Semantic Spaces:
A new Access Paradigm to Hypermedia Systems

Michael Klemme
Hypermedia Unit, Department of Computer Science
University of Auckland
Private Bag 92019, Auckland, New Zealand

September 2, 1996

Abstract

An increasing number of institutions and universities are using hypermedia systems for educational purposes. Unfortunately many of the systems available, with the WWW being a special source of concern, do not support advanced tools for navigation, study and collaboration.

In this paper I suggest a new tool, called Semantic Spaces *for this purpose and discuss its implementation into the Hyper-G hypermedia system. Following a general discussion on the potential of hypermedia systems for education, section 2 describes some of the the issues that arise of the use of hypermedia systems, especially spatial navigation (section 2.1) and organization of the information gathered while browsing. In section 3 the concept of semantic spaces is introduced as a medium for the organization of a users understanding of the contents of the hypermedia system. The implementation of these ideas using the hypermedia system Hyper-G is discussed in section 4.*

Keywords

Hypermedia, learning, spatial navigation, collaboration, Hyper-G, VRML

1 Hypermedia as a Learning and Working Environment

Hypermedia systems have recently attracted widespread interest. Driving factors for this development are to be found in the technical development, rather than in the development of genuine new applications. The proliferation of cheap networked computer systems and the hope for fast retrieval and efficient management of large amounts of data have helped hypermedia systems, and foremost the *World Wide Web*, to a user base of a size nobody dreamt of two or three years ago.

Unfortunately, most of the current development has been concentrated on the presentational side of hypertext and hypermedia. Although access to information is definitely one of the prerequisites for educational processes to take place, added educational value is expected from the "non-linear" structure of hypertext that enables multiple access to and a "multi-centric" view [20] of the material presented. Conklin

describes hypertext as a "computer-based medium for thinking and communication" [7]. The most promising aspects are:

- the integrative potential of hypermedia, allowing the integration of all materials used in the learning process,

- the availability of these materials for all forms and phases of learning, without needing to use different media in different situations.

- hypermedia's capability of being a repository for all learning situations, with the option of reusing hypermedia units in different situations and contexts.

This requires powerful tools allowing users to move through the hypermedia system, interact with it and other users, and customize it for their personal needs.

It is claimed that link structures in hypermedia systems map well onto the users' cognitive structures [16]. This hypothesis has not been proven however, and is questioned by many experts of the field.[1] It is based on the *syntactic* structure, i.e. the interrelationships defined on the document space by hyperlinks. This must not be confused with the *semantic* structure of the contents of hypermedia systems, i.e. the interrelationships of the concepts presented, since these are in most cases not the same [31, p. 63]. The insight into these structures is particular to a user and cannot be transferred to others. In particular non-experts need help finding out about the concepts covered in the system and are not interested in the *form* of presentation as represented in the link structure.

Learners new to a subject in particular have problems discovering the underlying structures and finding their way through the material. This problem is amplified in large interwoven systems like the World Wide Web [3] or Hyper-G [2], which span huge numbers of hosts and reference material of different contexts.

The understanding of a subject area can only be developed gradually. Students continuously extend their models as they find more information. These models are different from the relationships explicated by nodes and links in the hypermedia system, the students model can be completely wrong and will (hopefully) be corrected later.

Moving through the document space is normally associated with the metaphor of *navigation*, in recourse to travel and orientation in the real world. Unfortunately, it is not clear if the same skills used in real-life navigation can be applied to reading and using electronic information systems [22, pp. 65ff]. The efficient use of hypermedia systems has to be learned by users.

Several tools have been developed to help the user work with hypermedia systems. Some of the approaches are described in section 2.1. This paper later presents a tool for the spatial representation of the users' understanding of the semantic context and the manipulation of these visualisation as the users discover additional aspects of the topic.

---

[1] See [11], [27, p. 245], and [22, pp. 96ff].

2 Issues for the Use of Hypermedia Systems in Education

The following sections lists some of the issues that are important for the use of hypermedia systems in educational contexts but which are not sufficiently considered in most implementations. This list is not meant to be complete, but is a list of topics that arose when using hypermedia systems for teaching support at Universität–Gesamthochschule Paderborn, Germany.

2.1 Navigational Aids

Navigation in large hypermedia systems is considered a major problem in the usability of such systems.[2] The structure of nodes and links can be seen as a complex graph or a document space that—due to its irregular structure—cannot be visualized effectively. Utting and Yankelovich [30] described and tested different tools for solving this problem such as local and global maps. The proposed tools proved either to be insufficient in scope or too complex to be understood and efficiently handled by the users. Although most of the navigational features described below are not new, many systems like the WWW and its clients come with little or no navigational support. Most of these navigational aids display the *syntactical* structure, and are not sufficient for the understanding of the content: Users are trying to discover the *semantic* relationship of the information.

Global Maps These maps create a graphical representation of all documents of a system, or of a designated subset of the documents, and the links interconnecting them. They will display at least the titles of documents stored in the database, but may also contain additional information such as size and media type. These maps can easily be generated automatically and can give users an idea about the context of the documents visited.

However, global maps tend to be overcrowded. In particular, strongly interconnected document systems cannot be visualized without links intersecting other links or nodes. Displaying all links may obscure the high-level structure of the hyperspace. This problem cannot be solved easily. A classification of high and low-level links does not seem feasible, since the relevance of links might differ from different points of view or on different paths into the hypermedia system. Links to and from parts of documents cannot be displayed in a meaningful way.

Automatically generated maps do not retain their layout over time, they change as soon as the document base changes, to maintain consistency with the underlying database. Therefore, these maps cannot be used for long-term navigation through the database. Formatting global maps, it is computationally hard to visualize the "neighbourhood", i.e. the directly connected documents, in a meaningful way, due to the arbitrary interconnections of hypermedia documents. Again, relying on the physical link structure might mislead the user.

Local Maps, Fisheyes Local Maps reduce the complexity by only displaying the documents adjacent to the current or selected document. An overview of the structure

---

[2]Some authors maintain however, that this is more a problem of system design than hypermedia systems in general [27, p. 259].

3

of the hypermedia system is not possible. It is even difficult to judge the role of the centre document, since only links to and from this document are displayed. Again, local maps change as soon as the underlying system changes.

Fisheye views try to eliminate this disadvantage by additionally displaying "landmark documents". These are selected according to a relevance metric [13]. Unfortunately, those metrics are not easy to determine in an arbitrary graph.[3] Graphical fisheye views visualize the nodes and links with reduced detail if they are further away from a selected document [25]. This technique uses a fixed two-dimensional layout and is most usefull for the visualisation of planar graphs.

Overview Maps   Hand-made overview maps can be created by the author of the hypermedia systems as well as by the users themselves. They are usually produced using a standard drawing application and can be inter-linked with the referenced hypermedia documents. Overview maps were heavily used in *Intermedia* [32, 18]. These maps necessarily reflect their authors' view of the world and cannot normally be adapted to the users' personal needs. Hand-made maps can be based on a metaphor, representing a physical environment or a development in time. They also have to be manually adapted to changes in the underlying data as these happen.

As mentioned above, one of the educational advantages of hypermedia systems is the existence of multiple access paths to the material. It is therefore difficult to provide overview maps for all paths and all motivations.

History Functions   History functions record the documents visited by the users so far or the path of the users from their start document to the current document. The history allows users to evaluate their progress and to backtrack on their path if necessary to restart the search from a previous document. Many systems mark all visited documents and the links leading there. This applies to the links to visited documents as well as to the representation of documents in navigational overviews. The user can thus identify documents visited before re-entering them on a different path.

2.2   Spatial Navigation and Organization of Information

Studies have shown that people often recall the physical position of a piece of information within a text or a book [22, pp. 73ff] ("on the upper left side, in the first half of the boo"). They can use many physical clues available within a text to understand the structure and retrieve information, e.g. page numbers, pagination, indices, etc. Many of these clues are not available in electronic systems, others, like indices and tables of content, are available via the system software and not physically accessible .

The human skills of memorising and orienting in two- or three-dimensional worlds can thus not be utilized for understanding the hypermedia system. Research in human-computer interaction has shown the advantages of a spatial graphical interface that allows direct manipulation of objects [19, 28]. The introduction of the desktop metaphor into personal computing, with its associated tools and techniques, has significantly increased the usability of computer systems, especially for a more general audience. The

---

[3]Most of the examples given in the literature (e.g. [26]) apply to hierarchically organized hypermedia systems.

transfer and extension of this concept to hypermedia systems might help to mitigate the problems described above.

Location Feedback.   Tools for spatial navigation can be augmented and tightly integrated into hypermedia systems by visualizing the users' current position within the hypermedia system using these multi-dimensional maps as they move through the system [23, p. 44]. This helps the users understanding how the documents relate to each other and to the overall system. All changes of status are then reflected in one unified navigational system and can be evaluated in their overall context.

Other spatial Information Systems.   Several commercial programs exist to help users organize their World Wide Web and Internet references. One of them is *Web Squirrel* [4], which allows the user to arrange Internet pointers on a two-dimensional sheet and to define groups of entries. Filter operations can be performed automatically on these groups and other local resources. The authors refer to this as "information farming".

Dieberger, Pohl, and Purgathofer have developed a graphical interactive hypertext editor [10]. This tool provides a graphical overview of all documents and interconnection links. These documents can be placed arbitrarily on the screen to foster the organization of the writing process. The authors findings were that their students had problems formulating non-hierarchic relationships between their concepts.

Several spatial metaphors for hypermedia systems were proposed [8, 9]. They organize the information in terms of rooms and buildings that can be visited by the user. These metaphors may be useful in limited contexts, but it has to be evaluated case-by-case how far these metaphors support the information seeking task.

The value of spatial layout programs for the organization and correlation of knowledge was recognized for general learning strategies [15]. Fischer et.al. describe a computer-supported mapping system for the organization of facts and hypotheses into a medical diagnosis [12]. It is used for training medical student in collecting and organizing many facts and evaluating multiple possible solutions.

2.3   Personal Customization and Extension

Most of the hypermedia systems available so far lack tools for the customization to personal preferences and needs. Often there are no facilities to add personal annotations, record the significance of a document, or record discovered relationships by addition of personal links. This is perfectly sufficient for presentation systems (e.g. the WWW) but renders them nearly useless for any more intensive and engaged use as would be expected in learning environments. Browsing a hypermedia system might not be sufficient for more intensive study [31], active engagement is desirable in many situations [29].

Users are producers of knowledge as well. New information is gathered while traversing the hypermedia system. This information needs to be recorded and arranged for later use. Their knowledge relates to the information content of the hypermedia system and should therefore be integrated into the system and its navigational overviews. Some systems permit the addition of personal annotations and links to documents, Intermedia being again one of the most advanced systems [6], but this is still less

than paper-based materials allow, e.g. marking, marginal notes, transcription, etc. The added information tends to be spread out through the database and needs to be organized to be accessible.

As stated in section 1 hypermedia systems have the potential to integrate the materials needed for learning. In order to maintain consistency, the students' notes and products have to be included into the system as well.

2.4 Collaboration using Hypermedia Systems

The educational value of group interaction and collaboration has been discussed at length [27, p. 261]. Nevertheless, most hypermedia systems lack support for group activities. Kent Norman reports that special tools are needed for the communication of an actual position and knowledge of the structure of the material used [23]. This problem is significant when using a hypermedia system as a teaching support system for lectures and tutorials. "Pointing out" your current position or sharing your knowledge with other people is impossible without shared knowledge of the structure of the system used.

2.5 Filtering

Filter mechanisms can help the users reduce the complexity of navigational aids and allows for a more effective exploration, and thus understanding, of the document space, as already recognized in [14]. Dynamic query filters [1] allow users to adjust filter parameters dynamically. Immediate feedback helps refine the query parameters and allows visual scanning of the results. The *Film Finder* [1] might be considered as a proof of concept. However, the definition of filtering rules should not be too complex or difficult if all users are supposed to use them. Possible search parameters have to be identified beforehand and appropriate values added to all documents. Most multipurpose hypermedia systems have few attributes for their documents or do not provide efficient access or manipulation to them. Full text searches on text documents might alleviate this situation. Pattern recognition mechanisms might provide clues about non-textual documents, e.g. video [33], but more research in this field is needed.

3 Semantic Spaces

Use of *Semantic Spaces* may help to overcome some of the limitations described above. A Semantic Space is a tool for the exploration of hypermedia systems. Users can explicitly describe their personal understanding of the part of the system they have explored themselves. They start with an empty map or with a map that was prepared by a knowledgeable person. Then, as they move through the hypermedia system they can take notes on what they think is significant. They can arrange these pieces of information on two- or three-dimensional work sheets and group information as it might seem to be useful. Information may be arranged "around" identified key or landmark ideas (see Figure 2) or might be sorted according to metrics that the users feel to be meaningful in this context. Possible dimensions are time, space or different categories (see Figure 1). The distance between objects may represent the degree/strength of the relationship in between.

Figure 1 displays a example of a semantic space ordered according to some parameters. The hypermedia system contains a range of scientific articles on the relationship of computer development to the military. The semantic space is laid out according to the time and different categories [5]. Figure 2 displays a prototype semantic space created for an educational hypertext on the rivers of New Zealand.[4] Concepts are grouped according to their relationships in the database.

Users record the interrelationships between the concepts in the semantic space and the documents of the 'real' hypermedia system by adding specially labelled hyperlinks between concept and document. With the help of these links the position of the users within the hypermedia system can be visualised on the current semantic space when the user moves through the hypermedia system.

Links can be created interactively by naming the concept and target document or by dragging a document icon into the semantic space, creating the links from the semantic space to the documents automatically.

The term "Semantic Space" was chosen to be a contrast to the term "Hyperspace".[5] Hyperspace describes the multi-dimensional information space created by documents and links (and thus only covers the syntactical relationships of nodes and links).

By creating maps of their own, users can develop and visualize a spatial understanding of the hypermedia system. This intensifies the users' involvement and might thus foster their understanding of the overall context. The layout remains persistent until the user decides to change it because it no longer represents their understanding of the document space.

These concepts are not necessarily related to the nodes and links in the hypertext, but represent the user's current knowledge and understanding of the system. They will be altered, moved or deleted as the process of understanding continues. A set of tools, similar to those of standard desktop environments, for grouping, moving, deleting, etc. provide for rapid manipulation. Additional drawing and writing tools are needed for annotation and graphical markup (arrows, boxes, etc.). New "raw" semantic spaces can be effectively generated from the hypermedia database using dynamic query filters [1].

If users move within the hypermedia system beyond the scope of their current map they will want to extend the map to cover the new information found. This involves adding new information to the map as well as re-arrangement of the existing items.

Multiple semantic spaces can be produced per user and database/collection hierarchy. It remains to be shown how to switch between semantic spaces when the user transfers from one to the other. The user may pick a semantic space before starting the exploration and switch explicitly to a different environment.

Semantic spaces can be used for the design of new hypermedia contents, as a worksheet for testing the structure and organization of the materials to be put into the database.

Semantic spaces are also documents within the hypermedia system. They can be inter-linked with the other documents or be referenced by other semantic spaces, thus creating a mesh or, if required, hierarchy of semantic spaces. They are persistent over time and can be shared with other users.

---

[4] To be found at http://www.hmu.auckland.ac.nz/seakeepers/

[5] The term is used in the area of Artificial Intelligence and Natural Language Processing as well
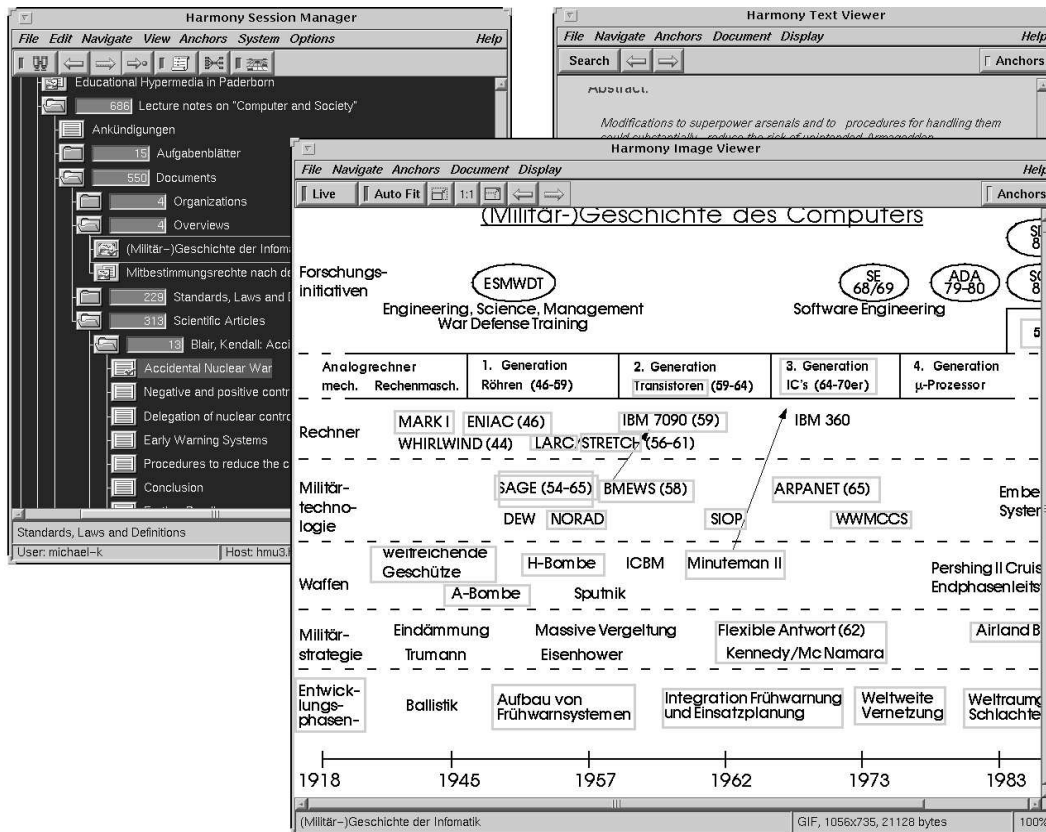
Figure 1: Mock-up of a semantic space. Concepts are ordered according to time and specified categories. The light grey rectangles connect this map to the documents in the Hyper-G server (as seen in the background).

Major changes to the hypermedia system (e.g. deletion, addition or substantial changes of documents) cannot be automatically reflected in the semantic spaces, since deleting information might render the map useless to the user. Rather I propose deleting the semantic link from the map to the document being deleted[6] and marking the link as deleted in the semantic space, e.g. by greying out its icon.

Semantic spaces solve some of the problems described previously. They integrate overview maps and make them extendable and adaptable by the user. The users can customize the view by grouping and arranging the information in a meaningful way and thus personalize the hypermedia system. These products show the users' view of the database and can thus be shared and discussed. Collaborative design of semantic spaces might foster understanding and learning in groups. Semantic spaces can be used as a visualization of a group's knowledge about the hypermedia system and as a starting point for further exploration.

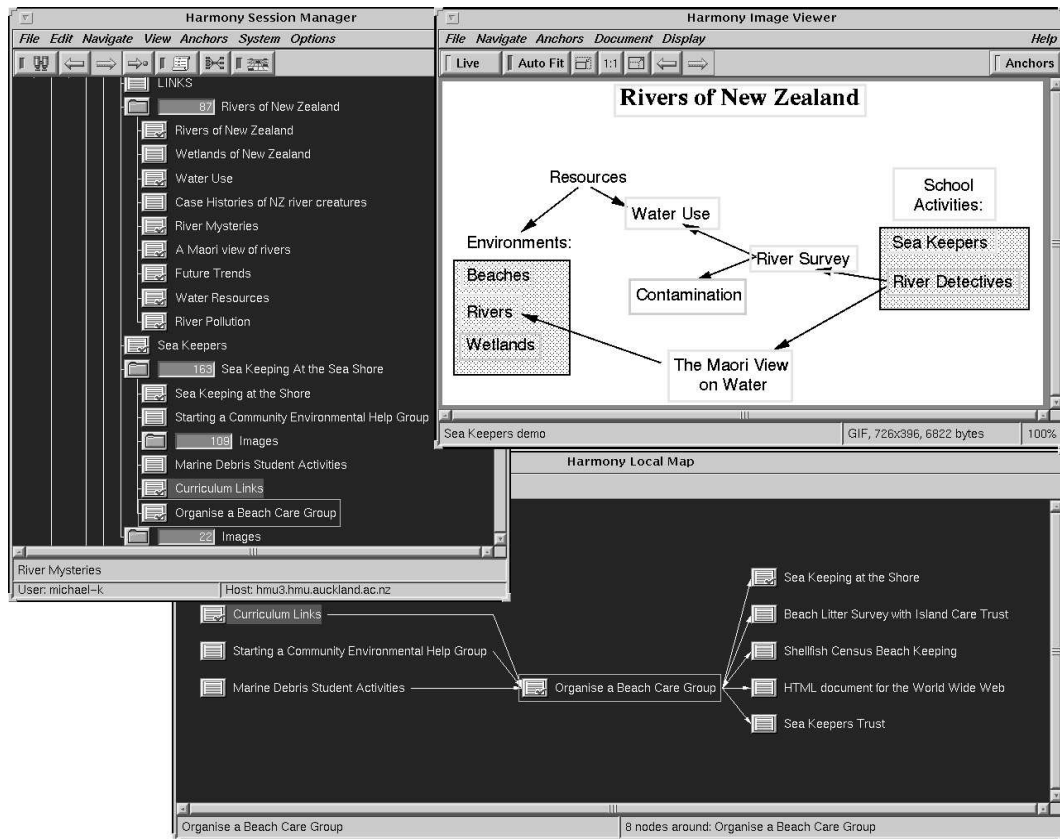[6] This operation is particularly trivial in Hyper-G.

Figure 2: Mock-up of a semantic space: Concepts are grouped according to their relationships in the database.

4 Proposed Implementation into Hyper-G

Hyper-G[7] is an advanced networked hypermedia system that offers many features that can be employed for an implementation of the concepts described above [2, 21]. The most distinctive feature compared to the WWW is the separate link database. This allows for bi-directional links and links to be defined on objects in all document formats. The system automatically keeps track of all documents and links and disables links whose destination documents are unavailable.

Furthermore, Hyper-G offers transparent access to and from the *World Wide Web* and thus should at least allow read access to semantic spaces from most computer platforms used today. Currently Hyper-G provides two editing clients "Harmony" for X-Windows and "Amadeus" for MS-Windows, which provide viewers for a wide range of media. Harmony will be the primary target for implementation.

A data format needs to be selected for this project. Due to the limited resources available, a custom format cannot be developed. Instead, available viewers will be modified for our purpose. Implementing filter mechanisms might jeopardize the use of

---

[7] Hyper-G is now marketed under the name *HyperWave*.

9

standard data formats, since they involve reducing the information content at presentation time.

Clients other than native Hyper-G clients should be able to display the spaces, even if the WWW clients cannot modify them.

The *Virtual Reality Modelling Language (VRML)* is a format for three-dimensional scenes developed by the World Wide Web community.[8] It is loosely based on SGI's GL and Open GL. A viewer for Hyper-G has been implemented and tools for the interactive modification of VRML are available. VRML defines links in and out of 3D scenes and to other WWW resources. WWW-Viewers for VRML exist for some platforms and plug-in modules for the popular Netscape Navigator allow VRML to be viewed from inside this program. Using these extensions, links from the semantic spaces could be seen from the WWW as well. Because VRML requires powerful machines for display, semantic spaces will initially be based on a two-dimensional subset of VRML (e.g. fixed viewing parameters).

Location feedback will be implemented as a special type of link, pointing from the space to documents. Semantic links can thus be defined by all registered users of the system. Due to Hyper-G's bi-directional links the system can follow the links backwards from the document to the specified region of the map and visualize the users position. The session control of Harmony must be modified to inform the space viewer of all changes of location.

Drag-and-Drop support for Harmony will be implemented for moving and copying of object within the collection hierarchy. Once implemented it can then be used to transfer object information from the collection hierarchy into the semantic spaces.

Additional editing functions for grouping, moving information, and inserting text, lines, etc., will also be needed.

5   Conclusions and Future Work

Semantic spaces have the potential to solve, or at least ease, some of the problems using hypermedia which have been identified in this paper. They allow for long-term spatial navigation and give direct feedback about concepts and structures of the hypermedia system as it is currently conceived by the user. They can be tailored by the users according to their personal needs and are therefore are a tool for the systematic and intensive exploration of hyperspaces.

The proposal will be be implemented into the Hyper-G hypermedia system in the very near future, and an extensive evaluation as part of a university course is planned.

---

[8] See http://www.w3.org/pub/WWW/MarkUp/VRML/ for further information.

References

[1] C. Ahlberg and B. Shneiderman. Visual Information Seeking: Tight Coupling of Dynamic Query Filters with Starfield Displays. In *Human Factors in Computing Systems, Proceddings of CHI 94*, pages 313–317. Association for Computing Machinery, April 1994.

[2] K. Andrews, F. Kappe, and H. Maurer. The Hyper-G Network Information System. *Journal of Universal Computer Science*, 1(4), April 1995.

[3] T. Berners-Lee, R. Cailliau, A. Luotonen, H. Frysyk, and A. Secret. The World-Wide Web. *Communications of the ACM*, 37(8):76–82, August 1994.

[4] M. Bernstein. The Eastgate Web Squirrel. *ACM SIGLINK Newsletter*, 5(1):6, February 1996.

[5] A. Brennecke and R. Keil-Slawik. Notes on the Alltagspraxis of Hypermedia Design. In H. Maurer, editor, *Educational Multimedia and Hypermedia Annual (Proc. of ED-MEDIA 95, Graz)*, pages 115 — 120. Association for the Advancement of Computing in Education, June 1995.

[6] T. Catlin, P. Bush, and N. Yankelovich. InterNote: Extending a Hypermedia Framework to support Annotative Collaboration. In *Proceedings ACM Hypertext 89*, pages 365–379, New York, 1989. Association for Computing Machinery.

[7] J. Conklin. Hypertext: An Introduction and Survey. *IEEE Computer*, pages 17–41, September 1987.

[8] G. Davies, H. Maurer, and J. Preece. Presentation Metaphors for very large Hypermedia Systems. *Journal of Microcomputer Applications*, 14:105–116, 1991.

[9] A. Dieberger. Spatial User Interface Metaphors in Hypermedia Systems, Workshop at the European Conference on Hypermedia Technology, September 1994, Edinburgh, Scotland. WWW: http://www.lcc.gatech.edu/faculty/dieberger/ECHT94.WS.toc.html, 1994.

[10] A. Dieberger, M. Pohl, and P. Purgathofer. Einige Fragen zur Entwicklung von Hypertext–Autorenwerkzeugen. In E. Schoop, R. Witt, and U. Glowalla, editors, *Hypermedia in der Aus- und Weiterbildung*, number 17 in Schriften zur Informationswissenschaft, Dresden, October 1994. Universitätsverlag Konstanz.

[11] T. M. Duffy and R. A. Knuth. Hypermedia and instruction: Where is the Match. In Jonassen and Mandl [17], pages 199–225.

[12] F. Fischer, C. Gräsel, A. Kittel, and H. Mandel. Ein computerbasiertes Mapping-Verfahren zur Strukturierung komplexer Informationen: Entwicklung und Evaluation. Forschungsbericht 57, Institut für Pädagogische Psychologie und Empirische Pädagogik, LMU München, October 1995.

[13] G. W. Furnas. Generalized Fisheye Views. In M. Mantei and P. Obeton, editors, *ACM Conference on Human Factors in Computing Systems III (CHI 86)*, pages 16–23, April 1986.

[14] F. G. Halasz. Reflections on Notecards: Seven Issues for the Next Generation of Hypermedia Systems. *Communications of the ACM*, 31(7):836–852, 1988.

[15] D. Jonassen, K. Beissner, and M. Yacci. *Structural Knowledge Techniques for Representing, Conveying and Acquiring Structural Knowledge*. Lawrence Erlbaum Ass, Hillsdale, NJ, 1993.

[16] D. H. Jonassen. Hypertext Principles for Text and Courseware Design. *Educational Psychologist*, 21(4):269–292, 1986.

[17] D. H. Jonassen and H. Mandl, editors. *Designing Hypermedia for Learning*, volume 67 of *NATO ASI Series F*, Berlin, Heidelberg, New York, 1990. Springer.

[18] P. Kahn. Linking together Books: Experiments in adapting published Material into Intermedia Documents. *Hypertext*, 1(2):111–145, 1989.

[19] R. Keil-Slawik. Gestaltung interaktiver Systeme — Ein ökologischer Ansatz. *Log In*, 12(5/6):18–27, 1992.

[20] G. P. Landow. Hypertext in Literacy Education, Criticism and Scholarship. *Computer and the Humanities*, 23:173–198, 1989.

[21] H. Maurer, editor. *HyperWave: The Next Generation Web Solution*. Addison Wesley UK, 1996.

[22] C. McKnight, A. Dillon, and J. Richardson. *Hypertext in Context*. The Cambridge Series on Electronic Publishing. Cambridge University Press, 1991.

[23] K. L. Norman. Navigating the Educational Space with Hypercourseware. *Hypermedia*, 6(1):35–60, 1994.

[24] G. G. Robertson, S. K. Card, and J. D. Mackinlay. Information Visualization using 3D Interactive Animation. *Communications of the ACM*, 36(4):57–71, April 1993.

[25] M. Sarkar and M. H. Brown. Graphical Fisheye Views of Graphs. In *Human Factors in Computing Systems, Proceddings of CHI 92*, pages 83–91. Association for Computing Machinery, May 1992.

[26] K.-H. Saxer and P. A. Gloor. Navigation im Hyperraum: Fisheye Views in HyperCard. In P. A. Gloor and N. A. Streitz, editors, *Hypertext und Hypermedia — Von theoretischen Konzepten zur praktischen Anwendung*, number 249 in Informatik–Fachberichte, pages 190–204. GI, Springer, 1990.

[27] R. Schulmeister. *Grundlagen hypermedialer Lernsysteme: Theorie - Didaktik - Design*. Addison Wesley, Bonn, 1996.

[28] B. Shneiderman. *Designing the User Interface: Strategies for Effective User-Computer Interaction*. Addison–Wesley, 2nd edition, 1992.

[29] B. Shneiderman. Education by Engagement and Construction: A Strategic Education Initiative for a Multimedia Renewal of American Education. In E. Barret, editor, *Sociomedia: Hypermedia, Multimedia and the Social Creation of Knowledge*, pages 13–26. MIT Press, Cambridge, MA, paperback edition 1994 edition, 1992.

[30] K. Utting and N. Yankelovich. Context and Orientation in Hypermedia Networks. *ACM Transactions on Information Systems*, 7(1):58–84, January 1989.

[31] P. Whalley. Models of Hypertext Structure and Learning. In Jonassen and Mandl [17], pages 61–67.

[32] N. Yankelovich, G. P. Landow, and D. Cody. Creating Hypermedia Materials for English Literature Students. *SIGCUE Outlook*, 19(2–4):12–25, 1987. ACM Special Interest Group on Computer Use in Education.

[33] H. L. Zhang, C. Y. Low, W. Smolinar, and J. H. Wu. Video Parsing, Retrieval and Browsing: An Integrated and Content-based Solution. In ACM, editor, *Proc. ACM Multimedia*, pages 15–24, New York, November 1995. ACM-Press.