Augmenting User Interfaces for Digital Libraries with Virtual Reality

Chaomei Chen
Department of Information Systems and Computing
Brunel University
Uxbridge UB8 3PH, UK
Chaomei.Chen@brunel.ac.uk

Abstract

This paper describes a generic approach to the development of a virtual reality-based user interface for a collection of digital documents. We emphasise the role of intrinsic inter-document relationships in structuring and visualising a digital library. We integrate techniques from information retrieval and structure modelling for visualising the ACM Hypertext Compendium. Salient relationships in the Compendium are extracted and visualised as associative networks. A virtual reality-based user interface is automatically generated in VRML. We show that the spatial layout of the virtual world highlights prominent structural patterns in the Compendium. We discuss design issues concerning seamlessly browsing the digital library in a 3D virtual world and how users can search and browse within a consistent spatial metaphor provided by the user interface.

1. Introduction

The widespread use of the World-Wide Web (WWW) has highlighted the demand for cost-effective ways of accessing information in large distributed hypermedia environments, such as digital libraries, electronic publishing and subject-specific knowledge repositories. It is a challenging task for users to navigate through a large number of documents and hyperlinks without an idea of how these documents are interrelated together. Users need to access a large information space in an intuitive and effective way.

Previous research in hypertext systems has recognised the significant role of graphical overviews and typed links in dealing with a complex network. In late 1980s, researchers started to incorporate the notion of hypertext into traditional information retrieval systems [7, 8]. More recently, the role of hyperlinks in improving the quality of information retrieval has been increasingly recognised [19].

The highly dynamic and distributed nature of the WWW highlights some fundamental issues concerning the development of large, distributed hypermedia systems. For example, traditional hypermedia systems are often statically constructed. A greater degree of flexibility and maintainability can be achieved by using dynamic node-link binding in which the structure of a hypermedia network is dynamically determined by distinct sets of link configurations. In fact, there is a rapidly growing interest in open hypermedia services [2, 11]. A fundamental issue is whether large link configurations can be efficiently generated to represent intrinsic characteristics of the underlying document space.

The ACM Hypertext Compendium [1] is a collection of papers from a series of international conferences on hypertext. In this paper, we introduce a generic approach to the development of a cohesive user interface for the ACM Hypertext Compendium. This paper is organised as follows. First, we introduce Generalised Similarity Analysis (GSA), especially the role of the vector space model and Pathfinder networks in GSA. Second, we describe the development of a virtual reality-based user interface for the Compendium. We expect that these user interfaces would allow users navigate seamlessly in a complex network of digital documents. Four, we examine actual documents associated with some predominant structural patterns in the Compendium network. Finally, we discuss the implications of this approach for designing user interfaces to digital libraries.

2. Organising Digital Libraries

In traditional libraries and bookshops, books are often arranged with a layout that would place books on similar topics near to each other; therefore people may find convenient to browse the same local area for books on related topics. In the development of digital libraries, it is desirable to have a similar spatial metaphor such that

1060-3425/98 $10.00 (C) 1998 IEEE

users can rely on the skills that they have developed in the real world.

When dealing with a large information space, designers traditionally face a fundamental challenge — how to present users with sufficiently local detail as well as maintain a meaningful context so that users can move through the information space efficiently. Information visualisation techniques, such as Fisheye Views [10], Cone Trees [17] and Perspective Walls [13], provide some solutions to the problem of balancing local detail and global context. Many of these techniques are based on explicit attributes of a document or a set of documents, such as file size, file names and file system structure.

Designers of interactive systems for digital documents need to address many fundamental issues. For example, what types of interactive systems are needed for users to access the vast amount of information on the Internet seamlessly? How can one take users' browsing patterns into account and incorporate the knowledge into the design of a new generation of user interfaces? What is the role of virtual reality in the design of new generations of interactive systems?

Automated link generation and virtual reality can be integrated together to enable users seamlessly navigate in an information space in association with a digital library. It is possible to provide a generic methodology for dynamically generating a cohesive view of a set of digital documents, which may well be scattered across several geographically distributed websites, or extended over a long period of time.

2.1 Framework

The central concept in our work is an integrated, iterative design framework for developing a dynamic, adaptable library of interconnected digital documents. It unifies information structuring, visualisation, and user modelling techniques (see Figure 1).

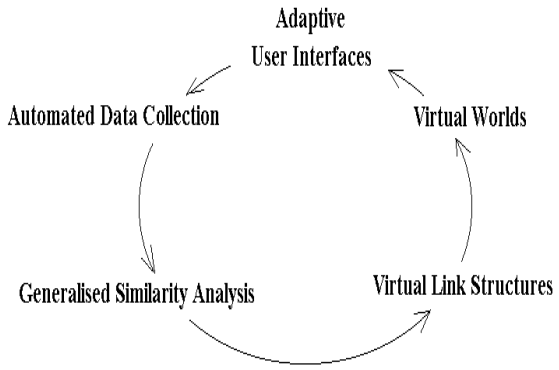

Figure 1. An iterative design process.

The framework not only provides mechanisms to model how users browse the information space, but also allows structuring and adapting the information space to match the interests of users. The integrated modelling and design process can be used iteratively.

The overall process consists of the following steps:

1. **Data collection**, gathering information such as the title, authors' names, the abstract and full-text of a digital document; if digital documents are stored in distributed sources, Internet traversal and retrieval agents, commonly known as spiders or crawlers, are typically used.
2. **Building similarity models based on content similarity**, **hypertext linkage and usage patterns**. In this paper, we concentrate on the content-based similarity model. Examples of other types of similarity models can be found in [4].
3. **Pathfinder network scaling**, extracting salient inter-document relationships and generating spatial configurations.
4. **Generating virtual reality models**, using spatial configurations as a blueprint to produce virtual worlds in VRML.
5. **Transforming and optimising the present virtual reality-based user interface**, adjusting existing models and spatial configurations based on usage patterns.

In the following sections, we introduce the basic concepts and procedures used by Generalised Similarity Analysis.

2.2 Generalised Similarity Analysis

Generalised Similarity Analysis (GSA) is a unifying framework for extracting structural patterns from a hypermedia information space [4]. A number of intrinsic interrelationships in hypertext, such as hypertext linkage, content similarity and browsing patterns, are consistently incorporated into the generic framework. In this paper, we apply GSA to the ACM Hypertext Compendium.

The architecture of the GSA framework consists of a number of computational models (see Figure 2). Each of these computational models generates a virtual link structure based on a distinct characteristic. A virtual link structure can be generated based on an integration of some or all the component models. One can incorporate additional inter-document relationships into the framework, for example, based on citation and co-citation counts between documents.

The similarity between two documents can be measured psychologically or statistically. In hypermedia systems, some fundamental relationships are hypertext linkage, content similarity and browsing patterns. In this

paper, we concentrate on similarity models based on terms, or words, used in a collection of documents.

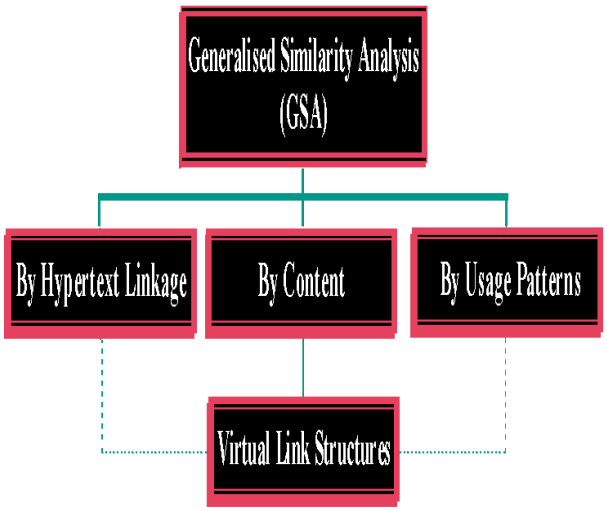

Figure 2. Generalised similarity analysis.

The vector-space model has been widely used in information retrieval [18]. In this model, each document is represented by a vector of terms. Terms are weighted to indicate how important they are in representing the document. The distance between two documents is defined as the distance between the two corresponding vectors. Salton *et al.* automatically generated semantically-based hypertext networks using the vector-space model [18].

In this study, we use the well-known $tf \times idf$ model, term frequency times inverse document frequency, to build term vectors as in [18]. Each document is represented by a vector of T terms. The weight of term T_k in the vector of document D_i, is determined by

$$w_{ik} = \frac{tf_{ik} \times log\left(\frac{N}{n_k}\right)}{\sqrt{\sum_{j=1}^{T}\left(tf_{ij}\right)^2 \times log\left(\frac{N}{n_j}\right)^2}}$$

where tf_{ik} is the occurrences of term T_k in D_i, N is the number of documents in the collection (such as the size of a WWW site), and n_k represents the number of documents containing term T_k. The document similarity is computed as follows based on corresponding vectors $D_i = (w_{i1}, w_{i2}, ..., w_{iT})$ and $D_j = (w_{j1}, w_{j2}, ..., w_{jT})$:

$$sim_{ij}^{content} = \sum_{k=1}^{T} w_{ik} \times w_{jk}$$

The HyPursuit system [20] also used a modified version of the vector-space model. However, the HyPursuit system focuses on generating hierarchical classifications of documents, whereas our purpose is to generate a network representation of the structure of digital documents. Pirolli *et al.* [16] used the vector-space model in a study of a Xerox's WWW server, but content similarities were only computed between documents connected by hyperlinks. In our study, content similarities are considered across the entire collection of documents in order to find out under-represented patterns.

In GSA, the spatial layout is determined by a spatial configuration, which maps similarities from the domain N×N to $\Re^2$ or $\Re^3$ as with multidimensional scaling (MDS). The difference is that a Pathfinder network is built on salient relationships determined by the triangular inequality condition, which is imposed to eliminate invalid links, whereas an MDS solution is derived from N×N measures of pair-wise relationships.

The graph layout of a Pathfinder network is determined by a force-directed graph drawing algorithm. Because the idea is simple and intuitive, force-directed graph drawing becomes increasingly popular in information visualisation [3]. In a spring-energy model, nodes are connected by springs corresponding to weighted links. These nodes are forced into place by spring energy transformed from the weights. As the overall spring energy in the system is minimised, the graph gradually takes shape. Resolving spring models usually requires the computational complexity of $O(N^2)$. As the number of nodes in the graph grows, more efficient solutions are necessary. One possible solution is to use the divide-and-conquer strategy: a large information space can be rapidly split into smaller clusters with simple classification algorithms until those computationally expensive algorithms can be effectively used.

2.3 Source Data

We use the ACM Hypertext Compendium™ [1] to build an experimental digital library, containing about 130 full-text papers from international conferences on hypertext. We describe the *Memex and Beyond* project [15] as a closely related effort towards a comprehensive digital library on hypertext. The long-term goal of the *Memex and Beyond* project is to provide a highly interconnected set of services which will allow users to access bibliographical information efficiently. It also proposes to support graphical and spatial hypertext user interface, although currently the textual access is the only available means of access.

The *Memex and Beyond* web site is a major research, educational, and collaborative web site integrating the historical record of and current research in hypermedia [15]. *Memex and Beyond* is an outreach website of the NSF Graphics and Visualization Center, an National

Science Foundation (NSF) Science and Technology Center [15].

The site aims to provide graphical, spatial, and textual representations of the relationships among the people, projects, institutions, publications, conferences, and themes that comprise the hypermedia community. The global index now contains bibliographic entries for the following conference proceedings:

- Hypertext'87
- Hypertext'89
- ECHT'90
- Hypertext'91
- ECHT'92
- Hypertext'93
- ECHT'94
- IWHD'95 (International Workshop on Hypermedia Design)
- Hypertext'96

The web site is to provide a rich mixture of content and navigation mechanisms. Three types of navigation mechanisms are planned: graphical, spatial, and textual. However, at the time of this study, the textual navigation mechanism is the only mechanism in place. The following description is from the *Memex and Beyond* site, which shows a clear interest in the direction that our approach may play a significant role:

The spatial navigators will be Java-activated imagemap versions of Eastgate Systems' Web Squirrel farms. These farms are 2D representations of a domain in which the elements can be annotated Web resources, such as URLs and email addresses, or text. The advantage of this representation over a list or outline is that logical relationships among elements can be graphically shown by their spatial relationships.

The ACM *Hypertext Compendium*™ (HTC) is a hypertext database containing literature from the field of hypertext [1]. HTC intends to collect a significant proportion of the literature of the field in a single hypertext database. The Compendium includes:

- Hypertext'87 full-text articles
- Hypertext'89 full-text articles
- ECHT'90 full-text articles
- About 30 articles from other sources related to hypertext, including CHI86-CHI90 and CSCW88 conference proceedings and *Communications of the ACM*.

The Compendium contains two versions: the KMS version and the text version. The KMS version contains 133 files requiring 12.5Mb disk space, including graphics and rich cross-reference links among these articles. The text version contains 128 files, taking 3.7Mb disk space.

In this paper, we focus on the ACM *Hypertext Compendium™*. The principal reason is that the KMS version of the Compendium contains a rich set of cross-reference links among these articles. Although the work described in this paper is based on the text version, the existing links in the KMS version will be explored in our next prototype, in which we plan to utilise the hypertext linkage model of GSA together with the content similarity model used in this paper.

3. Visualisation

The step-by-step application of GSA to the ACM Hypertext Compendium is described as follows.

3.1 Data Collection

The ACM Hypertext Compendium was obtained directly from the ACM on diskettes. A total of 126 documents from the text versions of the ACM Compendium were used. Two missing files were due to a corrupted diskette. It is intended that the next prototype will be based on the KMS version of the Compendium and make use of the existing cross-references and other types of links. The 126 documents were pre-processed and converted to a standard format for GSA.

3.2 Inter-Document Similarity Matrix

The inter-document similarity matrix was computed based on the vector-space model [18]. The vector-space model was based on the title, authors' names and the abstract of each document. If there is no abstract, the first 25 lines of the main text were used in the computation. A list of common English words was used as the list of stopwords. The similarity matrix was generated according to the dot product of corresponding vectors.

If the value of a pair-wise similarity was less than the threshold of 0.20, it was excluded from further analysis in order to reduce the level of noise. The resultant similarity matrix was submitted for Pathfinder network scaling. For detailed descriptions about Pathfinder networks, see [14, 6, 4]. The result of Pathfinder network scaling is a spatial configuration of the minimum cost network, including the co-ordinates of documents and links that satisfy the triangular inequality condition throughout the network.

3.3 Virtual Worlds Design

We develop a virtual reality-based user interface in Virtual Reality Modelling Language (VRML). The rationale is to represent the underlying semantic structure as a constellation of documents in a virtual world such that users can explore the structure of a complex network efficiently and intuitively. Major issues and experience

associated with our design rationale are discussed in subsequent sections.

In our user interface, each document is visualised as a sphere. The link between two connecting documents is visualised by a cylinder connecting the two spheres. The length of the cylinder is the semantic distance between the two documents connected by the cylinder. The diameter of the cylinder is proportional to the degree of similarity between the two documents. In this paper, spheres and cylinders in a virtual world are randomly coloured on smooth surfaces.

Table 1. Visualisation model.

Digital Objects	Geometric Model	Attribute	Semantics
document	sphere	radius	size
document	sphere	colour	creation time
link	cylinder	radius	similarity
link	cylinder	length	distance
query	cube	height	weight
query	cube	colour	coding
proximity	level of detail	range	focus

We have developed a series of computer programs to transform these representations into documents in VRML so that one can view these documents with a WWW browser such as Netscape together with a plug-in viewer for VRML. We use a Netscape Navigator plug-in VRML viewer Live3D with the virtual worlds.

There are several advantages for using a virtual reality-based user interface, especially one that has been integrated into a widely available WWW browser such as Netscape Navigator. For example, as a VRML plug-in to Netscape Navigator, Live3D supports several direct manipulation tasks on a virtual world, such as walk, spin, slide, look and point. VRML 1.0 supports a specification object called WWWanchor. When users click on a WWWanchor, the associated document, no matter local or remote, will be downloaded and users can access the document directly from their own browsers. Usability issues specifically associated with Live3D will be discussed later in this paper. We intend to produce an atlas of interconnected documents to allow users manipulate its components easily.

New ways of interaction become possible with the virtual reality-based interface, such as seamlessly walking through, which effectively overcomes the traditional focus versus context problem. Traditional information retrieval strategies are augmented with an intuitive direct manipulation metaphor. Users can move closer to the virtual world for more details as well as move further away for a global view.

We utilised a technique provided in VRML called Level of Detail (LOD) in order to provide an intuitive and seamless interface for users travelling across virtual worlds in various situations. As the user approaches to an object in the virtual world, the virtual world displays increasingly more information about the object. For example, by moving close to a cluster document, the full structure within the cluster is displayed in the virtual world.

Existing empirical evidence suggests that graphical overviews are significantly useful for users to browse easily and effectively [5]. Virtual reality-based user interfaces provide a promising means of helping users in accomplishing these cognitive needs. For example, by explicitly representing salient relationships between two documents in a virtual link structure, users are able to see the connectivity patterns in the entire information space. Virtual link structures of different natures, be hyperlinks, content similarity or navigation patterns, can be readily presented to users as virtual worlds for them to explore intuitively.

Users may benefit from the virtual reality-based user interface in a number of ways. For example, direct manipulation-based user interfaces are often easier to learn and to use than other types of user interfaces such as command driven and menu driven.

As noted by Fairchild et al. [9], a graphical user interface should be designed such that users will be able to identify individual elements in the large space easily, to make sense of the spatial structure in visualisations as a meaningful context, and to recognise salient relationships between elements.

On the other hand, it is not easy to measure the impact of a graphical representation of the structure of a complex network. It is difficult to determine the effects of such visualisations on the performance of users on a range of tasks. We are currently investigating the relationship between such visualisation and individual differences, especially the interaction between users' spatial ability, cognitive styles and various visualisation representations.

The design of virtual reality-based user interfaces is still at early stages. A widespread use of VRML 2.0 will add more interactivity to virtual worlds on the Internet. Based on our experience in this design using VRML 1.0, we have identified some problems that should be addressed in future work for refining the integrated and iterative design process.

4. Design and Verification

We use the following three examples to illustrate the design rationale of building virtual reality-based user interfaces for digital libraries.

4.1 Levels of Detail

One problem identified from the initial user feedback on an earlier prototype was that document spheres should have some labels that are displayed all the time in a virtual world. In the earlier prototype, the label was only displayed if the mouse cursor is on the document object. We have tried to cover the surface of an object with an image and the image contains some keywords or the title of the document. However, VRML blurs the display intentionally unless the view is straight to the surface.

Figure 3 shows a detailed local layout and inter-document similarities. For example, the overall similarity between paper *pp* and *pp2* is 0.69 and the similarity between paper *pp2* and *skc1* is 0.83, but *pp* and *skc1* are not directly connected, suggesting that they are only related because they both have some connections to *pp2*. Such information filtering is useful for users who need to determine the most salient relationships among different documents.

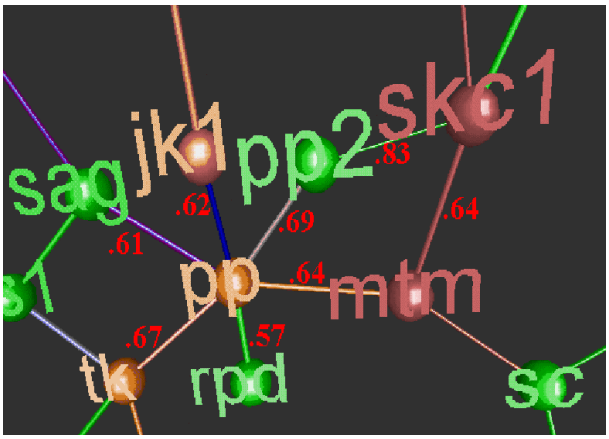

Figure 3. Local details in the spatial layout.

In Figure 4, the user is approaching to paper pp2. Figure 4.a is an overview map of papers in the proceedings. Figure 4.b is the local structure of paper pp2. In Figure 4.c, more information about paper pp2 becomes visible, including the paper's title, the authors' names and a list of keywords. Finally, in Figure 4.d, the user can see the abstract of the paper. The user can actually read the full-text version of the paper published on the ACM site by clicking on the cube in the centre of the screen.

4.2 Query-Specific Views

The user interface can visualise specific queries in a larger context. Figure 5 is a query-enhanced structural visualisation of papers in Hypertext'97 conference proceedings. It shows how strongly a query term is associated with a paper. In this example, three keywords were specified: *WWW*, *retrieval* and *visualisation*. For example, the long, red bar associated with the paper by Anderson in the map indicates the document is likely to be relevant to the WWW.

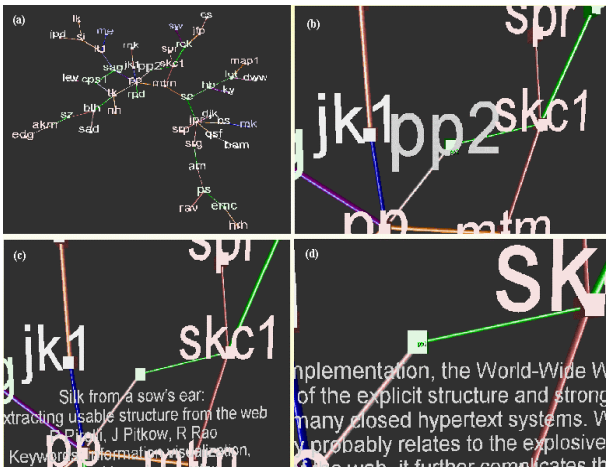

Figure 4. Varying details regarding paper pp2.

4.3 Design Issues

The Virtual Reality Modelling Language (VRML) is bringing widespread 3D graphics to the Internet. VRML is a format for sending descriptions of virtual worlds back and forth over the Internet. 3D virtual worlds become increasingly accessible to home computer users. Netscape Navigator, the world's most popular Web browser, has incorporated VRML viewers such as Live3D and Cosmo Player. For example, Live3D supports animation, rotating objects, streamed video and audio, and other multimedia effects.

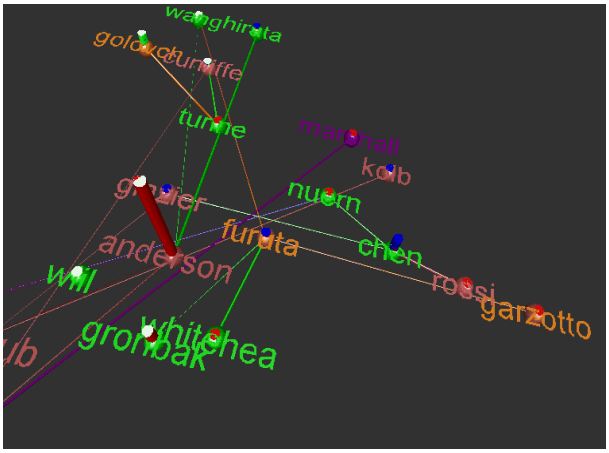

Figure 5. Hypertext'97 papers and queries.

For millions of Internet users, the greatest attraction is that they can access these graphical representations from various platforms that they already have to access the Internet. Users can easily incorporate VRML models into the design of their own Web pages.

A basic virtual reality is constructed by arranging 3D objects into scenes and worlds. A group of 3D shapes are

calculated and rendered on the screen from a particular point of view. The user can then change his or her point of view as if he or she is working through, flying over, or crawling beneath the objects. Since the objects are essentially mathematical or geometric representations, they can be rotated, shifted, or zoomed in a virtual world.

Currently, users must explicitly switch back and forth between browsing the WWW and browsing a virtual world. To walk through, fly through, or examine an object in the virtual world, the user must use the VRML mode, which consists of several actions such as walk, spin, look, point, and view. If the user wants to have a closer look at an object, he or she must select the point function and then point and click on the object. The pointed object will fly closer towards the user. On the other hand, if the user simply wants to follow the Web link attached to the object, then he or she must de-select the point function. A more seamless interface would be able to distinguish the two actions in a single mode, for example to use single clicks for following an http link and use double clicks for moving closer to an object.

4.4 The Compendium vs Random Graphs

The design of virtual worlds needs to provide an interface that users can draw upon the skills that they have developed in the real world [9, 12]. In the following sections, we compare the spatial layout of the ACM Hypertext Compendium and that of a random graph so as to identify the relationship between spatial patterns and the underlying structure.

The network in Figure 6 visualises the structure of 126 papers in the Compendium. Figure 7 is a Pathfinder network derived from a random graph with the same number of nodes. We expect to identify some distinct spatial patterns as the basis for our subsequent interpretation.

First, the proximity data in the Compendium and the random graph have different levels of mean and standard deviation (see Table 2).

Table 2. Similarity measures for the Compendium and a random graph.

Variable	Cases	Mean	Std Dev
COMPENDIUM	15625	.0138	.0666
RANDOM	15625	.5022	.2885

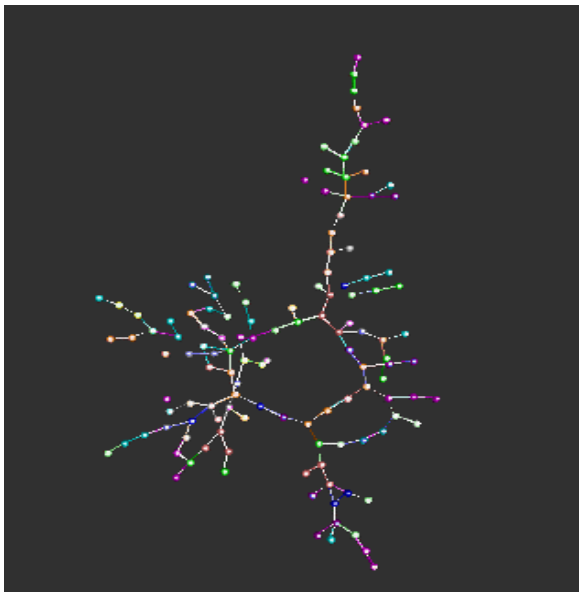

Figure 6. Pathfinder network for the Compendium (126 nodes, 132 links).

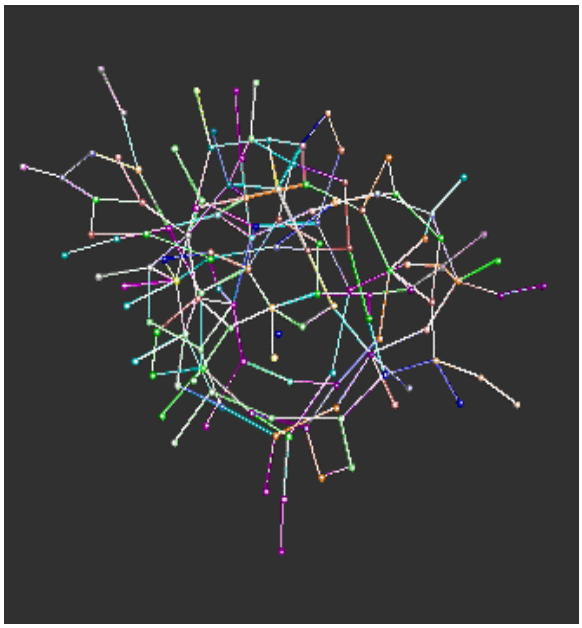

Figure 7. Pathfinder network for a random graph (126 nodes, 164 links).

It is clear from the statistics and the spatial layout, these two data sets have little in common. The spatial layout of the Compendium network appears to reflect some salient relationships among documents. In the next section, we will verify individual documents in association with some distinctive patterns in the network.

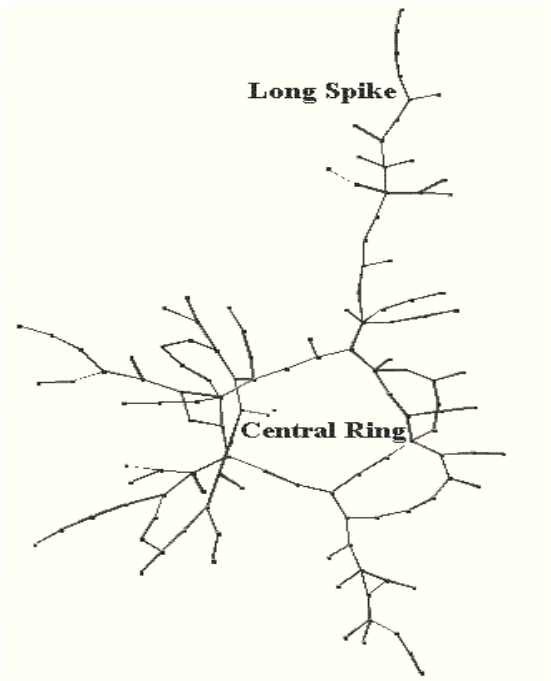

Figure 8. The long spike and the ring.

Second, there are 126 nodes and 132 links in the Compendium network. The link/node ratio is 132/126 = 1.048, suggesting that in this network one node is usually connected to one or two nodes only. Figure 6 shows a lower degree of interconnectivity in the Compendium network. A lower degree of interconnectivity makes it easier for users to identify the most salient connection between two arbitrary nodes.

In contrast, the random graph network has 126 nodes and 164 links. The link/node ratio is 164/126=1.302. In the random graph network, one node is more likely to be connected with two or more nodes than its counterpart in the Compendium network. Figure 7 confirms that the random graph network has a higher degree of interconnectivity. There are a considerable number of rings and it is more likely that two arbitrary nodes in the random graph network are connected by more than one path.

According to this analysis, long spikes in the Compendium network are more likely due to substantial similarities among different papers. On the other hand, although rings are more likely to appear in a random graph network, its appearance in a subject-specific network may suggest a group of highly interrelated documents; therefore, rings in the Compendium network may also reveal some latent relationships. To further investigate the nature of the resultant spatial configuration for the Compendium, in the next section, we examine documents associated with the longest spike and the largest central ring structure in the network.

4.5 Verification

The spatial layout of the Compendium network shows a degree of clustering in that some papers tend to be associated together as a group, although the boundary of such a group may not be clear.

In Figure 8, there is a ring of papers in the centre of the Compendium network; there is also a long spike pointing outwards from the ring. What could these spatial patterns tell us? We will examine documents associated with the ring and the long spike in the following section.

The longest spike contains more than 30 documents. We verified the first 12 documents starting from the remote end of the spike (see Table 3).

Table 3. Documents in the longest spike.

HTC	SOURCE	AUTHOR. TITLE
98	HT87	Smolensky, P. *et al*. Constraint-Based Hypertext for Argumentation
25	MCC87	Conklin, J. A Survey of Hypertext
105	HT87	van Dam, A. Hypertext '87: Keynote Address
42	HT87	Frisse, M. Searching for Information in a Hypertext Medical Handbook
46	HT87	Garg, P. Abstraction Mechanisms in Hypertext
86	HT87	Raymond, D. & Tompa, F. Hypertext and the New Oxford Dictionary
52	HT87	Halasz, F. Reflections on Notecards
53	CHI87	Halasz, F. *et al*. NoteCards in a Nutshell
36	HT89	Evenson, S. *et al*. Towards a Design Language for Representing Hypermedia Cues
103	HT87	Trigg, R. & Irish, P. Hypertext Habitats
104	CSCW88	Trigg, R. Guided Tours and Tabletops
111	HT89	Zellweger, P. Scripted Documents

It turned out that these documents include some of the most influential works in the hypertext literature, including Conklin's landmark survey, Frisse's Hypertext Medical Handbook, Garg's Hypertext Abstract Machine, the Oxford Dictionary conversion and Halasz' Seven Issues. The distinct shape of the long spike would easily attract the attention of users to these classic works, which would be a desirable feature for users who are not familiar with the field.

The central ring consists of 16 documents, starting from the root of the longest spike clockwise (see Table 4). Some connections are easy to detect. For example, documents 74, 75 and 96 are connected because they all address writing with hypertext systems. Similarly, documents 1, 54, 91 and 77 are grouped together because they describe models for hypertext systems.

Table 4. Documents in the central ring.

HTC	SOURCE	AUTHOR. TITLE
109	HT89	Yoder, E. & Wettach, T. Using hypertext in a law firm
99	ECHT90	Stotts, P. D. & Furuta, R. Hierarchy, composition, scripting languages, and translators for structured hypertext
22	HT87	Collier, G. Thoth-II: Hypertext with explicit semantics
20	HT87	Charney, D. Comprehending non-linear text: Reading strategies
74	HT87	Neuwirth, C. *et al*. The Notes program: A hypertext application for writing from source texts
75	HT89	Neuwirth, C. & Kaufer, D. The role of external representations in the writing process
96	HT87	Smith, J. *et al*. A Hypertext writing environment and its cognitive basis
1	ECHT90	Afrati, F. & Koutras, C. A hypertext model supporting query mechanisms
54	NIST90	Halasz, F. & Schwartz, M. The Dexter hypertext reference model
91	ECHT90	Schutt, H. & Streitz, N. HyperBase: A hypermedia engine based on a relational database management system
77	ECHT90	Ogawa, R. *et al*. Scenario-based hypermedia: A model and a system
84	ECHT90	Puttress, A. & Guimaraes, N. The toolkit approach to hypermedia
116	HT89	Ehrmann, S. *et al*. Panel: Hypertext and higher education: A reality check
122	HT89	Meyrowitz, N. *et al*. Panel: Confessions--What's wrong with our systems?
128	ECHT90	Streitz, N. *et al*. Panel: What's specific about user-interfaces for hypertext systems?
73	HT87	Theodor H. Nelson. Invited talk: Hypertext '87: Foreword

6. Discussion

We have described a generic approach to the design of a user interface for a digital library. On the other hand, there is still a lot of work to be done. Evaluating the usability of such design requires a re-examination of existing theories and methodologies of human-computer interaction. We have identified the following issues for further studies:

- *Scale up clustering and graph drawing algorithms*. To reduce the computational complexity in order to handle a large collection of documents.
- *Integrate with established information retrieval systems such as the SMART system*. To make use its evaluation facilities.
- *Build a comprehensive digital library on hypertext*. To build a digital library with high-quality documents and an intuitive structure map.

These tasks will be undertaken in our forthcoming two-year project funded by the Engineering and Physical Sciences Research Council (EPSRC) under the Multimedia and Networking Application Programme (MNA). The project is to develop a development and evaluation testbed for structuring and visualising large distributed hypermedia.

6. Conclusion

The virtual reality-based user interface described in this paper offers an intuitive and effective means of navigating in a digital library. The new user interface takes the advantage of the widely used Web browser and VRML viewers such as Live3D. Users who have access to the WWW can immediately access such interfaces from their existing computers. The techniques described here can be applied to other subject domains.

The design practice described in this paper provides some valuable experience and lessons for developing an intuitive representation of the complex structure of a digital library. In addition, a number of virtual link structures can be used together to provide users insights into the underlying structure from various perspectives. More work is needed in a number of areas, such as evaluating the usability of the new user interface, optimising clustering and graph drawing algorithms, investigating the role of individual differences in the use of spatial-oriented user interfaces, especially the spatial ability and cognitive styles. The approach described in this paper provides a promising way for building intuitive graphical user interfaces for users to browse seamlessly in a digital library.

Acknowledgements

The work was conducted while the author was at Glasgow Caledonian University. A copy of the Hypertext Compendium was obtained from the ACM.

References

[1] Akscyn, R. (Editor) *The ACM Hypertext Compendium™*, ACM Press, 1991.

[2] Carr, L., Hill, G., De Roure, D. and Hall, W. 'Open information services' *Computer Networks and ISDN Systems*, 28 (1996), 1027-1036.

[3] Chalmers, M. 'A linear iteration time layout algorithm for visualising high dimensional data' in *Proc. of IEEE Visualization* (San Francisco, Oct. 1996). http://www.ubs.com/ubilab/Publications/ Cha96a.html

[4] Chen, C. 'Structuring and visualising the WWW by generalised similarity analysis' in *Proc. of the ACM Hypertext'97* (Southampton, England, April 1997), ACM Press, 1997, pp. 177-186.

[5] Chen, C. and Rada, R. 'Interacting with hypertext: A meta-analysis of experimental studies' *Human-Computer Interaction, 11*, 2 (1996), 125-156.

[6] Cooke, N. J., Neville, K. J. and Rowe, A. L. 'Procedural network representations of sequential data' *Human-Computer Interaction, 11*, 1 (1996), 29-68.

[7] Croft, W. W. and Turtle, H. 'A retrieval model for incorporating hypertext links' in *Proc. of the ACM Hypertext'89* (Pittsburgh, PA), ACM Press, 1989, pp. 213-224.

[8] Crouch, D. B., Crouch, C. J. and Andreas, G. 'The use of cluster hierarchies in hypertext information retrieval' in *Proc. of the ACM Hypertext'89* (Pittsburg, PA), ACM Press, 1989, pp. 225-237.

[9] Fairchild, K., Poltrok, S. and Furnas, G. 'Semnet: Three-dimensional graphic representations of large knowledge bases' in *Cognitive Science and its Applications for Human-Computer Interaction*, R. Guindon, Ed. Lawrence Erlbaum, 1988.

[10] Furnas, G. 'Generalized fisheye views' in *Proc. of CHI'86*, ACM Press, 1986, pp. 16-23.

[11] Gronbaek, K., Bouvin, N. O. and Sloth, L. 'Designing Dexter-based hypermedia services for the World-Wide Web' in *Proc. of the ACM Hypertext'97* (Southampton, England, April 1997), ACM Press, 1997, pp. 146-156.

[12] Leung, Y. K. and Apperley, M. D. 'A review and taxonomy of distortion-oriented presentation techniques' *ACM Trans. on Computer-Human Interaction, 1*, 2 (1994), 126-160.

[13] Mackinlay, J. D., Robertson, G. G. and Card, S. K. 'The Perspective Wall: Detail and context smoothly integrated' in *Proc. of CHI'91* (New Orleans, LA), ACM Press, 1991, pp. 173-179.

[14] McDonald, J. E., Paap, K. R. and McDonald, D. R. 'Hypertext perspectives: Using Pathfinder to build hypertext systems' in *Pathfinder Associative Networks: Studies in Knowledge Organization*, R. W. Schvaneveldt, Ed. Ablex Publishing Corporation, Norwood, NJ., 1990, pp.197-212.

[15] NSF Graphics and Visualization Center 'Memex and Beyond' http://www.cs.brown.edu/memex/, 1996.

[16] Pirolli, P., Pitkow, J. and Rao, R. 'Silk from a sow's ear: Extracting usable structures from the Web' in *Proc. of CHI'96*. http://www.acm.org/sigchi/chi96/proceedings/papers/Pirolli_2/pp2.html, 1996.

[17] Robertson, G. G., Mackinlay, J. D. and Card, S. K. 'Cone Trees: Animated 3D visualisations of hierarchical information' in *Proc. of CHI'91* (New Orleans, LA). ACM Press, 1991, pp. 189-194.

[18] Salton, G., Allan, J. and Buckley, C. 'Automatic structuring and retrieval of large text files' *Commun. ACM, 17*, 2 (1994), 97-108.

[19] Savoy, J. 'An extended vector-processing scheme for searching information in hypertext systems' *Information Processing & Management*, 32, 2 (1996), 155-170.

[20] Weiss, R., Velez, B., Sheldon, M., Nemprempre, C., Szilagyi, P., Duda, A. and Gifford, D. 'HyPursuit: A hierarchical network search engine that exploits content-link hypertext clustering' in *Proc. of the ACM Hypertext'96* (Washington, DC. March, 1996). http://www.psrg.lcs.mit.edu/ftpdir/ papers/