

Design Issues for Multi-Document Hypertexts

Robert J. Glushko

Search Technology, Inc.
4725 Peachtree Corners Circle, Suite 200
Norcross, GA 30092

INTRODUCTION

Vannevar Bush conceived of hypertext as the "computer glue" that binds information from a wide variety of books, documents, communications, and other artifacts to enhance its accessibility and usefulness. However, most of the recent hyper-activity in research labs and in the marketplace falls short of Bush's vision. Most hypertext software is oriented toward hypertext as a new form of writing via incremental combination of bits and pieces of information. These hypertext programs typically provide little support for converting existing information from its more linear printed form. Where hypertexts have been created from existing text, they generally have been converted from a single encyclopedia ([Glus88], [Oren87]), a single reference document ([Fris88], [Perl88], [Raym88]), or a single system's documentation ([Egan89], [Walk88a]). Hypertexts that integrate the complete contents of more than one book or large document seem nonexistent, even though the expected benefits from such multi-document hypertexts were the original motivation for the concept.

In this paper I describe the design issues that have arisen in a project at Search Technology to build a multi-document hypertext whose major components are a 3000-page engineering encyclopedia and a 388-page standards manual:

- * **Document selection** – Which documents should be included?
- * **Extent of integration** – How tightly linked should the documents be?
- * **Locus of integration** – Where in the documents should the links be?
- * **Type of integration** – What kinds of links should there be within and between the documents?

DOCUMENT SELECTION

Selecting documents to incorporate in a hypertext is not a simple problem. The goal should not be to build a hypertext encyclopedia, a hypertext reference manual, or a hypertext version of system documentation. Instead, the goal should be to enhance the performance or enjoyment of people performing their jobs, using hypertext when it is appropriate to improve the access to and the usefulness of information. This view implies that the choice of documents to include in a hypertext should be dictated by a user and task analysis. Only when the intended users must combine information from more than one document to meet the information needs of the tasks they carry out is a multi-document hypertext called for.

For example, I use both a printed phone book and a printed software manual on my desk, and each would be more useful if I could browse and search them in hypertext form on my computer. Nevertheless, I need to strain to imagine a task for which I need hypertext connections between an electronic phone book and an electronic software manual. These documents seem too dissimilar in content, format, and intended purpose to justify or benefit from hypertext connections to each other.

So when should two documents be integrated in a multi-document hypertext? The motivation for combining more than one document in a hypertext is the conviction that each document adds value to the others, that the whole is greater than the sum of the parts. If two documents were written with the same purpose for the same audience with the same content, a user would only need one of them. Only when the documents differ in some significant way can they complement each other. However, to the extent that the documents differ, it must be true that from the perspective of the user of one document, at least some of the information in the other document is irrelevant some or all of the time. A final implication is that complementarity is not necessarily symmetric. A document can enhance the value of another without its use being enhanced in return.

After the task and user analysis has identified a document to include, it can be important to make similar (though certainly less thorough) evaluations about whether to include any of the documents that it cites. How well each cited document complements those already included is again the major consideration.

Case Study

These document selection issues can be explained more concretely using our current project as a case study. Our goal is to make it easy for designers to find and apply the information they need to build systems that are more consistent with the capabilities and limitations of the humans who operate those systems. Our initial project ([Glus88], [Glus89]) was to design an electronic version of the *Engineering Data Compendium* [Boff88], a 3000-page multi-volume engineering encyclopedia whose 1138 entries contain comprehensive information about human perception and performance. The *Compendium's* intended users are designers of systems with complex user interfaces who need detailed information about human perception and performance to make design tradeoffs but who typically have education and experience in another technical specialty [Linc88]. The *Compendium* entries are arranged in twelve major topic areas that roughly reflect a scientific taxonomy of human perception and performance; topic areas include vision, audition, other senses, perceptual organization, attention, and so on. In printed form, the *Compendium* is organized in three volumes that contain the entries and a fourth volume that contains a user's guide, index, glossary, and other aids for finding relevant entries.

The *Compendium* emphasizes the basic parametric data of human perception and performance, tabular summaries of data from research studies, predictive models, and perceptual and performance principles. For example, the *Compendium* takes 25 pages to relate the parametric details and predictive models for various phenomena of visual flicker. In contrast, the *Compendium* contains little explicit advice on the appropriate application of these basic data on visual flicker when designing visual displays. As a result, we sought documents that would complement the *Compendium's* comprehensive data with design guidelines in the context of specific applications.

Bill Cody [Cody89], Dan Sewell and others of us at Search Technology have been studying how designers identify, locate, and use information about human perception and performance to make design tradeoffs and meet system requirements. Among the documents used by designers is a 388-page military standard for human engineering design whose official title is *MIL-STD-1472-D, Human Engineering Design Criteria for Military Systems, Equipment, and Facilities* [DoD89]. *1472-D* is organized according to an equipment-oriented or task-oriented taxonomy, with sections for visual displays, audio displays, controls, control/display integration, anthropometry, and so on.

We chose to incorporate *1472-D* in our multi-document hypertext after in-depth analysis of several candidate documents and interviews with the users, editors, and authors of the *Compendium* and *1472-D* to determine the extent and nature of their overlap. *1472-D* is broader in coverage than the *Compendium*, contains more qualitative guidance on human engineering, and gives prominence to rules of thumb and expert opinion. *1472-D*'s intended users are specification writers, designers looking for constraints or ideas for a particular design problem, or testers who use it as an evaluation checklist to verify compliance with the original specification. Taken together, this means that *1472-D* complements the *Compendium* more than vice versa. A designer may make more sense of quantitative information on perception and performance because of the context provided by *1472-D*'s application guidelines, yet a user contractually required to follow *1472-D*'s guidelines may not be interested in the *Compendium*'s detailed scientific rationale for them.

1472-D poses a challenge for us with respect to the documents that it cites that we call the **tiering problem**. In the jargon of standards, the first tier is defined as the set of documents referred to by a given standards document. The second tier includes the set of documents cited by documents in the first tier, and so on. Standards documents are carefully modularized into many narrow documents to make them easier to maintain, but one standard may cite other standards and in effect treat them as if they were bundled into a single standard. That is, the first tier of references from the standard to other standards is often considered to be binding.

1472-D is explicit in its tiering requirement and lists 43 other standards documents that "form a part of this document to the extent specified herein (p. 3)." Therefore, if we choose not to include the cited sections of these standards in our multi-document hypertext, our system will be significantly less helpful for someone whose primary task involves using *1472-D*. At the same time, however, not including these documents cited by *1472-D* will have negligible impact on someone whose primary task in the multi-document hypertext involves using the *Compendium*.

The way in which documents complement each other in support of the intended users for their primary tasks should be the most important issue in selecting documents for multi-document hypertexts. But we did not discount more pragmatic considerations in choosing *1472-D*. The availability of computer-readable source files, the cooperation of its author, and its lack of copyright as a government document all contributed to our decision to use it.

EXTENT OF INTEGRATION: HOW TIGHTLY LINKED ARE THE DOCUMENTS?

While to some it may strain the definition of hypertext, it is obvious that a user may benefit from a multi-document hypertext in which each document contains links within it but no links to any other document; this is the situation I mentioned earlier with my hypothetical electronic phone book and software manual. Indeed, few or no links between documents may be a desirable goal of effective modular design in which document units are self-contained with minimal relationships to other units [Walk88b]. From a user's perspective at the user interface, it is easy to understand the extreme cases of zero integration and total integration.

Zero integration of two hypertext documents means that they appear totally separated and unrelated. The user selects a document to work with, but once this choice is made, no trace of the second document is ever seen as the user moves among different parts of the first document. Neither the entry points (like a Table of Contents or Index) nor the text units of the first document refer to parts of the second document. If the user conducts a search, the candidate list contains units from only one document. Once selected, a candidate text unit may contain cross-references to other parts of the same document, but not to the second document.

Total integration. At the other extreme, if two documents are totally integrated, the user has no separate view of each document. Instead, the user sees a collection of interconnected text units in which the idea of "belonging" to separate documents is not apparent. The user specifies some starting point in the text and sees no contextual boundaries that separate one document from another as different text units are viewed by traversing links between them. The starting point might be determined explicitly when the user starts to browse at some specified location (e.g., the "Home" card in *HyperCard*) or implicitly (e.g., by a query that generates a set of candidate text units), but in neither case is there a view that explicitly reflects the structure of separate documents. This extreme is the "seamless integration" of multiple documents expressed in visionary claims for hypertext like those of Ted Nelson's *Xanadu* system [Nels80].

There is a basic tradeoff here between the need for integration at entry points and the need for integration at the unit level. The more structure provided by the entry points, the more precisely the relevant text units can be located. For example, a Table of Contents that lists chapter headings and sub-headings will require the user to browse less within a chapter than one that lists only chapter titles. In contrast, a hypertext that contains an unstructured and coarse-grained entry point (like a "Home" card) will necessarily require complex links among units if readers are to find relevant information.

In between the extremes of zero integration and total integration is a wide range of choices about the extent of integration of multiple documents that appears in the hypertext user interface. Again, it is useful to begin by contrasting the clear endpoint cases before considering intermediate ones.

LOCUS OF INTEGRATION: WHERE ARE THE LINKS?

Integration solely in text units

Multiple documents have always been partly integrated solely by explicit interconnections of their component text units, with their entry points remaining totally separate. That is, documents may each have their own Table of Contents, but their component text units may contain cross-references to other documents. This is the model of integration of multiple documents used in most hypertext systems.

Integration solely in entry points

Multiple documents can be partly integrated solely by combining the entry points from which users generate a set of candidate text units for more detailed viewing. Conceptually, this is analogous to traditional bibliometric retrieval systems like *Dialog* or *Lexis*, in which different documents (or abstracts) are indexed together. A user query results in a list of candidates that satisfy the query. **The candidate text units are not themselves integrated in any way.**

A more sophisticated and user-oriented approach to integrating multiple documents solely by integrating their entry points would merge their Tables of Contents or the "back of the book" Indexes. It might be desirable to use the Table of Contents of the document with the broadest coverage as the "skeleton" into which the other documents are merged. A combined Table of Contents or Index might look like this:

- 1.0 Main topic A
 - document 1, section 1.2
 - document 2, section 2.4
- 1.1 Topic B
 - document 1, section 1.2.1
 - document 3, sections 6.5-6.9
- 1.2 Topic C
 - document 1, section 1.2.2
 - document 2, section 2.4.2
 - 1.2.1 Subtopic D
 - document 1, section 1.2.2.1
 - document 6, section 4
- 2.0 Main topic E
 - document 4, section 3
 - document 5, section 5.2.1
- 3.0 Main topic F
 - ...

In this example, the multiple documents are integrated to the extent that any candidate list a user obtains from the merged entry points contains text units from more than one document. But as with bibliographic retrieval systems, none of the candidate units contains any explicit cross-references or links to other units.

Integration both in entry points and in text units

Now it should be easy to envision hybrid forms of multiple document integration in which the user interface combines documents both at their entry points and in their text units. This is the approach we have chosen in our combination of the *Compendium* and *1472-D* in a multi-document hypertext. At the entry points, we will augment the existing Tables of Contents and Indexes with several additional taxonomies that span the topics contained in the two books. These are being developed by Janet Lincoln, one of the editors of the *Compendium*. We are also developing a framework for deciding which links to create within and between the two documents that is described later in this paper.

HOW MUCH INTEGRATION OF MULTIPLE DOCUMENTS IS DESIRABLE?

Our own experience and that of other hypertext designers has shown that usability is enhanced by use of the explicit structure of documents, so we reject the *Xanadu* hypertext model in which multiple documents are integrated to the extent that their separate structures disappear. The task and document analysis we conducted to design the hypertext version of the *Compendium* led us to exploit the explicit structure of the *Compendium* in various hierarchical browsers for the Table of Contents, Index, and other access points, and to emphasize these structured entry points over context-free Boolean search techniques ([Glus88], [Glus89]). Similarly, our evaluation of other compact disc encyclopedias showed us convincingly that failing to use the explicit structure of documents and relying solely on full-text search can make it extremely difficult to find information in them.

The limited experimental literature on hypertext suggests that excessive integration through large numbers of links creates unusable "spaghetti documents" [Foss88]. Excessive linking causes serious problems of disorientation and cognitive overload for the user because it destroys most of the structural and contextual cues [Conk87]. Several researchers are working on new representations of hypertext link structures that may help to solve these problems ([Conk88], [Fein88], [Furn86]), but limiting the links in the first place seems like a more practical solution.

Our design team has chosen a moderate degree of multi-document integration in which users can readily see that documents maintain their separate identity but are enhanced by

cross-references to information in other documents. The user selects one document as the primary focus of browsing or searching, but related information in other documents can be retrieved. The key difference is that choosing to view cross-referenced information from another document does not automatically switch the user's context to the second document.

Instead, in this middle range of integration the user has some control over the extent of contextual preservation or integration across different documents. Tom Coonan has been designing several alternative user interfaces that support this intermediate level of integration. The user can choose to display cross-referenced documents without changing context -- perhaps to view it in a pop-up window that temporarily overlays the current document. If the user chooses to change to the second document, he or she would then be prompted to decide how much of the current context should be carried along. For example, the "bookmarks" for the parts of the first document that the user has already viewed might be left behind or merged with the bookmarks for the second one.

I have presented the extreme positions on integration in this paper to make an intermediate approach seem reasonable at face value. However, the fact remains that the design question of **How tightly linked should the documents be?** is a legitimate research issue that has not yet been confronted seriously. Raymond and Tompa's [Raym88] innovative and thorough description of the cross-reference structure of the *Oxford English Dictionary* provides some foundations, however.

A FRAMEWORK FOR LINKING IN MULTI-DOCUMENT HYPERTEXTS

Much of the research in hypertext concerns the number and type of links among the units ([Conk87], [Hala88]), and we agree that link typing will enable great flexibility and power in hypertext representation. But from a practical standpoint, we have found it necessary to work at a coarser level of analysis and start with more basic questions about the kinds of links to create. First, we contrast **intra-document** links with **inter-document** links. Next, we contrast **explicit** links with **implicit** ones. Taken together, these two dimensions characterize four broad classes of links that I will now discuss in the order that we think they should be considered in the design of multi-document hypertexts.

Explicit Intra-document links

Links of this class explicitly connect two parts of the same document, and include footnotes, "See also" cross-references, and pointers to figures, illustrations, or other non-textual components. These links are the easiest to identify when converting existing texts to hypertexts and they are probably the most usable and useful as well. Since good writers use these structural and rhetorical devices in predictable ways, it follows that a reader of the hypertext will find them predictable as well. People who have never heard the word "hypertext" know what footnotes are for and have reliable expectations about the likely number, length, relevance, and value of footnotes¹. As a result, when printed documents are converted to electronic ones it is essential to exploit this sort of knowledge by capturing the explicit intra-document links first².

Implicit Intra-document links

These links are those that are part of the logical structure of the document but which may be impossible to make explicit in the printed form because of limitations in the medium, the nature of the writing and production process, or publishing conventions. For example, every explicit "See also" link in a printed document implies a "Cited by" link in the opposite direction. Likewise, the first appearance of a Glossary item in the text can readily be linked to its definition in the "back of the book" Glossary in electronic form. We think that links of this type pose little risk of disorientation or cognitive overload because they follow naturally from the printed version of the document, so we consider them nearly as important as the explicit intra-document links.

Explicit Inter-document links

Like the explicit links within a document, these links are easy to identify because they follow presentation conventions in the printed document and are often collected in reference or bibliography sections. Yet we think that they pose more challenges for the hypertext designer and reader than intra-document links, because it is much harder to predict the extent or usefulness of the information at the end of the link. An author may cite another document for many different reasons and the cited documents may add little value. To link or not to link hinges on the issue of complementarity that I discussed under the topic of document selection.

Implicit Inter-document links

These are the links that might be closest to the vision of hypertext, namely links that are not explicit between related documents but that can be extracted by careful and creative analysis of the two texts and the relationship between them. But as George Landow

¹ Expectations about footnoting conventions may vary significantly for different readers and different kinds of documents. Practicing lawyers depend on the fact that law review articles typically contain hundreds of detailed footnotes to all the relevant cases, and are not surprised if the footnotes comprise 50% or more of the total text. Other disciplines might reject this footnoting style as excessively pedantic.

² We warn, however, against slavish mechanical translation of footnotes into hypertext links, because footnote structure and semantics can be complex. For example, a document may contain footnotes to other footnotes (or even to the footnotes of another document), or may use footnotes as an indirect addressing technique to refer to a section of text (so that cross-references can be specified prior to page layout). We doubt that the design and implementation contortions required to preserve exactly this granularity and complexity of hypertext structure are worth it. See footnote 1.

[Land87] noted, there is no consistent "rhetoric of hypertext" that makes it easy for a reader to understand what such links mean and what is likely to appear at the link destination. Doland [Dola89] has cautioned that creating links is an ideological act that may limit "interpretive diversity."

When we first began working in hypertext several years ago, we expected that it would soon be possible to extract these implicit links automatically with natural language processing or clever indexing techniques (see [Fox88]), but we have been disappointed so far and we are starting to conclude that implicit intra-document links are best identified by the hypertext reader. Mark Weaver came up with an analogy to used textbooks that helps to explain why. Weaver noted that if we created links based on our understanding of the document, some of the links won't fit with the understanding and context brought to the hypertext by another reader. Like a used textbook, some of the highlighting and margin notes may be useful to another student, but may be distracting or misleading at other times. We have decided in our current project to provide functions that make it easy for readers to create private links and notes rather than try to create many of them ourselves.

SUMMARY

Hypertexts that combine the complete contents of more than one book or large document come closer to the spirit of hypertext than hypertexts that are created from a single source or created incrementally from document bits and pieces. In designing such a multi-document hypertext, our team has identified several key design issues and devised analyses and practical rules of thumb for resolving them:

- * Select documents to include based on a user and task analysis. The extent to which the documents complement each other for the intended users and tasks will determine the extent to which it makes sense to combine them with hypertext links.
- * Exploit the user's experience with the printed document by appropriately transforming the Table of Contents, Index, and other entry points into electronic equivalents.
- * Devise composite entry points like hybrid tables of contents or new taxonomies to provide an understandable perspective that spans the set of documents.
- * Emphasize the links within each document before trying to support the explicit and implicit links between documents.

Our experience in this and previous projects confirms to us the attractiveness of the hypertext vision in enhancing the accessibility and usefulness of information, but also emphasizes the need for a disciplined approach to "hypertext engineering" [Glus88] that rests on task analysis, document analysis, and careful consideration of lessons from the technical writing and information retrieval disciplines.

ACKNOWLEDGMENTS

This work was primarily supported by the Designer's Associate project funded by the Air Force's Armstrong Aerospace Medical Research Laboratory at Wright-Patterson AFB, Ohio under contract AF #F33615-86-C-0542. Dr. Kenneth R. Boff is the program sponsor. My colleagues at Search Technology, especially Bill Cody, Tom Coonan, Dan Sewell, Mark Weaver, and Brad Wiederholt helped shape the ideas in this paper. Jonathan Grudin, Janet Lincoln, Gary Perlman, Pamela Samuelson, and Jan Walker also made important contributions.

REFERENCES

- [Boff88] Boff, K. R., and Lincoln, J. E. *Engineering Data Compendium: Human Perception and Performance*. AAMRL, Wright-Patterson AFB, OH., 1988.
- [Cody88] Cody, W. Designers as users: Design supports based on crew system design practices. *Proceedings of the 45th Annual Forum of the American Helicopter Society*, Boston, MA., 1989.
- [Conk87] Conklin, J. Hypertext: An introduction and survey. *Computer*, 20(9), 17-41. September 1987.
- [Conk88] Conklin, J., and Begeman, M. gBIS: A hypertext tool for exploratory policy discussion. *CSCW '88: Proceedings of the Conference on Computer-Supported Cooperative Work*, 140-152, 1988.
- [DoD89] Department of Defense. *MIL-STD-1472-D, Human Engineering Design Criteria for Military Systems, Equipment, and Facilities*. Washington, D.C., 1989.
- [Dola89] Doland, V. Hypermedia as an interpretive act. *Hypermedia*, 1(1), 6-19, Spring 1989.
- [Egan89] Egan, D., Remde, J., Landauer, T., Lochbaum, C., and Gomez, L. Behavioral evaluation and analysis of a hypertext browser. *Human factors in computing systems. CHI '89 Conference Proceedings*, ACM: New York, 205-210, 1989.
- [Fein88] Feiner, S. Seeing the forest for the trees: Hierarchical display of hypertext structure. *Proceedings of the ACM Conference on Office Information Systems*, 205-212, 1988.
- [Foss88] Foss, C. Effective browsing in hypertext systems. *User-oriented content-based text and image handling. Proceedings of the RIAO Conference*, March 21-24 1988, 83-98.
- [Fox88] Fox, E., Weaver, M., Chen, Q., and France, R. Implementing a distributed expert-based information retrieval system. *User-oriented content-based text and image handling. Proceedings of the RIAO Conference*, March 21-24 1988, 708-726.
- [Fris88] Frisse, M. Searching for information in a hypertext medical handbook. *Communications of the ACM*, 31(7), 880-886, July 1988.
- [Furn86] Furnas, G. Generalized fisheye views. *Human Factors in Computing Systems*. ACM: New York, 16-23, 1986.
- [Glus88] Glushko, R. J., Weaver, M. D., Coonan, T. A., and Lincoln, J. E. "Hypertext Engineering": Practical methods for creating a compact disc encyclopedia. *Proceedings of the ACM Conference on Document Processing Systems*, 11-20, 1988.
- [Glus89] Glushko, R. J. Transforming text into hypertext for a compact disc encyclopedia. *Human factors in computing systems. CHI '89 Conference Proceedings*, ACM: New York, 293-298, 1989.

- [Hala88] Halasz, F. Reflections on NoteCards: Seven issues for the next generation of hypermedia systems. *Communications of the ACM*, 31(7), 836-852, July 1988.
- [Land87] Landow, G. Relationally encoded links and the rhetoric of hypertext. *Hypertext '87 Papers*. Proceedings of a conference held at the University of North Carolina, Chapel Hill. November 13-15, 1987, 331-343.
- [Linc88] Lincoln, J., and Boff, K. Making behavioral data useful for system design applications: Development of the *Engineering Data Compendium*. *Proceedings of the Human Factors Society 32nd Annual Meeting*, 1988.
- [Nels80] Nelson, T. Replacing the printed word: A complete literary system. *Proceedings of IFIP Congress 1980*, North-Holland, 1013-1023, 1980.
- [Oren87] Oren, T. The architecture of static hypertexts. *Hypertext '87 Papers*. Proceedings of a conference held at the University of North Carolina, Chapel Hill. November 13-15, 1987, 291-306.
- [Perl88] Perlman, G., and Moorhead, A. Applying hypertext methods for effective utilization of standards. *Proceedings of COMPSTAN'88 Conference on Computer Standards*, IEEE, 55-59, 1988.
- [Raym88] Raymond, D., and Tompa, F. Hypertext and the new Oxford English Dictionary. *Communications of the ACM*, 31(7), 871-879, July 1988.
- [Walk88a] Walker, J. Supporting document development with *Concordia*. *IEEE Computer*, 21(1), 48-59, January 1988.
- [Walk88b] Walker, J. The role of modularity in document authoring systems. *Proceedings of the ACM Conference on Document Processing Systems*, 117-124, 1988.

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.

© 1989 ACM 089791-339-6/89/0011/0060 \$1.50