

Collection Maintenance in the Digital Library

Mark S. Ackerman

Roy T. Fielding

Information and Computer Science

University of California, Irvine

Irvine, California, USA, 92717

{ackerman, fielding}@ics.uci.edu

<http://www.ics.uci.edu/CORPS/ackerman.html>

Once this [restoration of learning] was the task of princes, and it was the greatest glory of Ptolemy. But his library was contained between the narrow walls of its own house, and Aldus is building up a library which has no other limits than the world itself. -- Erasmus, Adages, 1508

I love the Web. You can find all kinds of interesting information. At least theoretically you can. Of course, things come and go, and sometimes you can't find what you're looking for. -- Professor of Computer Science.

ABSTRACT

Maintenance will be critical to digital libraries, especially those that promote broad access to diverse, informal materials. If ignored, maintenance issues within the digital library -- especially those relating to its materials -- will threaten its usefulness and even its long-term viability.

We perceive the maintenance problem to be both technical and institutional, and this paper considers the maintenance of the digital library as both institution and technology. The paper examines collection maintenance from several vantage points, including software architecture and the type of collection, arguing that digital libraries that contain informal and dynamic material will have substantially greater maintenance problems. The paper ends with an examination of potential technical solutions.

KEYWORDS: digital library, collection, maintenance, World Wide Web, organizational memory, usability

INTRODUCTION

As with any new technology-based idea, there has been considerable controversy over the definition and possibilities of the term "digital library." To the computer science community, the new technical possibilities beckon enticingly. However, as traditionalists in the library

community might point out, important issues are being ignored.

This paper promotes a view of collections and the long-term consequences of their operation, based in a consideration of digital libraries as social institutions. This runs contrary to the substantial body of digital library research that focuses on creating the initial collections and providing access mechanisms. We believe that the problematic must be recast to include long-term issues. By centralizing those issues surrounding the maintenance of institutions and their artifacts -- especially the library collection -- important considerations for the long-term success of digital libraries emerge.

To distinguish our concerns from traditional collection management, we call these concerns *collection maintenance*. We use "maintenance" to deliberately invoke "software maintenance" and its often ignored importance for software systems.¹ As will be discussed below, collection maintenance is likely to be a significant problem in the digital library -- more so than in traditional libraries or current organizational memory repositories.

The conception of collection advocated here includes access to informal and dynamic information, albeit with a strong caveat. In our view, informality and dynamism materially increase accessibility and content issues over the long run. Collection maintenance for this expanded notion of collection will be difficult, but absolutely critical since:

The necessities and problems of collection maintenance are intertwined with institutional viability over the long-run.

The traditional, or paper-based, library has established methods of maintaining access over the long-run. However, we do not yet have maintenance methods for the digital library, especially within the conception that includes informal and dynamic materials.

This paper begins by discussing the differing notions of the digital library, anchoring the issues in an analysis of institutional needs and practices. We then examine the various types of collections, including those that include dynamic and informal materials. This consideration of collection types and their control lends itself to analyzing the institutional arrangements and resulting maintenance issues for digital libraries.

This analysis proffers two conclusions. First, maintaining collections that are extensions of traditional collections (with delineated boundaries) not surprisingly requires only extensions of traditional methods. Existing institutional arrangements and resources can be modified to handle these requirements. Second, maintaining collections that include dynamic and informal information will be possible only with new technical solutions. We therefore end the paper with some prototypical software tools, using the World Wide Web as an example. One, called MOMspider, checks links within a defined area of the Web. The other, called Web:Lookout, checks for new links and information in previously viewed Web pages.

THE “LIBRARY” AND ITS BOUNDARIES

In this paper, we wish to discuss the long-term issues in maintaining a digital library as both technology and as institution. However, before doing that, we need to decide which conception of the term “digital library” to consider. This issue, one of social construction, has implications for the types of maintenance technologies that will be necessary and useful.²

The Broadly -Construed Library

As Levy and Marshall [17] note, there is the “digital library” narrowly-construed and more broadly-construed. In their paper, they argue for the broadly-construed digital library. Their vision argues for rich access to a variety of materials, all of which can be used in an interpretive and collaborative manner:

We argue ... that the design of digital libraries must take into account a broader range of materials, technologies, and practices -- transient as well as permanent documents, fluid as well as fixed materials, paper as well as digital technologies, and collaborative as well as individual practices. (p. 163)

In a related paper, Marshall, Shipman, and McCall [20] argue for fluid access and integration among many disparate sources. They further offer an idealistic conception of what technology could allow.

On the other hand, the narrowly-construed library is a single collection, defined sharply from materials not in the collection. As opposed to the broadly-construed library, the narrowly-construed library lacks integration with other

information sources and community practices. It presumably does include new possibilities of access, but perhaps only those pre-defined by the library designers.

Miksa and Doty [21] argue for the narrowly-construed digital library. This view idealizes the role of the collection and the additional indexing above the raw material:

...the idea of the library includes the construction of a set of arrangements that overcomes the disparateness of the individual sources by relating them to one another in terms of a single, operational, intellectually structured whole. (p. 4)

Both of these papers attempt to confront the boundaries of the collection, and those boundaries are in different places for the two papers. Miksa and Doty emphasize the collection and intellectual access to it. Levy and Marshall’s emphasis is on access and use of the collection by a community, and since their emphasis is on practice, the collection is used in conjunction with other information sources.

These Digital Library ‘94 papers offer visions of what is important in a digital library, and as such, offer considerable insight into two sets of important core beliefs. These two sets of beliefs are not completely separate; we tease them apart here because of their subtle implications for maintenance and the kinds of institutional arrangements required for that maintenance in the long-run. One view, more intrinsic to the library community, has found it important to consider collections, the intellectual access to them, and classification (e.g., [12, 16, 24]). Another view, held more strongly by the hypertext community, has an interest in direct-manipulation access and integration issues among heterogeneous materials (e.g., [22]). Although not completely separable, these views point to different emphases on what is important to maintain. One emphasizes selection from a bibliographic universe; the other, interaction among materials.

Multiple Libraries and Multiple Collections

The apparently dichotomous nature of the digital library architecture has broad consequences for its maintenance. However, most technical architectures will be some combination; it will be argued below that this is likely for institutional reasons. In this section, we show that some technical architectures muddy the question of broad versus narrow. It is entirely possible to have both types of libraries simultaneously.

Figure 1 shows some possibilities; we mean for these to be illustrative rather than definitive examples of technical possibilities. Figure 1(a) is the library narrowly-construed with an application that serves as a personal client. This client would function in much the same way as a Z39.50 (information retrieval) client, providing a satisfactory user

interface and user functionalities. Figure 1(b) shows that there could be organizational memories or other information repositories serving as group collections [2, 3]. Finally, Figure 1(c) shows the range of possible information sources. Here, the personal client interacts with organizational memories, group or intermediate information repositories, multiple digital libraries, and informal sources of information (such as the Web or Usenet).

These architectures suggest that by extending other metaphors, we can find other results. For example, if we extend “personal library”, we find a need for integrating information sources, not for all of society but for an individual. We can also extend “corporate library” or “organizational memory” to consider information repositories for groups or organizations, intermediate groupings between society and the individual.

Such architectures provide not just intriguing technical possibilities for computer scientists. They also provide ways for the narrowly-construed digital library to be incorporated within personal, group, and organizational information repositories. These intermediate architectures suggest that the idea of collection will be much more porous in the digital library: The intermediaries that are only implicit in the traditional library will be much more tightly tied into the digital library. However, the linking in and arrangement for multiple collections, intermediate layers, and informal materials makes some current institutional practices difficult, especially those related to the maintenance of the library over time.

Institutions and Their Practices Over Time

The library community has been very successful in

maintaining traditional libraries and has developed many practices and ideals to do so. Among them are the ideas of unified collections and access methods (including organized indices and catalogs). In fact, one could argue that many core library activities, such as circulation, technical services, and even shelving, were created by the library to maintain access to its collection over time.

In this institutional analysis, any set of practitioners develops normative structures over time to promote the maintenance and promulgation of their community (as in [11]). (This is the larger context of technical maintenance.) It is not just that the institutions themselves provide for their continued existence; the community members promote such activities for the institutions [5].

It is essential to note that this set of practices and ideals maintain the traditional library as institution and as community of practice together over time. This equilibrium is in the process of being disturbed or destroyed with the intrusion of new technical possibilities.

We therefore ask: *What will be required to maintain continued access and availability of the digital library over time?*

Only some of the practices of the traditional library will carry into the digital world, perhaps only for a narrow conception of the digital library. Furthermore, the hypertext and computer science communities have no such traditions and practices for the broadly-construed library.

If we wish either vision to be successful, the above question must be answered. Some of the answer will be institutional and some of it will be technical. The mix that

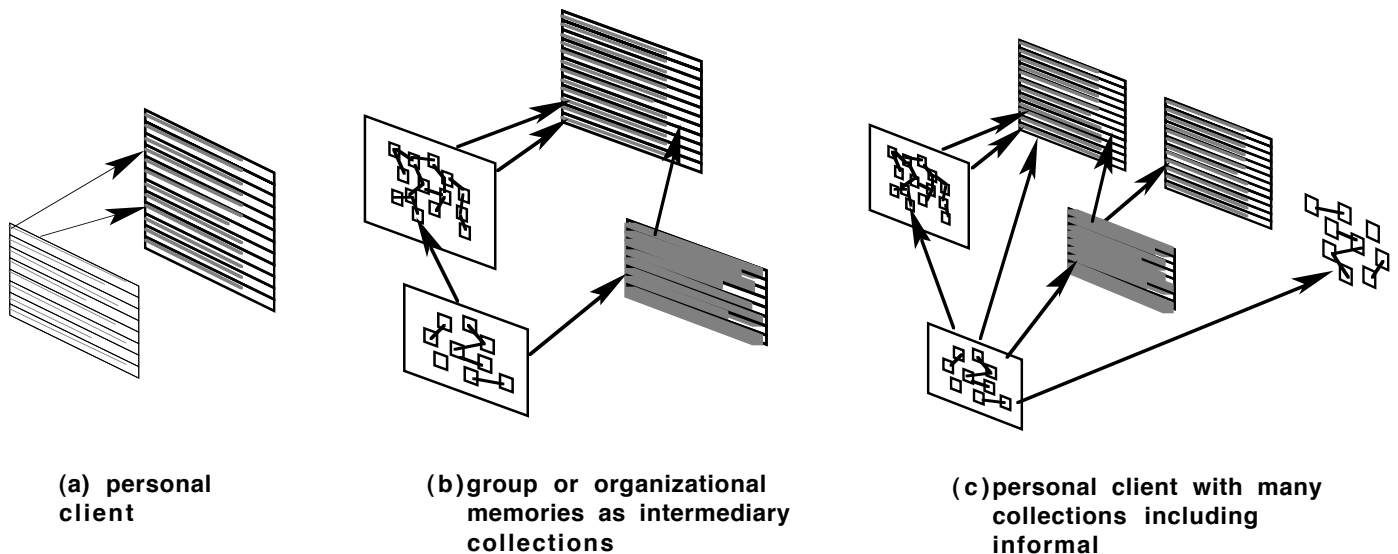


Figure 1: Potential configurations of personal clients, intermediate collections, formal collections, and informal sources

	digital library type			
	traditional library	organizational information repository (organizational memory)	World Wide Web	Usenet or other computer-mediated communication system
type of collection	monographs, serials, special collections	documents and files	text and multimedia nodes	messages
authoring agencies	authors and their publishers	organizational members or sanctioned individuals	anyone	anyone
collection control	selection by organizational members (e.g., bibliographers)	selection by organizational members. May have requirements for official approval	none (or individual)	none (or moderator)

Table 1: Collection control in digital libraries (by type)

is possible, however, will be dependent on the types of desired collections as well as the potential control mechanisms for those collections.

COLLECTIONS IN THE DIGITAL LIBRARY

In the management of most resources, managerial control can play many roles. In the digital library, different control mechanisms over the collection are possible, and these control mechanisms will influence the possibilities and necessities for maintenance. These control mechanisms are influenced by collection types, as an examination of different collection types demonstrates.

Collection Types and Control

Libraries have always managed their collections, selecting and removing items from their shelves. This has been viewed as a critical function of library management [12, 14, 24]. According to the collection management literature, the practices of collection management are dependent on the type of library collection. For example, Drew and Dewe [9] distinguish ephemera in special collections from other items, and Kovacs [16] distinguishes among types of libraries. We will argue here that new types of collections in the digital library will lend themselves to new types of maintenance issues.

Table 1 delineates four types of digital collections. This is not the only method of distinguishing among “digital libraries”; for example, we could have included access methods or network topologies. Additionally, an actual digital library could have elements of any (or all) of these

types. Nonetheless, distinguishing among types on the basis of control over the collection is important to the continuing operation of a collection.

In a traditional, or paper-based, library, there is considerable control over the collection. Versions of publications are stable (i.e., the contents of non-ephemera do not change from copy to copy). More importantly, because the collection is physically contained, there is considerable control over the collection. Library staff *can* decide what is and what is not in the collection. Maintenance of the collection is within the purview of the institutional members.

At one opposite is the World Wide Web. The Web nodes often change, in content, location, and even existence. On the other hand, the content in a location does not shift rapidly; it tends to remain relatively stable. There is no control by any given individual over the entire Web. An individual has control only over his nodes and the selection of pointers to others’ nodes (URLs) that provide the capability for extended collections.

At another opposite is Usenet or similar computer-mediated communication (CMC) systems. The locations (i.e., channel or topic) do not change, but the contents of any given location (e.g., comp.sys.laptops) change constantly. The control over the collection for this type is also very low for any given individual.

In the middle of both dimensions is the organizational information repository (organizational memory). Since the information resides within an organization and since the organization generally provides a sanctioned organizational member to manage the system, there is considerable control over its collection. Nonetheless, organizational memories tend to be more dynamic than traditional library collections; for example, they may include informal information, time-sensitive information, or bulletin board mechanisms.

Collection Maintenance

The variation in collection control determines the type of institutional and technical maintenance possible. Collections that are closer to those in traditional libraries can use more traditional control and maintenance mechanisms. Digital libraries that incorporate more individualistic, dynamic, and informal information may need to find new maintenance mechanisms.

Digital libraries, if self-contained in a manner similar to traditional libraries (i.e., the narrowly-construed library) will be able to draw upon traditional methods and practices. Traditional libraries, as noted above, have developed methods for maintaining their core set of institutional ideals, their community of practice, and their collections of materials. Although many maintenance practices will need to be adapted -- such as preservation and circulation control as extreme examples -- the existing

practices can serve as the bases for new norms and practices. We have argued above that such norms and practices are required for the long-term viability of the digital library as institution. Therefore, if the adaptation of current practices is successful, then the narrowly-construed digital library will stand a much greater chance of success in the long run.

These traditional practices, however, have had their limitations. They were based on a constrained collection; i.e., a selection from the bibliographic universe. A traditional library could never cope with much ephemera; it would require too many resources. To adequately index, catalog, and otherwise conform to the institutional ideals for the broad range of ephemera would be overwhelming. In the past, there were strong institutional reasons to constrict the bibliographic universe.

Going digital, however, changes the cost structure, and collection and maintenance costs need to be revisited. It is not necessary, for example, in a digital universe to hand-index all materials; there are other methods of access.

There is considerable evidence that this change in the cost structure is affecting organizational memories (see [15] for a popular summary). As Figure 2 shows, however, the dynamism and volatility of organizational memories tend to be close to that of traditional libraries. That is, organizational members can reliably return to the same location in the organization's information repositories to find the same materials. However, the widespread interest

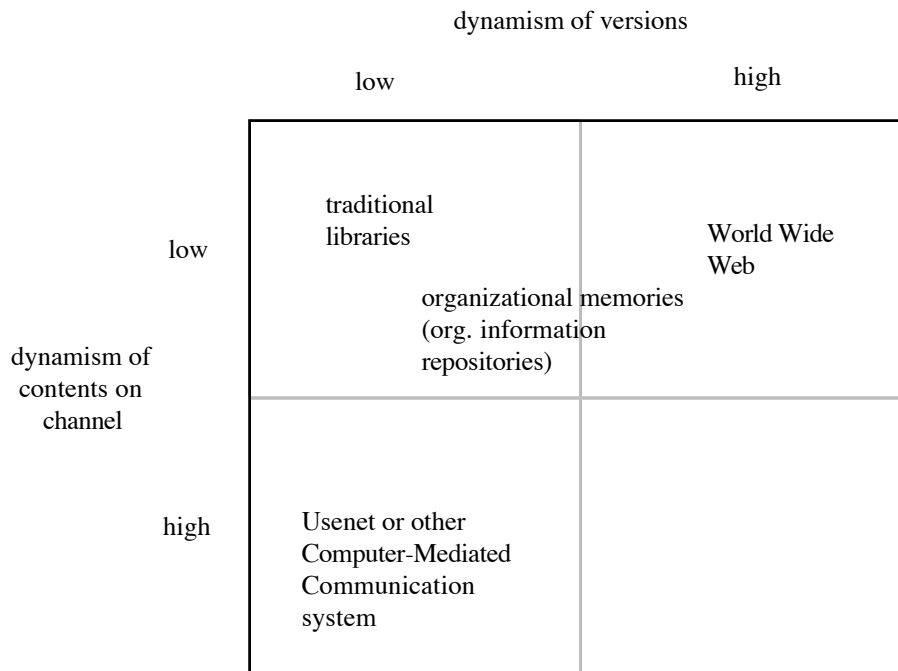


Figure 2: Dynamism (volatility) in types of digital library collections

in Lotus Notes, which combines electronic messaging and document storage, argues that combining formal organizational materials with informal is valuable (e.g., [23]). It will be interesting to see how existing maintenance procedures within organizations will adapt [7, 8, 25, 27].

Nonetheless, because institutional adaptation is much more manageable and amicable than extreme change, we expect that there will be pressure to have digital libraries as discrete and contained collections. The resource limitations of any social institution argue for defining a discrete collection as well [19]. Of course, it will be possible to incorporate these narrowly-construed libraries in architectures that promote broad access.

In addition, the broadly-construed digital library lacks practices that will maintain itself over the long run. The same incentive that promotes use of the Web or of Usenet -- namely that there is no intervening institution between the author and his potential audience -- also brings its own disincentives to systematic use. As discussed, both the Web and Usenet have high levels of volatility. The location of Web nodes can change without notice. Usenet groups (e.g., comp.multi-media) do not change locations, but the message traffic changes constantly. And from Table 1, there is little, if any, control over either of these informal collections.

The individualistic orientation of both the Web and Usenet argues against easy control over the collection. Without institutional control, each individual using the collection must deal with the maintenance issues himself. If the informal materials are dynamic or volatile, then there is a constant maintenance problem.

Requiring every user to provide for collection maintenance raises the costs considerably. Users must determine whether materials will continue to be available, accessible, and even understandable.

Collection maintenance, then, will be a critical limitation of the broadly-construed digital library. If the broadly-construed library, with its access to multiple sources, is to be viable, users cannot have a high maintenance cost over the long run. It will be necessary to reduce the cost of long-term use, and without mediating institutions, this reduction will need technical solutions.³ The following section examines two such maintenance tools.

MAINTENANCE AGENTS AND FILTERS

The World Wide Web

Collection maintenance tools will be necessary for the digital library. As examples of these tools, we briefly consider here several potential mechanisms for reducing the user's maintenance cost in his continued use of the World Wide Web.

These examples are limited. Although many of the same issues exist for other informal materials, the mechanisms for dealing with them may not. For example, trying to capture traffic on Usenet groups might require methods of summarizing series of messages. Still, the Web is a valuable testbed for dealing with maintenance issues. Additionally, since many computer science pre-press materials are available on the Web, the Web serves as a salient and immediately useful example.

The Web shares the same maintenance problems as most hypertext systems. For example, in a hypertext collection such as the Web, one needs to:

- Check that there are nodes at the other end of a link. In a distributed system such as the Web, the site may no longer exist, or the node may have moved. If a node or link has changed or has vanished, there will not be a local notification.
- Determine that all nodes have links into them (i.e., there are no orphans other than entry points into a local area).
- Check that the information in nodes is not obsolete (e.g., by examining expiration dates for objects).

Since standard Web clients do not provide type information for nodes or links, consistency checking for types is not feasible. Similarly, since versioning is not generally supported, no checking for the correct version can be made.

Below are several semi-autonomous agents that attempt to ameliorate Web maintenance problems. Neither agents nor the list of problems is without omission, and again, these mechanisms are meant to be suggestive of work that will be required for the broadly-construed library.

MOMspider

MOMspider (Multi-Owner Maintenance spider) can automatically traverse a focused portion of the Web. It is most valuable in checking a portion of the Web as a local collection.

MOMspider identifies items that require the owner's attention such as broken links, moved documents, modified documents, and owned objects with expiration dates near to the current date. It can also provide an index document that can itself be traversed as a hypertext document. MOMspider [10] is conceptually based on an earlier tool for the Answer Garden Substrate [1].

This automatic maintenance obviates the need for continuous, manual traversal of a Web collection. The Web may be traversed by owner, site, or document tree, and it is possible to mark areas for the spider to avoid. However, for efficiency reasons, MOMspider performs the traversal from the perspective of a site administrator. Furthermore, MOMspider cannot analyze changed materials to see if the changes are significant to the user. This centralized focus makes it harder for individuals to use MOMspider for personal collections and limits their ability to customize the maintenance process (and focus it) to their needs.

Web:Lookout

Web:Lookout [4] is a semi-autonomous agent that can periodically examine a compiled list of interesting Web nodes or areas. Web:Lookout can determine when links or nodes have been added, removed, or changed.

Like MOMspider, the purpose of Web:Lookout is to avoid manual maintenance. Its major use is to automate the examination of “hot lists” or “favorites pages”. It is fairly common on the Web for individuals to have nodes with links to other individuals’ interesting nodes. These interesting nodes may be associated with a given topic (for example, computer-mediated communication or Web agents), making them extremely valuable to researchers. Additionally, individuals may have pages with pointers to work in progress, pre-press or published papers, informal descriptions of projects, data, and the like. Again, these pointers can be extremely valuable to researchers.

Unfortunately, without examination, there is no way to know when an author has added in a new link. Similarly, one cannot know when an author has added nodes to his Web area, perhaps providing new descriptions or progress on a project. Instead of having to manually traverse these lists to look for new items, Web:Lookout allows users to automate this process.

Unlike MOMspider or other robots that periodically examine specific locations for any change, Web:Lookout notifies the user only upon interesting changes. Its heuristic examines nodes for the specifics of content and link changes. For example, if the user were looking for new publications by a colleague, he could request that Web:Lookout note whether a local link has been added to the colleague’s publication page. In this scenario, the user would probably not be interested that links have been deleted or text has been added; Web:Lookout would not notify the user of these things if so instructed. Web:Lookout can also determine whether the content has significantly changed (measured by a similarity metric which the user can set), whether a localized Web area has changed shape, and whether other individuals have

stopped finding a particular link interesting. Web:Lookout can also place its results in a Web page for others to use.

The ability to obtain automatic notifications of changes in other parts of the Web complements the social practice of providing interesting links. As such it serves as a useful social filter, similar in effect but not in mechanism to Maltz and Ehrlich’s pointer filter [18]. Web:Lookout is, in their parlance, an active filter; however, Maltz and Ehrlich require each author to notify potential users. Web:Lookout removes this notification bottleneck.

Limitations of Maintenance Agents

The major issue with informal material is its impermanence. While informal material has the advantage of being more timely, it also vanishes. To return to the material later, it must be captured in some permanent store. Neither agent presented above deals with this problem.

Permanently storing all informal materials within a central location would be prohibitively expensive. This essentially duplicates the traditional problem of selecting ephemera from the bibliographic universe. To select more than a handful would threaten the rest of the collection because of the cost.

However, the costs to any given individual user may not be large. It is possible that any given user would not require substantial storage, and a simple mirroring process might be adequate. Empirical investigation of user habits, however, is required to determine whether this is feasible.

A further limitation is the inadequacy of any computational method of examining content. Robots cannot fully examine the intellectual content of informal information, and any attempt to even partially examine contents may require extensive network resources. We have tried to consider this in the design of Web:Lookout, which uses no more resources than manual traversal. However, careful consideration will need to be paid to the tradeoff between efficiency and usefulness.

CONCLUSIONS

This paper examined some of the mechanisms required to maintain the digital library over the long-run. *Collection maintenance* is likely to be a significant issue in the digital library -- more so than in traditional libraries or current organizational memory repositories.

Following Levy and Marshall, we examined in this paper both the broadly-construed and narrowly-construed digital libraries. The narrowly-construed library is an analog of the traditional library where the collection has known boundaries. Because of the possibility of control over the collection in the narrowly-construed library, many of the

institutional mechanisms for maintaining collections can be assimilated from the traditional library.

In the broadly-construed digital library, users are able to access diverse materials. The inclusion of dynamic and informal materials in the collection, however, leads to serious control and long-run maintenance issues. Because of the lack of control over the collection (or collections), technical mechanisms will be needed for collection maintenance.

While this paper ended with several technical possibilities for collection maintenance, we wish to emphasize that we perceive this problem to be both technical and institutional. Many proposals for digital libraries remove social exchange and interaction, focusing narrowly on the technical mechanisms of information; however, a strictly technical emphasis will not lead to an adequate understanding of the long-run issues in digital library use. Accordingly we have tried to emphasize both the technical and the social perspectives. The digital library is more than a set of technologies; it is also a social institution with long-term needs and maintenance requirements.

ACKNOWLEDGEMENTS

These projects have been partially funded by grants from the UCI Committee on Research and NASA (NRA-93-OSSA-09), and the Advanced Research Projects Agency (MDA972-91-J-1010). The content of the information does not necessarily reflect the position or the policy of the U.S. Government and no official endorsement should be inferred.

This paper benefited from conversations with Steve Franklin, Neal Kaske, Rob Kling, Jonathan Grudin, and Lisa Covi.

REFERENCES

1. Ackerman, M. S. Answer Garden: A Tool for Growing Organizational Memory. Massachusetts Institute of Technology, Ph.D. Thesis, 1993.
2. Ackerman, M. S. Augmenting the Organizational Memory: A Field Study of Answer Garden. *Proceedings of Conference on Computer-Supported Cooperative Work (CSCW'94)*, 1994, 243-252.
3. Ackerman, M. S. Definitional and Contextual Issues in Organizational and Group Memories. *Proceedings of Twenty-seventh IEEE Hawaii International Conference of System Sciences (HICSS 94)*, 1994, 191-200.
4. Ackerman, M. S. Web:Lookout: A Social Filter for the Web. In preparation, 1995.
5. Berger, P. L. and T. Luckmann. *The Social Construction of Reality: A Treatise in the Sociology of Knowledge*. Anchor, New York, 1966.
6. Callon, M. Society in the Making: The Study of Technology as a Tool for Sociological Analysis. *The Social Construction of Technological Systems*. Bijker, Hughes and Pinch ed. MIT Press, Cambridge, MA, 1987.
7. Chalmers, A. M. Case Study: Records Management in a Pharmaceutical Environment. *Records Management Journal*, 1(1), 1989, 19-28.
8. Dollar, C. M. and T. E. Weir. A National Archives Strategy for the Development and Implementation of Standards for the Creation, Transfer, Access, and Long-Term Storage of Electronic Records of the Federal Government. National Archives National Archives Technical Information Paper No. 8, 1990.
9. Drew, P. R. and M. D. Dewe. Special collection management. *Library Management*, 13(6), 1992, 8-14.
10. Fielding, R. T. Maintaining distributed hypertext infrastructures: Welcome to MOMspider's Web. *Computer Networks and ISDN Systems*, 27(2), 1994, 193-204.
11. Foucault, M. *Madness & Civilization*. Vintage, New York, 1973.
12. Gardner, R. K. *Library Collections: Their Origin, Selection, and Development*. McGraw-Hill, New York, 1981.
13. Ghezzi, C., M. Jazayeri and D. Mandrioli. *Fundamentals of Software Engineering*. Prentice Hall, Englewood Cliffs, NJ, 1991.

14. Hall, B. *Collection Assessment Manual for College and University Libraries*. Oryx, Phoenix, 1985.
15. Kirkpatrick, D. Why Microsoft Can't Stop Lotus Notes. *Fortune*, 130(12), 1994, 141-148.
16. Kovacs, B. *The Decision-Making Process for Library Collections: Case Studies in Four Types of Libraries*. Greenwood Press, New York, 1990.
17. Levy, D. M. and C. C. Marshall. What Color was George Washington's White Horse? A Look at the Assumptions Underlying Digital Libraries. *Proceedings of Digital Libraries '94*, 1994, 163-169.
18. Maltz, D. and K. Ehrlich. Pointing the Way: Active Collaborative Filtering. *Proceedings of ACM Conference on Human Factors in Computing Systems (CHI'95)*, 1995, forthcoming.
19. March, J. G. and H. A. Simon. *Organizations*. Wiley, New York, 1958.
20. Marshall, C. C., F. M. Shipman and R. J. McCall. Putting Digital Libraries to Work: Issues from Experience with Community Memories. *Proceedings of Digital Libraries '94*, 1994, 126-133.
21. Miksa, F. and P. Doty. Intellectual Realities and the Digital Library. *Proceedings of Digital Libraries '94*, 1994, 1-5.
22. Nelson, T. *Literary Machines*. Private printing, Swarthmore, PA, 1981.
23. Orlikowski, W. J. Learning from Notes: Organizational Issues in Groupware Implementation. *Proceedings of Computer Supported Cooperative Work (CSCW) '92*, 1992, 362-369.
24. Osburn, C. B. and R. Atkinson. *Collection Management: A New Treatise*. 1991.
25. Robek, M. F., G. F. Brown and W. O. Maedke. *Information and Record Management*. Glencoe Publishing, Encino, CA, 1987.
26. Schach, S. R. *Software Engineering*. Irwin, Homewood, IL, 1990.
27. Walch, V. I. The Role of Standards in the Archival Management of Electronic Records. *American Archivist*, Winter 1990, 1990, 30-43.

¹In software engineering, maintenance is generally defined as any development, enhancement, correction, or adaptation after the initial release of a software product [13, 26]. It includes anything required to keep the system running or to continue access by users over time. According to Schach [26], maintenance consumes about 65% of a software product's cost.

²The development and continued existence of institutions and their technologies interact. This is something of a paradox in that we need to know the future in order to design the systems for it. It is worth considering Callon's tale of technological innovation [6]. In his analysis, designers of technical systems are forced to predict the future, since their designs succeed or fail based on their understanding of the potential future. While developing technical solutions, we are forced to confront the sociological issues. We would extend this to argue that the futures come to exist in interaction with the technologies. In short, the sociological and the technical possibilities and constraints are intertwined.

³It may be possible that new forms of control and maintenance will arise, but there is no evidence of this as of the present. Developing institutional practices without an institution would be difficult.