NoteCards in the Age of the Web: Practice Meets Perfect

Catherine C. Marshall
Microsoft Corporation
1 Microsoft Way
Redmond, WA 98052
cathymar@microsoft.com

Abstract

Frank Halasz's "Reflections on NoteCards: Seven Issues for the Next Generation of Hypermedia Systems" was a remarkably prescient analysis that continues to influence the international hypertext research community. Meanwhile, the Web has offered a basic reality check on the seven issues and has given us, as a community, an opportunity to learn from many and diverse hypertext practitioners. In essence, the Web has brought hypertext out of the realm of research and into the realm of the everyday, the ordinary, the practical. In particular, I would like to introduce three major themes that come from observations of the Web in use: (1) The growing heterogeneity of hypermedia genres, uses, and users; (2) the need to acknowledge the distinct role of hypermedia readers and, more specifically, provide hypermedia readers with tools for personal annotation, re-retrieval, gathering, contextual access from mobile devices, and collaborative reading; and (3) the recurring tension between formal and informal hypertext structures and representations.

Keywords

hypertext/hypermedia, use, user interface design, World-Wide Web, reading, retrieval, formality

Introduction

"Reflections on NoteCards: Seven Issues for the Next Generation of Hypermedia Systems" (Halasz, 1988) prescribed a research agenda that profoundly influenced the direction and thought of the international hypermedia community for more than a decade. Meanwhile, the Web, a technology that was still on the drawing board when Halasz introduced his seven issues, gathered a force and momentum that surprised even the most fervent hypertext aficionados. By the mid 1990s, hypertext was no longer confined to the lab or academia; instead, hypertext practitioners—people who read, designed, and otherwise used the technology in their day-to-day work—had become commonplace, an acknowledged part of what was hailed as the new information economy. The term *hypertext* (or more specifically, *hypermedia*[1]) no longer needed Halasz's

ACM Journal of Computer Documentation, 2001,25:96-103

carefully crafted opening explanation, "Hypermedia is a style of building systems for information *representation* and *management* around a network of multi-media nodes connected together by typed links" (Halasz, 1988, p. 836). The fellow in the next airplane seat knew that hypertext meant clicking on a link to get to a new place (and that reading hypertext was, in a word, surfing).

Thus the world had simultaneously expanded and constricted for those of us who had participated in the development of what Halasz referred to as the second generation of hypermedia systems: Intermedia (Yankelovich et al., 1985), NoteCards (Halasz et al., 1987), KMS (Akscyn et al., 1988), Hypercard (Goodman, 1987), Storyspace (Bolter, 1991), and many other sophisticated efforts to make useful and usable computer tools of this sort, tools that by and large supported intellectual work. The world had expanded in that our experiences supporting small homogeneous user communities and their tightly woven hypertext networks could be brought to bear on solving broader problems and seeding user communities that were much larger in scope and numbers. It had constricted just by virtue of the fact that, once everyone knew what hypertext was, it became much harder to extend the view of hypertext to include, for example, links that were non-network structures (Halasz, 1991) or to imagine hypertexts as we did in which links were implied by action, by spatial and temporal characteristics, and could be more fluid and ambiguous (Marshall and Shipman, 1995).

Given Halasz's prescient discussion of the technological issues that came to light during the course of the NoteCards effort—many of which still form the basis of research projects and chart a space of important unsolved problems—it would be easy to interpret the Web as a limiting force, one that narrowed the field to a simple page-based model of hypertext with embedded links. Unlike the original vision of hypertext, at the Web's outset link-making and annotation were the province of the pages' authors. Nor could nodes be organized and structured in multiple ways; instead, embedded links represented in HTML described a single way of structuring a hypertext. Furthermore, many of the efforts to introduce database-like transactions into hypertext were thwarted by the Web's intentional statelessness; in fact, much early Web research focused on ways of reintroducing state to the Web.[2]

However, the Web offered something else that transcended any limitations it introduced: the ability to learn from many and diverse practitioners. It gave us, as a research community, the most basic reality check on the seven issues, and gave all of us the ability to reconsider our earlier reflections. What I attempt to do here is to bounce critical parts of the seven issues off the characteristics of the Web and our observations of the Web in use. What aspects of the seven issues still dog us today, and what new questions have arisen through this unparalleled opportunity to see hypertext in widespread use? It is only fair to take under consideration Halasz's revisions to his original assessment of what the important issues are (Halasz, 1991); these revisions prefigure some of the observations in this paper. In particular, I would like to introduce three major themes: (1) The growing heterogeneity of hypermedia genres, uses, and users; (2) the need to acknowledge the distinct role of hypermedia readers, and more specifically, tools for using hypermedia beyond search; and (3) the recurring tension between formal and informal hypertext structures and representations.

Heterogeneity of Hypermedia Genres, Uses, and Users

One of the most noticeable assumptions underlying Halasz's seven issues is an unproblematic notion of the genres, uses, and users of hypermedia. The Web offers us unprecedented heterogeneity along all three of these dimensions. While Halasz recognized by 1991 that there was considerable diversity in our field,[3] he did not anticipate the webmasters, amateur publishers, advertisers, hucksters, and people of many walks of life who have taken up hypermedia as their chosen vehicle for expression and influence, but have no strong philosophical or design stance on many of the research issues that we take to heart (whether or not scrolling destroys the navigational metaphor of hypermedia, for example).

The second generation hypertext systems were designed with intellectual work in mind. We envisioned users to be people like ourselves—students and teachers, lawyers and policy makers, researchers and (later) knowledge workers, museum-goers and information-seekers, writers and critics. We certainly did not envision shoppers, pornographers, advertisers, entertainment-seekers, and media conglomerates as NoteCards users. Because intellectual work involves exploration and follows a model of readers as writers who respond to texts with annotations, links, and new hypertexts of their own, we inhabited an ethereal plane apart from

many of the day-to-day activities of people. Hypermedia in the abstract was a populist, participatory medium, quite apart from our experiences in the paper document world. Even in Halasz's revised account of the seven issues in his memorable keynote address at Hypertext '91, he still used intellectual work (authoring, knowledge representation, interactive fiction, computer-assisted software engineering, computer-aided instruction, and so on) as the foil for his reflections on his past analysis.

Instead, today a diverse spectrum of Web users have incorporated hypermedia into their work, their entertainment, and indeed their lives, and wonder more appropriately how hypermedia fits into the metaphors offered by their own activities, rather than how their activities fit into the model offered by hypermedia. For example, a librarian might wonder how a Web page should be cataloged given the Anglo-American Cataloging Rules, and an on-line shopper might wonder whether an item has really been removed from her electronic shopping cart if she's hit the Back button on her browser. A home user might wonder how he can find a porn site that caters to his most esoteric fetishes, and whether or not his kids can find the same one by looking through the browsing history. In a word, we are now confronted by a broader scope of intentions and by metaphors derived from a much broader range of associated activities.

Let's look at how this increased scope changes our take on the issues. By 1991, there was significant work on Halasz's first issue, search and query. It is clear that search is indeed a fundamental element of effective navigation. In fact, when we see people using the Web, they often use a search engine (as opposed to a portal or known URL) as a way of finding a starting point, switching to link-based navigation only when they're "close" to what they want.[4] And, as Halasz anticipated, hypermedia structure comes into play in some of the more successful search engines, albeit not in the specific way we all expected, through user description of the desired structure, but instead through how many other pages link to this one. However, significant social factors have clouded the clear picture of search and query offered by the hypertext community.

What do I mean by social factors? As early as 1995, it became clear that the one key battle was for readers' attention. That year, newspapers like the *San Jose Mercury News* carried articles about Web content developers spoofing the then-popular search engines by using false metadata to describe what a given page was about or by using words hidden in page backgrounds to introduce new terms or exaggerating the importance of others:

> If you want to attract people to your Web site, make sure you've spiked your location with the hottest search keywords on the Net. So you may be selling widgets, but if you sprinkle a heavy dose of words such as "sex," "nude" and "naked" into your site—in ways that aren't necessarily visible to users—you'll have the best-read widget site. Guaranteed. (*San Jose Mercury News*, 1997)

Hence, Web pages seeking reader attention might then use an entire dictionary's word list, so that every term brings the searcher to the site, or they would appeal to particular of the reader's prurient interests. Or an automobile dealership might have the word *car* duplicated 500 times in the metadata tag to affect the page's ranking, in much the way the Yellow Pages contained listings for AAA Bail Bonds or A-Aardvark Pest Control. From then on, a unique dialectic arose between search engine developers (who sought to give their users what they were looking for so they would continue to use the search engine) and content developers (who simply sought readers' attention).

Furthermore, intention comes into play even without deliberate misrepresentation. For example, Google uses the Web's link structures to bring factors like page connectedness (and hence social evaluation of its worth) into its evaluation of which pages best match a user's query. Unfortunately, this strategy does not work for certain genres of hypermedia that are visited and sought, but that do not have many incoming links (again, pornography sites—a significant genre on the Web—provide good examples of this effect). Thus search engines have become invisible intermediaries between readers and authors of the hypermedia on the Web in a way that "Seven Issues" never anticipated.

The Distinct Role of Hypermedia Readers

In the salad days of hypermedia systems, we purposefully ignored the divergence of needs between readers and writers. Readers, we reasoned, were on equal footing with writers. At any time, they could fruitfully switch roles; in constructing meaning, a

reader could transform into an annotator, and the annotator into a contributor (Michalak and Coney, 1993).

Since then, it has become clear that there are many more readers than writers and that they occupy quite different niches in the hypermedia ecology (Pitkow, 1998). Readers are, in essence, gatherers (Rosenberg, 1996). Michel de Certeau is quoted by the historian Roger Chartier as offering the following almost poetic distinction between readers and writers:

Far from being writers—founders of their own place, heirs of the peasants of earlier ages now working on the soil of language, diggers of wells and builders of houses—readers are travellers; they move across lands belonging to someone else, like nomads poaching their way across fields they did not write, despoiling the wealth of Egypt to enjoy it themselves.... Reading takes no measures against the erosion of time (one forgets oneself and also forgets), it does not keep what it acquires, or it does so poorly, and each of the places through which it passes is a repetition of the lost paradise (Chartier, 1994, p. 1).

How does this distinction play out in Halaszian terms? "Seven Issues" identified no technologies for readers (save the acknowledgment that readers might interact with a separate presentation layer and that hypermedia materials might adapt to a reader's specific interests). In "As We May Think," Vannevar Bush (1945) identified at least one new role and an associated technology for readers, that of way-finding or path-making to link together formerly disparate bits of information (displayed on a double-screen Memex). I would contend there are at least five other reading-oriented practices that bear some investigation when one looks at hypermedia in use: (1) personal annotation (as opposed to published annotation); (2) re-retrieval; (3) gathering (the intentional amassing of post-retrieval content with a specific purpose in mind); (4) contextual access given, for example, mobile devices; and (5) collaborative reading. Each suggests an approach to defining new technology.

Personal Annotation

Hypertext is fundamentally annotative, a characteristic manifested in the many mechanisms and systems for collaborative writing and scholarly commentary. But these mechanisms are not always appropriate for personal annotation, an unselfconscious practice that goes hand in hand with many reading tasks (not leisure reading, but reading that is done with a purpose in mind). Nor is there much thought about how personal annotations might be useful to a community of hypertext readers in much the same way that Amazon uses the reading and purchasing habits of anonymous readers like oneself to make recommendations.

O'Hara and Sellen found three characteristics of personal annotations on paper that distinguished them from their counterparts online (O'Hara and Sellen, 1997). First, annotation on paper was smoothly integrated with reading; on-line annotation was distracting. Second, paper better supported marking "on" the source document; people wanted their marks to be distinguishable from the source document. Finally, the subjects of their experiment also took notes on separate sheets or in other windows; note-taking on paper was more closely interleaved with reading. In addition to these findings, my own field studies showed personal annotations to be highly varying in their status and value to other readers (as well as to the original annotator) and may require new strategies and abstractions to make them useful to other readers (Marshall, 1998).

Pen-based systems running on tablet computers that use free-form digital ink promise to help readers annotate hypertexts in a familiar way (Schilit et al., 1999). Paper-like user interfaces for hypertext readers support natural interaction with textual and graphic content, while they go beyond paper to support hypermedia capabilities (e.g., link traversal, the display of multimedia content, search and query, and so on).

Re-retrieval

There has been a sea change in the way people interact with digital resources in general. The emerging philosophy seems to be, "If I've found it once, I can find it again." People are far more willing to toss paper rather than collect it. Pitkow (1998) also cites a statistic that roughly half of the Web content is requested more than once by the same client-side computer. On the flip side, as an increasing number of digital documents cross our computer desktops—many times as hypertext ephemera, rendered bits from the network that appear briefly as we traverse links to new sites—we are moved to wonder, "I know I've seen that somewhere. Where did I read that?" Hence re-retrieval becomes an important part of interaction with hypertext, a partner to search and query. In fact, in a recent project of ours, we found that one of the most important questions members

of our user community wanted to answer for themselves was, "What did we say about this last time? Why did we say that?"

Right now, we mainly treat re-retrieval as a problem to be addressed either via conventional search engine mechanisms or via browsing history lists. Additional research is necessary to know how we should handle those documents that fall through the cracks between our personal information stores, where there is much promising work that transcends the limits of today's file-system-based approach (Dourish et al., 1999), and the growing hypermedia resources that are available online through the Web, where there is tremendous push to develop effective search and query mechanisms.

Gathering

Related to our last reading-related issue of re-retrieval is a phenomenon we might think of as gathering, the amassing of materials in service of a particular task (be it writing, research, or analysis, to name a few such tasks). Of course, gathering is not a new practice. Many of us maintain extensive physical filing systems for our personal archives, or work in an office in which every level surface is covered with semi-sorted stacks of working papers. But gathering is increasingly performed electronically, especially given the noticeable shift toward writing as bricolage, a practice in which writers appropriate parts of existing documents[5] and hypermedia networks, and incorporate them into new material. Hypermedia is an ideal medium for such bricolage, since it already has mechanisms, most notably links, that support the properly attributed re-use of materials. But, as Halasz noted (1988), the structure of hypermedia materials in an exploratory task is often unknown at the outset of the task; rather, it emerges over time.

The observation that structure is emergent and evolves over time led Halasz to conclude that virtual structures are the necessary remedy (Halasz, 1988). Virtual structures work in some situations in which a presentational structure is the goal, but not in others, in which a reader is exploring and interpreting materials, and needs to be in control of the structural fluidity him or herself. Readers' use of search engines is an interesting illustration of this distinction. A search engine returns a simple virtual structure—a ranked list—that is generated algorithmically. However, when the reader turns into a gatherer, he or she might also create something like a ranked list, but it is a handmade evolving structure that reflects the current state of his or her understanding of the materials and how they fit the task. To address this requirement, we turned to spatial hypertext in which the structures are ambiguous, fluid, and implicit, but not necessarily computed (Marshall and Shipman, 1995).

Gathering tools, as they are construed today, allow readers to gather references to Web documents rather than the documents themselves. This kind of indirection is link-like and stays true to the hypermedia paradigm, but introduces an element of ambiguity. Is the reader referring to the specific page? Or is the reader referring to the place (thus allowing content updates to occur)? The answer is usually a rousing "it depends." Careful gatherers may distinguish between the two cases, and choose to make a version of the content local when they want the content to remain fixed.

Is this merely, then, a recapitulation of Halasz's versioning issue? Perhaps, but when we consider gathering as a reader activity, what we are really concerned with is archiving (i.e., keeping the version that the reader is referring to intact). It is only the reader's reference that makes the node worth preserving, rather than automatic tracking of all of the author's changes. This switch in perspectives may seem like a nuance of interpretation, but on the Web, we can all think of very real examples of authors or publishers working at cross-purposes to readers when a Web page is removed from the hypertext network. Who controls node fixity and persistence is essentially a copyright/fair use issue; Samuelson and Davis (2000) discuss these sometime-divergent intellectual property interests. Right now it seems like there is no simple answer. If we apply Levy's sensible thoughts on archiving (1998), the fact that the page is in use makes it worth saving, especially taken in tandem with Pitkow's characterization of the life span of Web documents as being 50 days (Pitkow, 1998). Thus, from an author's point of view, versioning is a service that is most appropriate to meet the requirements of specific applications, for example, software engineering; from a reader's point of view, archiving is a service that may well be necessary for gathering.

Contextual Access

Where do we read? Throughout Halasz's "Seven Issues" discussion, there is an underlying assumption

that we will go to the hypertexts; we will access them through workstations and kiosks. In fact, dating back to Ted Nelson's SilverStands, the local outlets for Xanadu (Nelson, 1981), the hypertexts are accessible through stationary portals.

In the current technological climate, there is a trend toward mobile access to the Web, to finding the information you want where and when you want it, in the context that is most appropriate. Of course, the mobile telephony-based interfaces to the Web have proven to be frustrating and haven't taken off the way the mobile access providers have tried to convince us they would. Cellular phones and two-way pages are still mostly used for communication.[6]

On the other hand, if we look at devices with slightly larger form factors, devices more capable of displaying the kinds of hypermedia one finds on the Web—Pocket PCs, Palm Pilots, electronic book devices, and laptop computers—the story changes. Then contextual access to hypertexts and digital libraries through wireless connections begins to make sense to some of the traditional mobile knowledge worker constituencies such as journalists and legal professionals (Fagrell et al., 2000; Elliott, 1995).

Design challenges then begin to enter into the picture: how can we design hypermedia for radically different modes of access? How does the reading experience change with the form factor? For example, many of the overview techniques authors and designers use to maintain reader orientation (see the navigation issue, Halasz, 1988) are no longer appropriate; display limitations, coupled with the nature of the users' tasks, make them irrelevant.

Collaborative Reading

Collaboration is one of the issues Halasz calls out as a central research direction, and rightly so. In many ways, the Web itself can be thought of as the realization of collaborative hypertext. But I would like to single out collaborative reading as a phenomenon that begs further attention; much of Halasz's rhetorical force was aimed at collaborative authoring, and much of the work we (my collaborators and I) have focused on has involved explicit instances of collaborative interpretation and shared information spaces. Of course, the only way the Web can handle these tasks (collaborative authoring and shared interpretation) is with additional infrastructure. The substance of Issue 6 thus still stands as one of my own key concerns.

Yet we also see much lighter weight collaborations arising on the Web. People read together, navigate together, recommend sites to each other, and form implicit shared understandings of worthwhile resources. In fact, there is frequently a meta-narrative that arises when people read hypertexts together (either in face-to-face situations or using another communication mode like the telephone or IRC). It is through this narrative that hypertexts become mutually intelligible. The research here might well be to help record and reuse some of the transient forms that now surround collaborative reading (see, for example, Marshall et al., 1999).

Formality as a Leitmotif

Would the Web have taken off if it were as complicated as the hypertext systems and models that Halasz envisions in "Seven Issues"? Halasz himself notes (in a section entitled "NoteCards in Use") that links used in an idea structuring and authoring application were "as a rule ... 'See' or 'Unspecified' links and were placed at the end of the card of the card's text" (Halasz, 1988, p.839). In fact, much of the use I observed (especially at sites I supported outside of Xerox PARC) did not involve typed links. Link-typing schemes, if adopted, were frequently abandoned. The automatically maintained hierarchical FileBox links were a far more common way of introducing structure to the hypertext. Spatial means of structuring in which links were implied rather than declared were common as well (Marshall and Shipman, 1993).

Technologically, then, typed links—in fact, relations as full-fledged objects—seemed desirable. Our post-NoteCards system development effort incorporated relations rather than links (Marshall et al., 1991). But practically our experience with use showed them to be problematic (Marshall and Rogers, 1992) and led us to develop a much simpler hypertext model with implicit links; in this system, visual and spatial properties stood in for more explicit notions of typing (Marshall and Shipman, 1995). What we were led to believe from a larger set of experiences—ours and those of other researchers—was that, when formal structures such as typed links showed up in the user interface, interaction difficulties often followed. In fact, users would reject the system outright (or simply ignore the formalisms) unless the system demonstrated a clear benefit from the adoption and expression of these structures (Shipman and Marshall, 1999).

Indeed, what we have all experienced on the Web is that much explicit markup is directed toward creating pages with the right appearance. Far less of the markup reflects a page's functional structure (e.g., heading level). It is easy to envision that markup which reflects content semantics (as, for example, specified by the Text Encoding Initiative) would be very rare. Most professionally designed Web sites devote significant attention to getting the right look (that is, a look that is appropriate to their genre), rather than to expressing structure that would, say, be useful to an automatic summarizing program.

This basic tension between formal structures and informal practices has yet to be resolved and brings us to a further provocative question: What will happen as the Web standards become increasingly complex, as markup tilts toward XML and CSS? As link representations become richer (e.g., XLink)? Will the payoff be sufficiently great to lure the extended community of hypertext authors into using them? Will they use them consistently, and in a manner that fits readers' needs?

As was true with Halasz's "Seven Issues," it is through use over time that this new set of issues will be illuminated and potentially resolved.

Endnotes

[1]I use the terms interchangeably here, since the original notion of hypertext was not intended to preclude non-textual media types.

[2]Subsequent open hypermedia developments made a concerted effort to get around all of these limitations—especially the intertwining of structure and content—but, relying on the most basic markup and protocol specification, the Web is a stateless mix of structure and content.

[3]Halasz specifically suggested a philosophical distinction between Navigators and Architects, between Literalists and Virtualists, between Card Sharks and Holy Scrollers, and between the Literati and the Engineers.

[4]In fact, it is informative (and occasionally entertaining) to watch one of the so-called search engine voyeurs as it ticks past, for example, http://www.excite.com/search/voyeur/ or http://www.infotiger.com/voyeur. It is easy to see that people look for everything under the sun on the Web.

[5]In a corporate environment, this seems to be manifested by reuse of individual PowerPoint slides.

[6]Even the popular text-based i-mode cellular phones in Japan or SMS phones in Europe are still primarily about communication, not about information access.

References

Akscyn, R.M., McCracken, D. L., and Yoder, E. A. (1988). KMS: A distributed hypermedia system for managing knowledge in organizations. *Communications of the ACM,* 31(12), 820-835.

Bolter, J. D. (1991). *Writing space: The computer, hypertext, and the history of writing*. Hillsdale, NJ: Lawrence Erlbaum and Associates.

Chartier, R. (1994). *The order of books.* Stanford, CA: Stanford University Press.

Dourish, P., Edwards, K., LaMarca, A., and Salisbury, M. (1999). Presto: An experimental architecture for fluid interactive document spaces. *ACM Transactions on Computer-Human Interaction*, 6(2), 133-161.

Elliott, M. (1995). Digital library design for organizational usability in the courts. [http://edfu.lis.uiuc.edu/allerton/95/s3/elliott.html].

Fagrell, H., Forsberg, K., and Sanneblad, J. (2000). FieldWise: A mobile knowledge management architecture. in *Proceedings of CSCW 2000*. New York: ACM Press.

Goodman, D. (1987). *The complete HyperCard handbook.* New York: Bantam Books.

Halasz, F. G. (1988). Reflections on NoteCards: Seven issues for the next generation of hypermedia systems. *Communications of the ACM,* 31(12), 836-852.

Halasz, F. G. (1991). "Seven issues": Revisited (keynote address, Hypertext '91 Conference, San Antonio, Texas, December 18, 1991. [http://www.parc.xerox.com/spl/projects/halasz-keynote/].

Halasz, F. G., Moran, T. P., and Trigg, R. H. (1987). NoteCards in a nutshell. In *Proceedings of the ACM CHI + GI Conference,* New York: ACM Press, pp. 45-52.

Levy, D. M. (1998). Heroic measures: Reflections on the possibility and purpose of digital preservation. In *Proceedings of ACM Digital Libraries '98*, New York: ACM Press, pp. 152-161.

Marshall, C. C. (1998). Toward an ecology of hypertext annotation. In *Proceedings of ACM Hypertext '98*, New York: ACM Press, pp. 40-49.

Marshall, C. C., and Shipman, F. M. III. (1995). Spatial hypertext: Designing for change. *Communications of the ACM*, 38(8), 88-97.

Marshall, C. C., and Shipman, F. M. III. (1993). Searching for the missing link: Discovering implicit structure in spatial hypertext. In *Proceedings of Hypertext `93*, New York: ACM Press, pp.217-230.

Marshall, C. C., and Rogers, R. A. (1992). Two years before the mist: Experiences with Aquanet. In *Proceedings of the ACM Conference on Hypertext (ECHT `92)*, New York: ACM Press, pp. 53-62.

Marshall, C. C., Halasz, F. G., Rogers, R. A., and Janssen, W. C. Jr. (1991). Aquanet: A hypertext tool to hold your knowledge in place. In *Proceedings of Hypertext `91*, New York: ACM Press, pp. 261-275.

Marshall, C. C., Price, M. N., Golovchinsky, G., and Schilit, B. N. (1999). Collaborating over portable reading appliances. *Personal Technologies*, 3(1).

Michalak, S., and Coney, M. (1993). Hypertext and the author/reader dialogue. In *Proceedings of the fifth ACM conference on Hypertext*, New York: ACM Press, pp. 174 – 182.

Nelson, T. H. (1981). *Literary machines*. Swarthmore, Penn.

O'Hara, K., and Sellen, A. (1997). A comparison of reading paper and on-line documents. In *Proceedings of CHI'97*, New York: ACM Press, pp. 335-342.

Pitkow, J. E. (1998). Summary of WWW characterizations. *Web Journal* 2(1-2), 3-13.

Rosenberg, J. (1996). The structure of hypertext activity. In *Proceedings of Hypertext '96*, New York: ACM Press, pp. 22-30.

Samuelson, P., and Davis, R. (2000). The digital dilemma: A perspective on intellectual property in the information age. Presented at *The 28th Research Conference on Communication, Information and Internet Policy* (Alexandria, Virginia, September 23-25, 2000).

San Jose Mercury News (1997). As online worlds collide, dueling innovations net unexpected results, March 2.

Schilit, B. N., Price, M. N., Golovchinsky, G., Tanaka, K., and Marshall, C. C. (1999). As we may read: The reading appliance revolution. *IEEE Computer* 32(1), 65-73.

Shipman, F., and Marshall, C. (1999). Formality considered harmful: Experiences, emerging themes, and directions on the use of formal representations in interactive systems. *Computer Supported Cooperative Work (CSCW)*, 8(4), 333-352.

Yankelovich, N., Meyrowitz, N., and van Dam, A. (1985) Reading and writing the electronic book. *IEEE Computer*, October 1985.