovation. By a work-oriented approach to digital libraries, we do not mean an approach that focuses solely on people’s work, but rather one that evaluates library collections and technology in relation to the work that is being done with them.

Assumptions About Digital Libraries

What are digital libraries? In what ways will they be like today’s libraries—and in what ways unlike them? We believe that the current conception of digital libraries is derived in large measure from a widely held idealization of libraries, which might be called “the traditional library,” or perhaps just “the library.” This idealization is not so much wrong as incomplete—a simplified and simplistic image of the local and academic libraries most of us have used throughout our lives. By failing to question this idealization or to acknowledge the ways in which the current conception of digital libraries is based on it, we risk creating digital libraries that are unnecessarily limited, and, in the worst case, entirely fanciful and unusable.

We will next examine three characteristics digital libraries are assumed to have (Figure 2)—characteristics of their collections, of their technologies, and finally of the work they are meant to enable.

1. Documents: *Digital library collections contain fixed, permanent documents.*

Two important characteristics of documents and collections of documents [9] are:

- **Rate of change:** Over a given period of time, a document may be fixed (unchanging) or fluid (changing). Printed books are typically fixed during their useful lifetime, while address books and personal calendars are more fluid.
- **Duration:** A document may have a long or short useful lifetime—that is, it may be permanent (of long duration) or transient (of short duration). Books typically have a long lifetime (quasi-permanent) while Post-it® notes have a short lifetime (highly ephemeral).

Although often conflated, these are independent dimensions. Thus, a Post-it note may be unchanging during its very short lifetime, while a favorite cookbook may be subject to much change (i.e., through annotation) over many years of use.

Our idealized image of a library imbues it with




Figure 1. Three perspectives on libraries: the collections of documents they house, their enabling technologies, and the work conducted through, in, and by libraries.

qualities of fixity and permanence [15]. This is hardly surprising, since the library is considered to be the Home of the Book, and books are by and large one of the more fixed, more permanent types of documents. This association seems to have been carried straightforwardly and unreflectively into specifications for digital libraries. Many proposed or actual projects are oriented toward the management of relatively fixed, relatively permanent collections (e.g. Cornell's Class project [1] or ARPA's CSTR program¹). Even Carnegie-Mellon University's Informedia Digital Video Library project (see Stevens et al. in this issue), which is concerned with time-based media, appears to be oriented toward the access and management of relatively unchanging, archival materials.

But the reality of current library collections stands in marked contrast to this idealization. Libraries have always contained materials other than books. Special collections and archives are filled with unbound and handwritten ephemera—correspondence, photographs, and so on. Moreover, nothing in the nature of digital technologies mandates that digital libraries should include only rarely changing, long-lasting documents. On the contrary, the ease of modification afforded by digital technologies means there are, and will be, whole classes of digital documents that do not fit the traditional profile for library inclusion. To what extent do we want to consider collections of *listserv* messages, wire service articles, preprints, and other quickly changing and/or ephemeral documents as appropriate materials for digital libraries?

It is also worth noting that assumptions of fixity and permanence are likely to creep unseen into other corners of the digital library apparatus, even where people are explicitly trying to deal with fluid, ephemeral materials. Thus, current library cataloging practices are very much directed toward slowly changing materials and do not deal well with versions or custom documents. Current attempts to “catalog the Net,” for example, are largely based on these existing practices and are unlikely to be sufficient for more fluid, transient materials [10].

2. Technology: Digital libraries are based on digital technologies.

Nothing would seem to be more self-evidently true than that digital libraries will, and should, be based on digital technologies. It is an unquestioned assumption that digital libraries will contain digital

materials, just as traditional libraries have contained books—that is, materials based in the technologies of paper and print. But, in fact, traditional libraries have long contained a diversity of technologies and media; today these include film and video, microfilm and microfiche, vellum and papyrus. Why should we expect digital libraries to be any less heterogeneous?

The tacit assumption is that digital libraries will contain *only* digital material. There are two ways in which this could come to be true: if digital-only libraries came to exist independently from nondigital libraries or if all libraries came to be digital. Today, given the huge amounts of paper in use, the only way to create digital-only libraries would be to construct them specifically to exclude nondigital materials. In the short run, this would mean that digital-only and nondigital libraries would coexist, side by side, in effect. But in the long run, all libraries could be digital-only if the vast archives of paper documents were digitized and all newly created documents were in digital form.

We question both of these scenarios. While no one can actually predict future paper use, it seems unlikely that paper—so flexible, portable, inexpensive, and easily annotated—will simply go away.


It is unimaginable that the entire contents of our libraries and archives will ever be digitized—not just because there is too much paper, but because libraries contain materials that do not naturally lend themselves to digital representation, such as the gifts from heads of state to be housed in the George Bush Presidential Library. (The boundaries between libraries, museums, and archives, although intuitively clear, are not so easy to draw in practice.)

We are then left with two possible futures. One is that digital-only libraries and nondigital libraries will coexist. The other is that libraries will contain digital *and* nondigital material—in which case “digital library” is a misnomer, and certainly not synonymous with “the library of the future.” In either event, we can expect we will continue to deal with a broad range of heterogeneous materials, including those outside computational reach. A narrow focus on digital technologies is unlikely to be sufficient.

3. Work: Digital libraries are to be used by individuals working alone.

What about the work that users of libraries do (research) and the work that librarians do to support their users (service)? In both of these realms the preva-

Figure 2. Three assumptions underlying current digital library efforts.



¹ <http://cs-tr.cs.cornell.edu/info/cstr.html>. ARPA-sponsored CSTR project

lent image is of solitary work—the scholar cloistered with his or her materials, the librarian walking the labyrinthine stacks. Either implicitly or explicitly, much of the current work on digital libraries assumes this idealized model of use: the lone researcher sitting at a workstation, browsing, scanning, searching, retrieving, reading, and writing. Yet this idealization is at odds with observed work practice in a range of settings (e.g., [6, 12]). Libraries are meeting places where joint research is carried out; research is a highly col-

lection and its use. But even where the end user has not been forgotten, there may still be a tendency, as we have already suggested, to idealize users and uses, projecting or inventing an incomplete or even inaccurate picture of the real work being done.

A partial antidote to these tendencies is to take a work-oriented perspective: to focus directly on the user community and to use ethnographic techniques [2] to observe the work being done and the documents and technologies that support it. We turn now

The highest priority of a library, *digital or otherwise*, is to serve the research needs of its constituents.

laborative activity, even if this is not officially acknowledged. Service work is highly collaborative, too: collections are maintained through the joint efforts of library staff participating in a broader community.

Even information-seeking, the digital library activity apparently most consistent with the idealized image of solitary work, is more collaborative than generally realized. People seek information by communicating with members of their communities; they look not only for materials and specific answers, but for corroboration, new interpretations, and new methods of finding information. We have observed this in our own use of the Internet, where even the simplest of information-seeking activities, such as viewing a remote document whose address we already have, sometimes involves consultation with colleagues on site as well as email communication with distant, possibly unknown consultants. All this suggests that support for communication and collaboration is as important as support for information-seeking activities and that, indeed, support for the former is needed to support the latter. If the exploding use of the Internet has shown anything, it is that users want communication at least as much as they want information access [16].

A Work-Oriented Perspective: Information Analysts

The highest priority of a library, digital or otherwise, is to serve the research needs of its constituents. The development, maintenance, and extension of its collection and its technologies must be supportive of, and subordinate to, this primary objective. Still, this priority may at times be lost in the midst of more immediate, and apparently more pressing, tasks. Among librarians, there is sometimes a tendency to focus on collection acquisition and maintenance and to lose sight of the library's role in supporting the community's research. Among technologists, there is sometimes a tendency to focus on the details of technology implementation and maintenance, losing site of technology's instrumental role in supporting the

to consider one community of users, a group of information analysts we have worked with for a number of years [11]. These analysts work for government agencies in which they are responsible for analyzing and making sense of complex situations in the world in order to respond to research requests of policy decision-makers. They communicate their understanding to policy-makers and others, largely in the form of written reports. We have found the analysts to be an excellent source of insights into libraries because their work is so dramatically research-intensive.

The particular study we refer to here consisted of a series of informal workplace interviews with analysts, their managers, information assistants (retrieval experts and database builders), and several technology providers who develop systems and software for the analysts. The analysts we interviewed were members of two organizations, each in a different physical location. Interview topics included questions about the analysts' physical setting (the organization of materials, people, and technology in their offices); the electronic and paper documents they retrieved, scanned, read, and wrote; the systems and software that were available to them and that they used; and some aspects of their social setting—how they worked together and functioned within their organizations.

The study reveals how, in at least this case, researchers engaged in highly information-intensive tasks use a broad range of materials, many of which fall outside the narrow bounds of idealized library collections. It also reveals how, in using this range of materials, they create more fluid, transient, and nondigital materials, constructing and maintaining local collections, which can then be shared with others. Finally, the study highlights the role of informal collaboration and communication in performing this kind of work.

First, let's consider the kinds of documents the analysts use in their day-to-day work. Analysts are voracious, pragmatic readers. They not only have spe-

cific assignments to complete; they are also more generally concerned with keeping up on the literature in their areas of specialization. In addition to using traditional library materials like books and journals in their work, analysts access commercial online information services like Dialog and Nexis. They also receive a continuous flow of email, cable traffic, newswires, and internally published journals made up of their colleagues' reports. Further reference resources may be drawn from other analysts' extensive files; these resources, while they may appropriately be thought of as comparable to library resources, are personal and local. Thus, an analyst's sources may range from the public archival materials we currently associate with libraries, through internal publications and shared collections, to personal files, collected, organized, and annotated for individual use. Moreover, these materials are a heterogeneous mix of digital and nondigital media.

Who maintains these materials? Information providers acquire and maintain materials at one end of the spectrum, and at the other end, individuals take care of their personal collections. But we have also observed an interesting middle ground: reading rooms—physical places set up to hold collections of reference materials that act as a resource for a group of analysts. They are not quite libraries—they are not usually organized according to library conventions, and their maintenance is frequently left to their users—but neither are they quite the same as an individual's files. Like an individual's paper or computer files, reading rooms are a locally constructed, locally available, locally controlled and maintained resource. Like libraries, reading rooms contain archival materials, journals and reports, which are, or course, limited in scope and more highly tailored to the tasks of

Analysts generally do not take notes by writing their observations down on a separate sheet of paper or in a text editor (although under pressure they may combine the activity of reading source materials with the concomitant activity of writing a new report). Instead, they mark on the documents themselves. In the case of books, photographic imagery, and paper archival materials, they do this marking in a nondestructive way—they use Post-its and stick-on signals (little colored dots). In the case of digitally delivered documents, they print copies and use them to contextualize notes. They highlight segments of text (sometimes whole paragraphs) and they scribble marginalia, sometimes noting where what they have seen in the text differs from what they would expect to see (“Not true!”). They also print automatically marked text, documents retrieved from databases that have the keywords that triggered retrieval or filtering explicitly marked (usually underlined). These marking practices increase the value of the documents to the analysts and form the basis for their personal and shared files.

In spite of organizational efforts to make all sources available through digital means and all composition and final production digital, analysts still make extensive use of paper as the principal interpretative medium. Documents retrieved from digital sources, even when they are used as the substance for a digital product, are frequently printed during the course of their use. It seems that from the point of view of this user community, paper is a valuable medium for recording many types of annotations not readily recorded in a digital medium. It is also manipulable in a manner that is not afforded by digital documents: analysts can express nuances of meaning by simply juxtaposing paper documents on their desks. It is common for analysts to spread out

In spite of organizational efforts to make all sources available through digital means and all composition and final production digital, analysts still make extensive use of paper as the principal interpretative medium.

their immediate constituency. It is through reading rooms that we see the most direct connection between products and sources—analysts' reports are made formally available and can become sources for other analysts; materials move in and out of the reading room with some regularity.

What do analysts do with these materials once they have them in hand? We found that annotation is a key means by which analysts record their interpretations of a particular document; in fact, annotation often acts as the mediating process between reading and writing,

their working papers over every available surface and to shuffle them around to reflect various alternative organizations (for example, chronology or subtopic categories). While the digital technologies the analysts use afford desktop-like manipulation of documents, in practice, for many reasons, analysts find nondigital media and real, physical, 3D spaces a more convenient way of working.

Practitioners of information-intensive intellectual work mostly will, if asked, assert that they work alone, discounting both their own reliance on their peers

and the contributions they themselves have made to the work of others. Yet observation of their work reveals many kinds of collaborative practices, most of them informal and most of them institutionally unrewarded. Analysts working in different media (one in the scientific literature, another in satellite imagery) might confer over the phone, looking for corroboration of an interpretation or asking, "What do you make of this?" Occasionally these informal collaborations grow into coauthorships, in which the analysts write a report together, usually by passing drafts back and forth. But more commonly, informal collaborations remain just that, and are an underappreciated part of intellectual work.

Analysts also share interpretative structures and partial interpretations of documents through mutual access to a set of files; one person's files are another analyst's well-tended and well-shepherded reference library. Occasionally, a group of analysts covering a shared topic or area will gather materials into a structured database or shared filing cabinet to be referred to by a whole group. More often, the materials they share are just those they've pulled out of an information resource and filed individually (sometimes in a container as informal as a shoebox). They rely on each other to act as librarians in this situation, evaluating the authority of and providing access to collected materials. Analysts are usually aware of the kinds of materials they can rely on each other to collect.

The work of information analysts thus clearly shows how a broad range of materials (fluid as well as fixed, transient as well as permanent, paper as well as digital) are used in a collaborative fashion. For people performing this sort of information-intensive intellectual work, a digital library, narrowly construed, would be highly inadequate if it were the sole or the primary information source.

We can reach this same conclusion through another route by focusing on the materials analysts create and share locally. Imagine, for the sake of argument, a future world in which all the analysts' primary sources were somehow available in all-digital libraries of fixed, permanent materials. They would still need, as an essential part of their analytic work, to create locally annotated and otherwise modified versions of these primary sources. Some of these more fluid, more transient materials would find their way into locally maintained and shared collections. In other words, the use of digital library materials would create, in effect, a ring of fluid, transient, and nondigital materials immediately outside the bounds of the digital library. Although this material would lie outside the digital library by definition, users would clearly want support for managing its relations with the

library collections from which it was derived. Thus, the problem of integrating this nonlibrary material would remain. Might we not be better off drawing the boundary more broadly from the start, in which case the issue is not digital libraries as separate entities, but libraries that encompass a broad range of materials, digital and nondigital alike?

Some Implications for Research and Development

It might be argued that we already have just such an expanded conception of digital libraries in the Internet and the World-Wide Web. While there is some truth to this, it misses several important points: First, the Net and Web technologies may well be useful as part of a library infrastructure, but an infrastructure in and of itself does not constitute a library. Second, if we include the materials on the Net/Web—the content, that is—these do not constitute a collection, at least not in the sense of a selection of items organized for a particular clientele; it is a bit like calling all the books published by American publishers a collection. Third, the Web and the Net lack the crucial institutional services, such as collection development and cataloging, by which collections are stabilized for ongoing use. Finally, there has thus far been little attention to the integration of digital and nondigital material. None of this argues that the Net/Web infrastructure cannot or should not be used in the construction of digital libraries; but rather that the Net/Web is not in and of itself a digital library in any interesting sense.

What additional work, then, needs to be done to broaden the scope of digital libraries research and development to take into account some of the issues addressed here? We will briefly mention four:

Media integration: The continuing use of paper and other nondigital materials challenges us to effect a rich integration of media. How must our protocols, our naming schemes, our search procedures be broadened if some of the references are not—and will never be—in digital form? How must our document architectures be modified to accommodate hybrid documents, parts of which are in digital form while others are on paper?

Versioning: Current commercially available digital versioning schemes are inadequate for many document-intensive tasks. Computer operating systems, for example, provide little or no help in keeping track of versions of files or documents. To date, far more attention has been paid to version management in CASE tools to support software development than in document management systems. One community where research on versioning is active is the hypertext community [14]. If future collections will include more fluid and transient material, as we expect, then

richer schemes for naming, finding, and manipulating versions will be critical—all the more so if collections include hybrids of digital and paper forms.

Tools for collaboration and communication: If library use is highly collaborative, then the tools we provide must not make unwarranted assumptions about single-user browsing and access. Tools will be needed to support a range of collaborative activities, such as shared annotation and the maintenance of local sub-collections of materials, as well as the communicative behaviors that underlie all work practice. Email, for example, has already proven to be an essential tool for users of the Web. While there has already been much discussion of the role of intelligent agents [4] in supporting information access, little attention has been devoted to the additional tools that *human* agents—e.g., on-line reference librarians—would require to provide service over the Internet.

Service practices: In addition to developing technologies, digital libraries will also need to develop new work practices and procedures to handle the cataloging and maintenance of digital collections. Librarians are already investigating whether existing library practices, largely designed to handle books, can be extended to deal with digital materials, or whether radically new procedures will be needed [8]. This is one area where librarians and computer scientists will need to work closely together in order for technologies and practices to evolve harmoniously.

The Future of Libraries: Conservation and Innovation

So what will digital libraries be like? It is too early in the process of speculation, imagination, research, and development to answer this question. But whatever they turn out to be, they will inevitably share many properties with current libraries and will differ from them in innumerable ways as well. In technological and social change there is always an interplay, a tension, between the forces of conservation and innovation. Cultures and communities do not, and should not, let go lightly of structures and practices in which they have invested heavily. The task in the years ahead will be to decide which existing practices and structures to cast off and which to retain, and which innovations to reject or adopt.

The current direction in digital library research and development is, we believe, quite conservative and in some respects either unnecessarily or inappropriately so. Certain features of existing libraries are being unreflectively conserved, as are certain features of an idealized and unreal past (if it is appropriate to use “conservation” for the holding to a misperception). The current focus on fixed, permanent materials can be traced to a preoccupation with books as the central elements in and the organizing


principle behind earlier libraries. The singular focus on digital documents parallels an imagined homogeneity in today’s libraries. It corresponds more to myth than to reality. And a focus on the individual holds to an idealized image, a fiction, of scholarship as solitary work. But to point out these limiting assumptions is also to point beyond them to a larger field of research and development. Digital technologies offer the possibility of creating and managing more fluid, dynamic documents and sharing these on a scale previously unimaginable.

The academic and public libraries most of us have grown up with are the products of innovation begun approximately 150 years ago.² We would find libraries that existed prior to this time largely unrecognizable. It is certain that the introduction of digital technologies will again transform libraries, possibly beyond recognition, by transforming the mix of materials in their collections and the methods by which these materials are maintained and used. But the better word for these evolving institutions is “libraries,” not “digital libraries,” for ultimately what must be preserved is the *heterogeneity* of materials and practices. As library materials and practices of the past have been diverse—more diverse than idealized accounts allow—so they will no doubt remain in the future.

Library developments ought to be grounded in a solid understanding of past and present practices. Without this, we risk losing still relevant structures and practices while maintaining an allegiance to mythical and irrelevant features of an unrealized past or an idealized present. But to understand and take account of the realities of practice does not mean to be overly constrained by them. On the contrary, seeing things as they are can provide abundant opportunities for grounded innovation. Doing this, however, will require collaboration among technologists, librarians, and library users, as well as other relevant constituencies.

To participate most fully, librarians will need the help of technologists to better understand the possibilities being created by digital technologies, and technologists will need the help of librarians to appreciate the richness of traditional librarianship and to identify the aspects of it that are most relevant to the continuing evolution of libraries. Both of these groups must find ways to attend to the needs of multiple communities of users. There is a lot to talk about... and a lot to learn.

Acknowledgments

Thanks to Bill Anderson, Michael Buckland, Ed Fox, Ted Metcalfe, Fran Miksa, Susan Newman, Ramana Rao, Mark Stefik, and Lucy Suchman for their help in thinking through and articulating the ideas in this article. 

² Miksa, F.L. Personal communication, November 1994.

References

1. Anderson, W.L., and Crocca, W.T. Engineering practice and codevelopment of product prototypes. *Commun. ACM* 36, 4 (June 1993), 49-56.
2. Blomberg, J., Giacomi, J., Mosher, A., and Swenton-Wall, P. Ethnographic field methods and their relation to design. In *Participatory Design: Perspectives on Systems Design*, D. Schuler and A. Namioka, Eds. Hillsdale, N.J., 1993, pp. 123-154.
3. Buckland, M. *Redesigning Library Services, A Manifesto*. American Library Association, Chicago, 1992.
4. *Commun. ACM* 37, 7 (July 1994). Issue on Intelligent Agents.
5. Ehn, P. *Work-Oriented Design of Computer Artifacts*. Erlbaum, Hillsdale, N.J., 1989.
6. Ehrlich, K. and Cash, D. Turning information into knowledge: Information finding as a collaborative activity. In *Proceedings of Digital Libraries '94* (College Station, Tex., June 19-21), J.L. Schnase, J.J. Leggett, R.K. Furuta, and T. Metcalfe, Eds. Hypermedia Research Laboratory, Texas A&M Univ., College Station, Tex., 1994, pp. 119-125.
7. Greenbaum, J. and Kyng, M., Eds. *Design at Work: Cooperative Design of Computer Systems*. Erlbaum, Hillsdale, N.J., 1991.
8. Horowitz, L.R. CETH workshop on documenting electronic texts. Tech. Rep. 2, Center for Electronic Texts in the Humanities, New Brunswick, N.J., 1994.
9. Levy, D.M. Fixed or fluid? Document stability and new media. In *Proceedings of ECHT '94* (Edinburgh, Sept. 18-23). ACM, New York, 1994, pp. 24-31.
10. Levy, D.M. Naming the nameable: Names, versions, and document identity in a networked environment. In *Scholarly Publishing on the Electronic Networks: Filling the Pipeline and Paying the Piper. Proceedings of the Fourth ARL/AAUP Symposium* (Washington, D.C., Nov. 5-7, 1994). ARL, Washington, D.C., 1995, pp. 153-159.
11. Marshall, C.C. A multi-tiered approach to hypertext integration: Negotiating standards for a heterogeneous application environment. In *Proceedings of the Hypertext Standardization Workshop* (Gaithersburg, Md, Jan. 16-18), J. Moline, D. Benigni, and J. Baronas, Eds. National Institute of Standards, Gaithersburg, Md., 1990, pp. 167-177.
12. Marshall, C., Shipman, F., and McCall, R. Putting digital libraries to work: Issues from experience with community memories. In *Proceedings of Digital Libraries '94* (College Station, Tex., June 19-21), J.L. Schnase, J.J. Leggett, R.K. Furuta, and T. Metcalfe, Eds. Hypermedia Research Laboratory, Texas A&M Univ., College Station, Tex., 1994, pp. 126-133.
13. Miksa, F.L. and Doty, P. Intellectual realities and the digital library. In *Proceedings of Digital Libraries '94* (College Station, Tex., June 19-21), J.L. Schnase, J.J. Leggett, R.K. Furuta, and T. Metcalfe, Eds. Hypermedia Research Laboratory, Texas A&M Univ., College Station, Tex., 1994, pp. 1-5.
14. Osterbye, K. Structural and cognitive problems in providing version control for hypertext. In *Proceedings of the ACM Conference on Hypertext* (Milan, Italy, Nov. 30-Dec. 4), D. Lucarella, J. Nanard, M. Nanard, and P. Paolini, Eds. ACM, New York, 1992, pp. 33-42.
15. O'Toole, J.M. On the idea of permanence. *American Archivist* 52, 1 (Winter 1989), 10-25.
16. Piller, C. Dreamnet. *Macworld* (Oct. 1994), 96-105.
17. Yates, J. *Control Through Communication*. The Johns Hopkins University Press, Baltimore, 1989.

About the Authors:

DAVID M. LEVY is a member of the research staff in the Systems and Practices Laboratory of the Xerox Palo Alto Research Center. Current research interests include digital libraries, document reuse, document standards, and the integration of digital and nondigital media. **Author's Present Address:** Xerox PARC, 3333 Coyote Hill Road, Palo Alto, CA 94304; email: dlevy@parc.xerox.com

CATHERINE C. MARSHALL was a member of the research staff at Xerox Palo Alto Research Center for six years. She is presently an associate research scientist at Texas A&M University. Current research interests include digital libraries, hypermedia, and computer support for information-intensive intellectual work. **Author's Present Address:** Dept. of Computer Science, HR Bright Building, Texas A&M University, College Station, TX 77843-3112; email: marshall@cs.tamu.edu

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.