Spatial Hypertext as a Reader Tool in Digital Libraries

George Buchanan[1], Ann Blandford[2], Matt Jones[3], Harold Thimbleby[2]

[1]Middlesex University, Bramley Road, London, N14 4YZ, UK
g.buchanan@mdx.ac.uk
[2]UCL Interaction Centre, 26 Bedford Way, London
{a.Blandford, h.Thimbleby}@ucl.ac.uk
[3]Waikato University, Hamilton, New Zealand
always@acm.org

Abstract. Visual Interfaces may facilitate human to computer interaction as well as computer to human communication. In this paper, we introduce Garnet, a novel visual interface for interaction between humans and Digital Libraries. Garnet provides a visual workspace in which the user can structure and organize documents of interest. This structure is then used to organize and filter further documents which may be of interest, such as search results. Spatial hypertexts are introduced as a framework for creating DL interfaces, and Garnet is compared to existing DL and Spatial Hypertext systems.

1. FOREWORD

Visual interfaces for Digital Libraries mediate communication between the user and the system. Many sophisticated visual interfaces exist for complex information resources. Most of these exploit the visual cognitive powers of humans to make the structure of all or part of the library more accessible to its users, borrowing particularly from the field of information visualization [2].

However, focusing on the visual system as a powerful means of communicating to the user is to lose half the story. Visual interaction is used to powerful effect on the desktop to express instructions from the user to the system, and visual spaces – physical and virtual – can also be powerful resources for collaboration between users, and for tasks such as organizing information [11].

This paper presents one approach, instantiated in a prototype system which exploits a visual workspace in this second manner, centered on the user's task rather than the data and structures of the digital library itself.

2. INTRODUCTION

Information seeking [10] is a demanding and complex task which forms the driving motivation behind the adoption of digital libraries. There is typically unequal support for information seeking and the subsequent information structuring process [11, 12]. Information seeking is well supported through searching and browsing facilities, for example, whereas there is little or no support for organizing and collating the discovered documents into user-generated structures – a process which improves comprehension and reflection. (Note that we use "document" to refer to a discrete item of information recorded in any medium: a journal paper; video excerpt etc.)

Our system, Garnet, is intended to assist users in these latter phases of document collation and organization. This task is the target of a family of tools known as Spatial Hypertexts. The concept of spatial hypertext was introduced by Marshall and Shipman – a virtual workspace in which each document is represented by a shape, and cues such as its position, color or form are used to indicate document relationships, purposes, etc.

Marshall and Shipman's systems, VIKI [8, 9] and the later VKB [13], have both informal and formal structuring features. Informal structuring facilities include the use of color and proximity of document shapes. Formal structuring facilities are represented by the ability to create a formal hierarchical structure of document collections, just like folders or directories in a filing system, for example. Garnet's structuring facilities and interactions are similar to those of VKB. However, unlike VKB, Garnet is deeply interconnected with a number of information systems.

Marshall and Shipman intended, and expected, a number of benefits in the use of spatial hypertexts. One was that, by using the hypertext, information workers could articulate their findings, expectations and conclusions to others. Another was that it provided the opportunity for users to clarify their own thinking through the process of organizing, selecting and rating documents. Their studies into the use of VIKI within information work had positive outcomes, so it appeared that a similar approach was appropriate in supporting DL users in document organization.

VIKI has itself not been connected directly to an information discovery system such as a digital library, so its own support for information seeking as a whole has a complementary role to existing digital library systems. This, therefore, suggested that an integrated system would bring benefits not found in two, disconnected, digital library and spatial hypertext systems.

In addition, for many information workers information seeking and information structuring are long-term tasks. If documents are organized and re-organized in a workspace over a period of time, they become an evolving reflection of the information work the worker is engaged in.

We wondered whether the implicit information in the workspace of a spatial hypertext could be extracted and used to amplify the effectiveness of the user at a later date through, say, improved precision or recall in retrieval, or assistance in the organization of documents. This is similar to the economic concept of the positive externality, where one person's actions indirectly benefit another, but in our case a user's past activity benefits their present work.

Therefore, in Garnet we are introducing features which exploit the implicit knowledge which can be discovered in the contents of a workspace, combining that knowledge with the advantages of a direct connection between the workspace and digital libraries. The first of these is a feature we call "scattering", in which the organization of the user's workspace is used as a filter over search results, both improving the identification of documents relevant to the user's interests and facilitating consistent placement for documents into the existing organization of the workspace. We will now illustrate the use of this feature and its benefits.

3. SAMPLE SCENARIO

A pilot version of Garnet has been created, which is connected to the New Zealand Digital Library Project's Greenstone software [16]. Greenstone is a comprehensive Digital Library software system, supporting common actions such as full-text and index searching, and browsing in category hierarchies.

We will now walk through the system in use, starting from a "bare" hypertext. The library material we will use is the Humanity Development Library of the United Nations, one of the widely available examples of a Greenstone library collection, which consists of several thousand pages.

3.1 Overview

In Figure 1, we see a Garnet user session in progress; a *window* appears inside the main application window. This is a *collection* of materials which the user has recorded in the current, or a previous, session. Each document is represented by a rectangle containing some text, as indicated in the diagram, which we term a *label* for simplicity. The user can create a hierachy of as many collections as they wish.

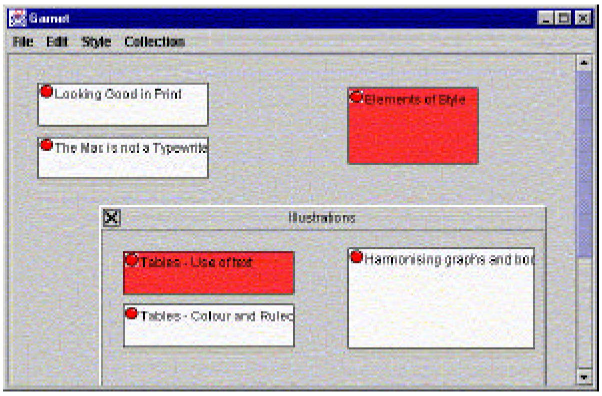

Fig. 1. A Garnet Client in use

Within a collection, the user is free to place, size and color each document label as they see fit – the space is entirely freeform. Labels can be moved and/or copied between collections in the usual way for similar direct manipulation environments. Document labels can be added explicitly by the user or through interaction with a digital library's facilities, e.g. search.

Therefore, the user is free to use the document labels both in freeform structures of their own making inside collections, and in a more formal organization by using the explicit hierarchical forms of a set of document collections. Taking the example above, we have a collection called "Illustrations", which has a column of documents on the left-hand side, and a single document on the right. The column is a structure created by the user's exploitation of space – it is not a feature enforced by the system.

The column idiom can also be seen in the root collection – again on the left-hand side. Some use of color can be seen here, but the relationship is not clear to a reader "from the outside", though it may be perfectly clear to the creator of the hypertext.

Larger examples, including examples of other structures such as tables (grid-like arrangements), piles (documents placed apparently randomly) and composites (groups of documents with repeated internal structures) can be seen in Marshall and Shipman's papers on the subject [8, 12].

3.2 Example search

Let us now follow a simple sequence of interactions, starting with an example search. For our purposes, we are going to investigate snail farming, in an attempt to discover whether we have the appropriate resources to consider that form of agriculture. With Garnet loaded, we start a new search in the Greenstone system (Garnet also supports searching via Google), and we enter the simple query "snail". As shown in Figure 2, a simple collection window appears with a number of document labels appearing one beneath the other, similar to a typical search result list.

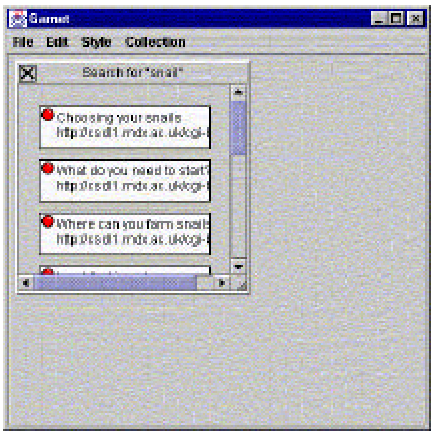

Fig. 2. A simple search

On reading the first two documents (achieved by a simple double-click on the appropriate documents), we decide that we'd like to keep the second document ("What do you need to start?"), and we move it to our root workspace window – simply dragging the document from the "Search for 'snail'" window onto the main Garnet window.

The first document, however, seems a bit advanced, and we can delete it from the list, clicking on the small red 'circle' on it's top left corner (figure 3). As a result of this, the later documents move upwards. Should we wish to return to the search results at a later date, by default these changes would be retained.

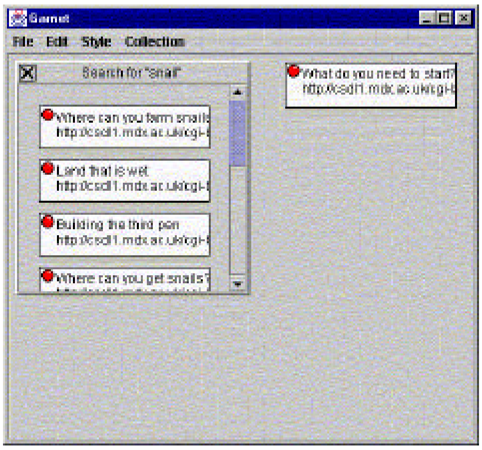

Fig. 3. The workspace after initial alterations – note changes in search list c.f. Fig. 2

3.3 Demonstration of "scatter results"

In the previous search, we performed a plain search. Garnet, however, can exploit the organization done by the user in a novel manner. We can "scatter" a set of documents (including search results) over the existing layout of documents in the workspace. "Scattering" places the search documents near to groups of existing documents with which they have a strong similarity.

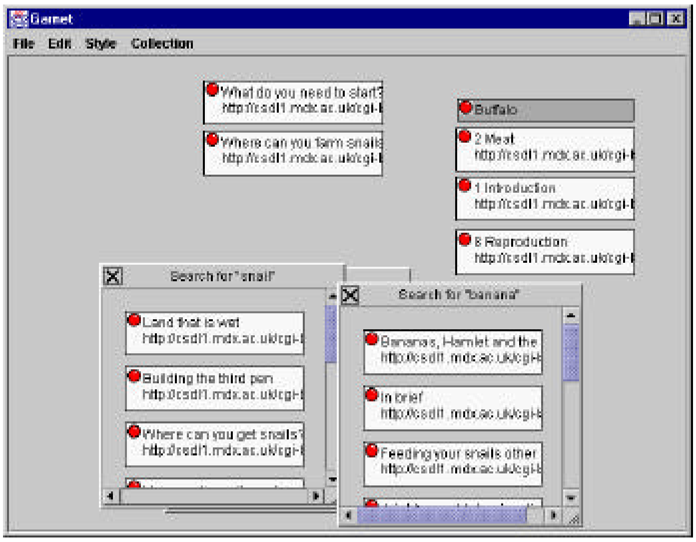

Fig. 4. Situation before "scatter"

Continuing our previous example, we have now selected a variety of documents of interest, but let us suppose that a couple of questions remain unanswered.

Suppose we have a plentiful supply of bananas which we would like to use, but we are not sure whether this food would be appropriate. If we were to do a naïve search, on “banana”, the initial results do not well match our particular interest (Figure 4).

In fact, documents which relate to our interest can be found in both the ‘snail’ and ‘banana’ searches. However, these documents of interest may not appear at the very top of either list. Normally, we would have to try and re-work our query manually to make it more targeted. In the case of Garnet, we could use the ‘scattering’ feature to discover any material similar to documents we have already selected for storage. Or, in other words, Garnet can generate existing search terms or filtering to represent our user’s interests, based on the workspace layout they have already created.

Viewing Figure 4 again, note the third item from the top of the list “Feeding your snails …” (for clarity in this example we’ve chosen something that is visible). If we now do a “scatter”, (Figure 5), a subset of the “banana” search results appear on the main collection. This small subset, which appears in a light gray below, has been found by Garnet to be a close match to the existing pair of documents, which appear in white. Suggestions are always displayed in this gray color, and below and to the right of the group of documents which they are believed to be similar to.

We can now investigate the two suggested documents which are similar to the previously selected pair. As it happens, these documents would confirm that ripe bananas can indeed be used to feed snails. If we wanted to permanently add one or other suggestion to the workspace, we can click on the ‘circle’ which appears on the top right corner of each of the suggestions.

If we no longer wish to see the existing suggestions, or when another set of documents is scattered, the current suggestions are cleared.

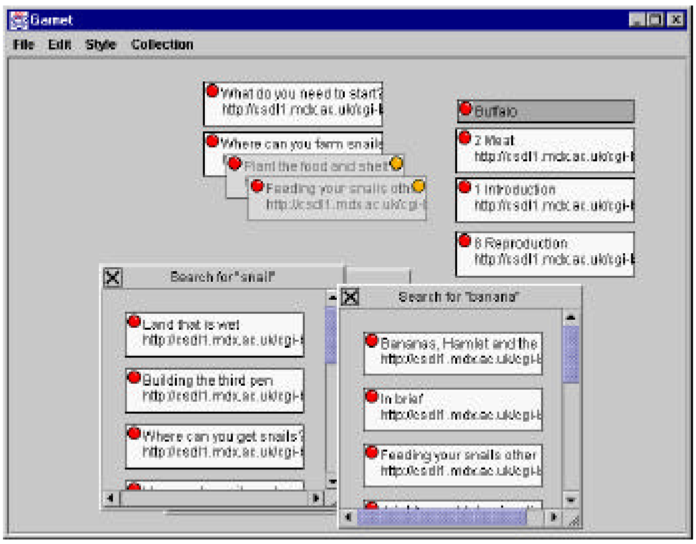

Fig. 5. After “scatter” – note the shaded document labels added in comparison to Fig. 4

We will now discuss the underlying technologies that support this behavior, and the general operation of Garnet.

4. SUPPORTING TECHNOLOGIES

At the heart of Garnet is a spatial parser [7,12], which identifies visual patterns within the arrangement of shapes (i.e. based on shape/position similarity rather than textual similarity). Our spatial parser simply identifies groups of closely placed documents; it does not distinguish visual patterns within groups of documents such as stacks or rows. Though proximity is a simple measure, it permits a wide variety of complex shapes to be successfully identified without recourse to a special recognizer for each pattern. For each identified group, the contents of the corresponding documents are used to generate lexical classifiers to represent the common topic of the documents.

The lexical classifiers are used in matching documents to the existing document groups. So, returning to our example above, the two documents on the root collection in Fig. 4 are seen as one 'group' by the spatial parser. The common topic of the group is calculated by the lexical analyzer, and a corresponding classifier produced. When the 'scatter' action was executed, Fig. 5, the classifier found a match for the two documents which were suggested. The spatial parser was again used to determine where the new, matching document labels were placed.

Importantly, a document does not yet have to appear in the workspace to match a 'scatter'. For instance, we only display the first ten matches for a search in the corresponding collection in a workspace. When the search is 'scattered', matching is done against the full returned search results, not just the top ten matches. Therefore, scattering can filter large result sets which would otherwise be laborious to browse.

To build the lexical classifiers, Garnet needs to profile each document in its workspace, and then agglomerate these individual document profiles into those used to represent each group of documents. For web documents, found through Google, Gamet builds this profile itself. In the case of digital library documents from Greenstone, this information is passed through an extension of Greenstone's digital library protocol. For other digital library protocols, such as Z39.50, Garnet could be altered to again building its own profile of documents. The profile of a document is a list of the most common words in a document (excluding stop-words such as 'and', 'the' etc.), and the profile of a group is a selection of the words a group has in common using a variant of Zamir and Etzioni's clustering method [17].

Many other capabilities can be generated from this combination of spatial and lexical analysis, such as mapping between different workspaces, or consolidating visually isolated yet topically related groups.

One significant and predictable problem may arise from groups of documents which are highly heterogeneous (e.g. those yet to be sorted and reviewed by the user) – in such circumstances, the corresponding classifier would be of little use, often matching many disparate documents. These "miscellaneous" sets were often observed by Marshall and Shipman, and occur due to the provisional nature of the organization and collation task. As our classifiers are based on clustering techniques [4] which include the identification of incoherent clusters, we use this common clustering function to avoid building active classifiers for heterogeneous groups. Clustering techniques can also permit Garnet to assist in the organization of disorganized groups by providing a "first-cut" at partitioning them. Our use of Zamir and Etzioni's clustering algorithm [17], which does not use word weighting, is also a distinct advantage when working across collections or when details of word weights is not available (which is often the case).

When the workspace is changed, the spatial parser must re-generate the corresponding classifiers. At present, this evaluation is done eagerly, but this could cause Garnet's response to slow down when the workspace is changing rapidly; therefore, we are moving towards lazy evaluation. Users often temporarily place individual documents arbitrarily during the ongoing flow of their work, and detecting this is problematic. The matching algorithm again helps, being highly robust to noisy data, but delaying classifier construction will improve accuracy as well as speed.

5. PREVIOUS WORK

Having described some features of Garnet, we will now review a small sample of systems similar to Garnet and compare them to each other and Garnet. Firstly, we examine a number of existing visual interfaces for Digital libraries, and then we will study some spatial hypertext systems. Finally, we will discuss some issues of interest which have arisen from implementing Garnet.

5.1 Visual Interfaces for Digital Libraries

A number of visual interfaces specifically designed to support extended work in Digital Libraries already exist, e.g. DLITE [2], NaviQue [5] and SketchTrieve [6]. All three are intended to give coherent access to a number of information services (e.g. search) and sources (different collections and libraries), and represent separate searches as discrete objects in a 2-d workspace. The workspace supplies a persistent context for storing and recalling previous activity (e.g. searches) across a long period of time, with the intention being to reduce unnecessary repetition of work, and to assist in the co-ordination of repeated searches across numerous information sources.

From the perspective of Spatial Hypertext, a major question would be the range of expressiveness which these systems give in the organization and annotation, implicit and explicit, of documents. Therefore, we will now consider how much control the user has over the appearance of objects in their workspace, the relative significance of different objects in the workspace, and (considering our immediate interest) in what manner they can use the objects in their workspace to perform further work.

The representation of individual documents varies considerably. In DLITE documents appear as small text-less icons, and only color can be set by the user. SketchTrieve's representations are larger, include text, and the user has complete control of their appearance. NaviQue uses a similar representation but it emphasizes zooming as a method of browsing large-scale areas (a property inherited Bedersen et al's Pad++ [1]). A consequence of zooming is that when a wide area is visible, the impact of individual documents is very small – often just as single points of color.

Structuring facilities are generally weak. In DLITE the ordering of documents within sets is system-controlled, and sets cannot be organized hierarchically, only as a set of peers. NaviQue lacks formal document groups – the user identifying sets of documents ad-hoc by explicitly selecting each document. SketchTrieve falls between these two positions: document sets do exist, but only as the results of a search – sets therefore cannot be used for organizing document groups. In comparison, Garnet has formal sets which can be organized hierarchically. Documents are added to a set at the user's discretion, and the internal organization of each set is fully within the user's control, except when the system has created the set (e.g. search results).

The focus of these interfaces has been connecting the information work-space to an information system. Here, DLITE is very extensible, though most of its standard facilities are traditional (e.g. browsing), and the core remains search. NaviQue is more developed, possessing a 'similarity engine' which permits the comparison of documents or sets of documents. Though similar to the underpinnings of Garnet's "scatter" facility, in NaviQue the principal benefit is navigational assistance, rather than support for organization and structuring.

Overall, therefore, these interfaces facilitate traditional digital library actions and exploit a few spatial hypertext idioms to assist long-term work. However, compared to Garnet and Spatial Hypertext systems, the scope for emphasizing, organizing and structuring documents is weak. Consequently, such features have not been leveraged further, as we have done with Garnet.

Thus, in its inheritance from spatial hypertexts, Garnet contains strong structuring and representational facilities compared to 'traditional' visual interfaces. In addition to this, Garnet's document comparison and matching methods, and its evaluation of the user's workspace have no direct equivalent in existing interfaces.

5.2 VIKI, VKB and Spatial Hypertexts

As already stated, VIKI and other spatial hypertexts have generally had a limited connection to information systems such as digital libraries. This limitation has consequences, no matter the benefits obtained when organizing documents.

VIKI, like Garnet, performs spatial parsing – identifying document groups by their visual properties (though the implementations differ). In VIKI this analysis is used to support visual interaction and direct manipulation in the workspace. As VIKI's connection to digital libraries is at best limited, no related textual analysis can be performed. This feature is found, however, in NaviQue and Garnet.

Though VIKI has been extended to obtain search results from the Web, when compared to the range of facilities accessible through D-LITE or Garnet, the connection to the digital library is basic, lacking browsing, metadata information, etc.

A recent paper on VIKI's successor, VKB [15], introduces a "suggestions" facility, similar to Garnet's "scatter". However, the main role of suggestions in VKB is proposing visual properties for documents already in the workspace, whereas scattering in Garnet can introduce documents not yet displayed in the workspace

Visually, Shipman et al's approach is also different to ours – we place the suggestions into the workspace near to matching documents, whereas in VKB suggestions appear in text in a pop-up modal dialogue. Thus, the ability of our spatial parser to place suggestions does not appear to occur in VKB. The text matching algorithms used are not described, so a detailed technical comparison is not possible.

5.3 Summary

When one reviews the comparison of Garnet to existing visual interfaces for DLs, and to existing Spatial Hypertext research, some common themes emerge. This may note be surprising, as the papers on SketchTrieve and NaviQue refer to spatial hypertext research (VIKI in the former case, Pad++ in the latter). So, spatial hypertexts certainly influence research on visual workspaces in digital libraries.

Frank Shipman recently identified seven outstanding issues in spatial hypertext research [14], which also provide an interesting focus for visual interfaces for DLs. For example, Shipman highlights the expressiveness of different representations for documents and groups and as we have seen this is also a substantial issue for visual interfaces. In the DL interfaces discussed here and in Garnet, both user and system can place objects into the workspace – raising issues of ownership and control which Shipman notes in multi-user spatial hypertexts, but which are also present when the system plays a more active role. Similarly, Shipman questions how to connect (technically and interactively) spatial hypertexts to the general information environment and this has been an enduring issue in DL interfaces.

The most novel facet of Garnet – the evaluation of user workspace, and the subsequent use of that interpretation to assist the user in document placement and interpretation – has not previously been seen in either field, and represents one development closely related to one of the most underdeveloped areas highlighted by Shipman – computation in and over spatial hypertext.

6. FUTURE WORK: USER INTERFACE ISSUES

Having successfully implemented a working version of Garnet, we have performed initial usability investigations with a small number of expert subjects using talk-aloud protocol and interviews, to elicit some of the impacts of combining spatial hypertext and digital library systems, and particularly to identify and remove any particularly impeding features before performing a formal evaluation. Subjects were observed interacting with Garnet after an initial introduction to the concepts of spatial hypertext and the particular features of Garnet such as scattering. The general spatial hypertext features were easily adopted, and subjects were able to perform scattering operations successfully. Nonetheless, a number of design questions have emerged regarding the user interface. Firstly, there is the question of user versus system control of certain features; secondly, of how to represent information to the user; and thirdly, there are potential problems of metaphor dissonance, as discussed below.

On the question of control, with the scatter feature, we chose to place the control of this feature fully into the hands of the user. If scattering were applied automatically, document labels representing the suggested position of the documents could appear and disappear continuously, removing the user's sense of control and continuity. A related question was also which search results should be scattered at one time. Again, placing this within the user's control should provide a more user-centered approach.

One representational problem is the placing and appearance of scattered documents in their suggested places. We currently use a very simple technique (using a shaded label and placed at the bottom left of the group), but many other possibilities exist. This choice has been informed by the observed choices of human subjects in Marshall and Shipman's work when representing uncertainty and doubt regarding the role of a document – it remains to be seen how universally intelligible this is. One unresolved issue is how to demonstrate the strength of the correspondence; the current display gives no indication of this. A further difficulty is how to inform the user of the positioning of scattered documents which are in areas of the workspace that are not currently visible.

A second representational problem is how to provide browsing of the digital library. The integration of the simple structures found in search facilities is reasonably straightforward, but browsing structures are often much more complex.

Some common desktop/GUI program metaphors do not seem to hold well, creating metaphor dissonance. For instance, in the walk through above (Figures 2 and 3) we deleted a document from the search results. We don't normally destroy documents in a library, or remove items from a list 'owned' by the system, so the 'delete' metaphor is inappropriate. Furthermore, should we bar the document from appearing in the current result set only, or from appearing in later result sets (i.e. "blacklist" it)?

These many questions are only a sample of the many aspects which offer scope for further investigation. Initial responses to the suggestions facility and the general interaction of the system have been favorable, so our next steps will be to more systematically evaluate our novel features and the best means of resolving these implementation details.

7. CONCLUSIONS

At present we have a simple pilot system working; our initial system has shown the potential for making sensible suggestions within the context of our trial Digital Library sources. A number of the issues faced so far have been discussed above. Some parts of the system are as yet incomplete, but we will soon perform a formal evaluation of the system through a full user study, and will then develop Garnet further to discover the issues which arise within a co-operative, multi-user context.

Beyond our own project, we believe that the idioms of spatial hypertext and the understandings of the spatial hypertext community represent a useful input and framework to future research into DL interfaces. Similarly, much can be investigated in visual DL interfaces which will extend the understanding of spatial hypertext. Some of these synergies have been covered in this paper, but others have yet to be fully explored.

8. ACKNOWLEDGMENTS

Our thanks to the University of Waikato's New Zealand Digital Library Project, especially David Bainbridge, Stefan Boddie and Ian Witten. This work was funded by Middlesex University.

REFERENCES

1, Bederson, B. and Hollan, J. Pad++: A zooming graphical interface for exploring alternate interface physics. In Proceedings of the ACM Symposium on User Interface Software and Technology (UIST ’94, Marina del Rey, CA, Nov.) ACM Press, New York, NY 17-26.
2. Card, S. K., Mackinlay, J. D. & Schneiderman, B. “Readings in Information Visualization – using vision to think.” Morgan Kaufmann Publishers, Inc. 1999.
3. Cousins, S. B., Paepcke, A., Winograd, T., Bier, E., and Pier, K. The digital library integrated task environment (DLITE). In Proceedings of ACM Digital Libraries '97 Conference. ACM Press, New York, NY, 142-151.
4. Cutting, D. R., Karger D.R., Pederson, J.O. and Tukey, J.W. Scatter/Gather: a cluster-based approach to browsing large document collections. In proceedings of ACM/SIGIR 1992, 318-329.
5. Furnas, G., and Rauch, S., Considerations for information environments and the NaviQue workspace. In Proceedings of the Third ACM Conference on Digital Libraries (DL ’98. Pittsburgh,PA, June) ACM Press, New York, 79-88
6. Hendry, D. G., and Harper, D. J., An informal information-seeking environment. Journal of the American Society for Information Science, 48(11):1036-1048, 1997.
7. Lakin, F., Visual Grammars For Visual Languages, ProceedingsAAAI-87, Seattle, Washington, July 12-17, 1987, pp 683-688
8. Marshall, C., Shipman, F. and Coombs, J., Spatial Hypertext supporting emergent structure. Proceedings of the 1994 ACM European Conference on Hypermedia Technology (ECHT ’94, Edinburgh, UK) ACM Press, NY, pp.13-23.
9. Marshall, C. and Shipman, F., Spatial Hypertext and the practice of information triage. In Proceedings of the Eighth ACM Conference on Hypertext (Hypertext ’97, Southampton, UK) ACM Press, New York, NY pp.124-133.
10. Marchionini, G. Information Seeking in Electronic Environments. Cambridge Series in Human Computer Interaction, Cambridge University Press, Cambridge, UK, 1995
11. O'Day, V. and Jeffries, R. Orienteering in an Information Landscape: How Information Seekers Get From Here to There, InterCHI'93, ACM Press, Amsterdam, 1993, pp. 438-445.
12. Shipman, F., Marshall, C., and Moran, T., “Finding and Using Implicit Structure in Human-Organized Spatial Layouts of Information”, Proceedings of Human Factors in Computing Systems (CHI '95), 1995, pp. 346-353.
13. Shipman, F., Hsieh, H., Airhart, R., Maloor, P., and Moore, J.M., “The Visual Knowledge Builder: A Second Generation Spatial Hypertext”, Proceedings of the ACM Conference on Hypertext, 2001, pp. 113-122.
14. Shipman, F, “Seven Questions on Spatial Hypertext”, 2nd International Workshop on Spatial Hypertext, Aarhus, Denmark, 2001.
15. Shipman, F., Moore, J.M., Maloor, P., Hsieh, Akkepeddi, R., “Semantics Happen: Knowledge Building in Spatial Hypertext”, Proceedings of the ACM Conference on Hypertext, 2002, 25-34.
16. Witten, I., McNab, R., Boddie, S., Bainbridge, D. “Greenstone: A Comprehensive Open-Source Digital Library Software System”. Proceedings of the Fifth ACM Conference on Digital Libraries, June 2000, ACM Press, pp.113-121.
17. Zamir, O., Etzioni, O., Mandani, O. and Karp, R. M. `Fast and Intuitive Clustering of Web Documents'. Third International Conference on Knowledge Discovery and Data Mining,August 14-17. AAAI Press, Menlo Park, California. 1997, pp..287-290.