Expression and Interpretation in Spatial Hypertext

George Buchanan
Middlesex University
Bramley Road, London
N11 2NQ
+44 20 8411 5939
g.buchanan@mdx.ac.uk

ABSTRACT
In this paper I present Garnet, a spatial hypertext integrated to a digital library, and NoteTwig, a spatial hypertext for small-screen use. In both cases, the scope for expression in the hypertext system is limited by usability constraints – potential ambiguity in the first case, and navigation in the second.

The lessons learnt to date from these systems, and the problems which I am attempting to identify and address through them are described.

General Terms
Design, Experimentation, Human Factors.

Keywords
Spatial Hypertext, Digital Libraries, Small Screen, Usability.

1. INTRODUCTION
Spatial hypertext supports the organization of a set of documents into groups by the user. This organization is usually done using a visual drag-and-drop interface, manipulating objects that each represent one or more documents.

The first focus of my research is upon identifying topical patterns in spatial hypertext – i.e. the topic common to a visually related group of documents – and exploiting those topical patterns to support later work, either within the hypertext or elsewhere.

A second focus is on providing spatial hypertext in an integrated information seeking and structuring environment, particularly in how to integrate information seeking and structuring facilities, and providing spatial hypertext where display space is limited.

Together these twin foci result in three separate yet interlinked questions:

1. What properties of spatial hypertexts are key to the structuring performed by users?
2. Can relationships between visual representation or organization and topical themes be identified in spatial hypertexts?
3. Are the topical themes extracted from a spatial hypertext useful for supporting later work?

Item 1) is being studied through the use of different spatial hypertext metaphors. Much of my work on 2) has been focused upon identifying the use of position in two-dimensional spaces. The final question is being addressed through the application of information retrieval measures to human-organized spatial hypertexts.

2. APPLICATIONS OF SPATIAL HYPERTEXT
I now briefly describe two spatial hypertext systems in turn.

The first explores the use of spatial hypertext where display space is restricted – this may be in the course of integration into a larger tool that supports information work, or when used on a small display device such as a PalmPilot or PocketPC PDA.

The second is spatial hypertext when used as an interface to a digital library. In this case, I am using a spatial hypertext not only as a repository for the user's own work and organization of documents, but also as a filter through which unseen material can be interpreted.

2.1 Small-Display Spatial Hypertext
Spatial hypertext has traditionally been used on a standard, desktop-sized PC display. Commonly, the spatial hypertext editor has taken much of the available screen space when in use. However, my existing experience of small-screen information work, e.g. [8], and experiments towards creating an integrated information environment which contained a small document organization panel has led to work on a specifically small-screen spatial hypertext editor.

2.1.1 Adapting for the Small Screen
Spatial hypertext is not alone in using spatial workspaces for information tasks. For instance, concept maps are a popular tool for reasoning about intellectual material. A 2d workspace is used to separate and relate ideas and topics. Unlike spatial hypertexts, which are used to organize documents, concept maps are used to organize ideas. However, the use of space to express meaning and separation is common to both systems.

PicoMap [6] is a concept mapping tool for small screen use. PicoMap uses a traditional 2d concept map presentation, with adjustments to the interface that are intended to make interaction easier in a small display surface. Evaluation of PicoMap against a desktop equivalent indicated a number of problems emerged in the creation of maps on a small screen that were not present when a full desktop display was used. For instance, only half of the concept maps created on a handheld were considered readable, whereas all maps created with a large display were. Most of the problems were a direct consequence of the limited display space – e.g. a significant increase in overlapping causing occlusion.

As we demonstrated in [2], effective modification for small screens typically involves changes in interaction and presentation. The PicoMap system adopted our suggestions regarding interaction changes, but did not modify the present-ation. Our experience indicates that presentation is the more significant factor, and the PicoMap results support our hypothesis.

Considerable research has been done into how to view traditional geographical maps on small displays. Techniques such as halo-ing [2], fish-eye views [4] and speed-dependent zooming have been used to try and combine clear representation of a limited area with communication of the neighboring context. We have successfully used the principles introduced by Furnas [4] in supporting the searching and browsing of documents [2].

2.1.2 Spatial Hypertext for the Small Screen

I am adapting the outline-mode of presentation used in [2] to explore whether it gives adequate expressiveness for information structuring [See Fig. 1]. The outline mode of presentation reduces horizontal scrolling, which has proven particularly problematic on small display surfaces. This mode has proven an effective replacement for traditional hyperlinking on small-screens for cross-page and within-page navigation. Given the limited use of the second, horizontal dimension, outline modes are often described as being a "1.5d" representation.

The use of a highly atypical spatial hypertext form raises some more fundamental questions. The first of the three questions raised in the introduction is in many ways a re-assertion of one of the Seven Directions for Spatial Hypertext raised by Shipman in 2001 [11]. In that paper, Direction 5 is "The impact of Topology and Expressiveness on Usability and Usefulness". This question, or Shipman's direction, is particularly important when display space is limited.

My immediate focus has been on the use of position. Shipman and Marshall [10] have identified visual patterns such as piles and grids. However, in limited-capacity displays such as small screen environments, or where a spatial hypertext appears as within a panel of an integrated digital library system, positional cues are likely to be more limited than would appear to be the case in traditional spatial hypertexts. An outline-style presentation has very limited scope for visual organization – perhaps the textual appearance of an item, its level of indentation and its vertical position in the linear form of the outline. In the system I implemented – NoteTwig – horizontal bars, bullet points and other decorations were added to improve the scope for expressing the content of a document. Whether this paradigm is a usable or useful one for spatial hypertext at all is at present an open question.

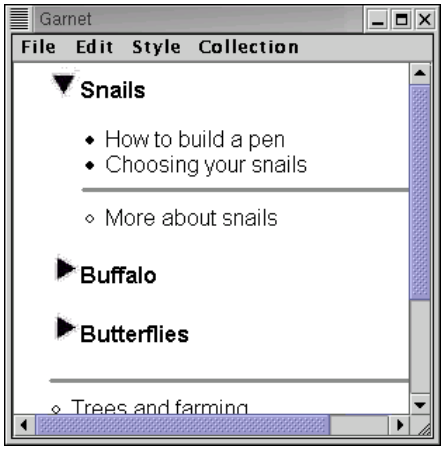

Fig. 1: A display of documents in NoteTwig, an outline-mode spatial hypertext for small-display contexts.

An immediate goal of this research is to identify ways of widening the degree of expressiveness available to the user, in order to mitigate against the loss of a fully 2d workspace. User studies will be undertaken to compare the effectiveness of traditional 2d workspaces against a 1.5d outline workspace in both full-screen and small-screen situations.

A benefit of the outline-style presentation is its consistency with our WebTwig interface for web searching and browsing on small-screen displays. Adopting the outline topics found within NoteTwig to provide a further context for user browsing may prove useful. Furthermore, consistent tools across different problem spaces have proven highly effective in many usability evaluations over the last two decades.

Alternative strategies include zooming. In that case, the visibility of individual documents becomes a potential problem. This difficulty is already present in traditional spatial hypertexts such as Pad++ [1], but is more immediately problematic in the case of small displays. Again, user studies are planned to compare the effectiveness of alternative navigation techniques where the display size is limited.

2.2 Integrating Spatial Hypertext and Information Seeking

In the Garnet Spatial Hypertext system, I have explored the integration of spatial hypertext and digital library facilities. This was primarily done to see how useful the implicit and explicit structuring performed by a user in a spatial hypertext was in supporting information seeking.

Visual interfaces to digital libraries have been influenced by spatial hypertext [3, 5, 7]. However, until recently the degree of genuine spatial hypertext support in each was extremely limited. Where a digital library and spatial hypertext are integrated, a number of approaches are possible. In Garnet, the visual library tools are presented inside the hypertext workspace. This means that objects in the workspace may be part of the library – e.g. a the results of a search. This opens some potential ambiguity as to what the user can or cannot do with a given object in the workspace. For instance, a user can clearly delete a document from their own part of the workspace, but not a book from the library. Color is frequently used as a cue to indicate the role of a document or other object in the workspace, but users can still choose identical or similar colors to those used by the library.

In addition to basic digital library access, Garnet supports the use of the user's workspace as a filter through which documents in a digital library can be selected. This requires the identification of visual patterns in the spatial hypertext, followed by a textual modeling of each identified visual cluster of documents. Finally, the textual representation of each document cluster is used to select documents in the digital library – be that from the whole library, or a subset such as the results of an individual search.

2.2.1 Interpreting Spatial Hypertext

Garnet attempts to interpret a spatial hypertext by identifying visual patterns of documents in the workspace and then creating textual summaries for each document group. The challenge here is the second question raised in the introduction. In addition, understanding the features that are most salient to user expression in a given form of hypertext will allow us to better focus our efforts on interpreting spatial hypertexts. Through Garnet, I am addressing the second question raised in the introduction to this

paper. Given the significant investment in effort that a mature spatial hypertext represents, any means by which a computer system can better extract context from it opens the door to providing a return on the time spent on information structuring.

The issue of a correlation between visual patterns and topical themes is problematic. Documents in close proximity may well have a common theme. However, any internal structure to that group is probably going to be more difficult to extract. Also, a common topic identified through textual analysis may not correlate well with the description that the author of the hypertext would use, nor with a description given by a reader of the hypertext.

Spatial parsing provides the basis for identifying visual patterns in spatial hypertext. However, as I wish to identify related topical patterns, traditional spatial parsing alone is not enough. Textual, or other semantic, analysis needs to be performed upon the identified set of documents within a visual structure to identify any common theme or characteristics.

In my initial work on a simple hypertext for digital library use, I have found that a simple spatial parser can be used in tandem with some simple textual analysis to identify topical themes. When compared to the topical themes identified by a textual clustering algorithm [13], the user-created groups prove of similar quality when evaluated using information retrieval analyses for textual consistency. These user-created groups also prove a useful filter on documents not found in the hypertext workspace.

2.2.2 *Experimental Experience*

The Garnet spatial hypertext system has been evaluated in user studies [2]. One concern with combining spatial hypertext and information seeking tools was that the resulting interface would be confusing. Thankfully, this proved not to be the case.

However, the participants in the experiments had little experience with Garnet, and may not have used its visual controls (e.g. color) sufficiently to provoke any potential problems. Further studies are planned, continuing with some existing participants, to identify if any long-term effects are observable.

In the introduction, I raised three questions. The section above briefly addressed the second question, and my user experiments gave an opportunity to discover whether any use of Garnet's "interpretation" of the spatial workspace was beneficial – the third and final question.

Features of Garnet that exploited the organization done by users in their workspace were positive received by participants. The "scatter" feature (see. Fig. 2) allows the user to select a search result set, and have documents from it placed next to documents with a similar topic on the main workspace area. In the user study I conducted on Garnet, the matching done by the scatter tool was rated highly except when the user used a single large miscellaneous pile or list.

However, this was not a direct evaluation of the textual model of the groups my users created. A future study needs to be undertaken to systematically compare the computer summarization with that suggested by the authors, and the acceptability of the computer summarization to the authors. Existing work in information retrieval suggests that the correlation tends to be weak, but the contexts in which these comparisons has been done to date do not directly correlate with spatial hypertext, and so differences may be found.

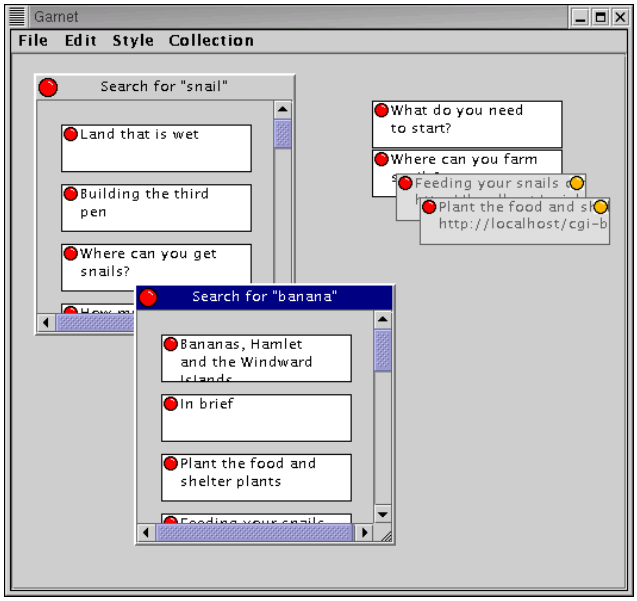

Fig. 2: Garnet's "Scatter" feature in use; the gray document labels seen on the right-hand side have been found to be similar to the two documents in white which they are positioned over. Both gray documents are originally found in the search result list in the centre of the screen.

2.2.3 *Further Exploiting User Organization*

In the experiments performed so far with Garnet, users have been accessing the digital library only through a workspace they themselves have created. Many questions remain unanswered.

For instance, the classic question of the intelligibility of hypertexts to readers other than the author [9] applies. Would one user find benefit in using the hypertext created by another?

The degree of consistency of users' organization of a common set of documents also remains to be explored. Garnet can produce representative texts for a group of documents, and again how much that set correlates with what the original author of the hypertext, or a reader, would suggest is unknown.

In each case, I am doubtful that we will find consistency. If consistency is not found, how to bridge the gap of comprehension will become a key question.

Existing experiments with Garnet have evaluated the use of the existing workspace structure for supporting post-query filtering of search result sets, and also the use of existing structures to initiate later queries. However, a further avenue for research is evaluating user's organization for the purposes of browsing. Substantial research has been done to compare the use of computed and manually assigned classification for browsing in digital libraries and other information repositories. The use of a user's own context, as encoded into their workspace organization, offers a third approach that can be compared to the other two.

2.2.4 *Computation over Spatial Hypertext*

Garnet represents a rare example of computation over spatial hypertext. Recent changes to the VKB system integrate it with a digital library system, and VKB also seems to use computation over the structures found in its workspaces. I am currently working to compare the differences between these two systems.

Computation over spatial hypertext needs the support of effective spatial parsers when informal structures may be found. More work needs to be undertaken to refine existing spatial parsers and to generalize the patterns that they can identify. Different metaphors of spatial hypertext – e.g. the use of a concept map paradigm – may result in new patterns not previously seen.

If hybrid systems that include spatial hypertext features increase in number, new challenges may also emerge. One simple example of this occurs in the Garnet spatial hypertext. Where a search result set is placed close to a pile of documents, it is not considered part of the document group. This decision was made as a result of observations during user studies we performed whilst piloting Garnet's design – users often making this placement with no intent to relate the search result and documents.

In addition to extracting information from user organization of documents, my experience with Garnet is that poor structuring of material can be determined using information retrieval metrics. Though users do deliberately use "miscellaneous" groups for documents with no natural place in the current organization, some users maintain large lists with no internal structure. A combination of clustering tools with the sort of assistance seen in VKB [12] may prove useful.

Similarly, document groups may readily have labels suggested – a common action seen in our user group. This can be achieved by exploiting the topical features identified through the textual analysis performed by Garnet.

2.2.5 What is Spatial Hypertext?

My experience in the digital library community is that spatial hypertext is quickly identified as a visualization of a library. I believe this to be in error, as visualizations are generated by computer, not as the informal workspace of a user. With Garnet, where the informal workspace is used to interpret other documents, the distinction between visual workspace and visual presentation becomes blurred, but my emphasis is still upon information structuring rather than visualization.

Others have described concept maps (also known as MindMaps™) as being Spatial Hypertext, though documents may not appear in the diagram at all. On the other hand, would the outline bookmarks of a web browser, very similar to our NoteTwig tool, be properly seen as a spatial hypertext? Though documents are certainly organized, is the scope of expression sufficient for a true spatial hypertext? Thus, the initial question from the introduction returns to the centre of attention. These three questions interact with each other in revealing and supportive ways.

Where, as with Garnet, the computer can play a role in placing material into the workspace, the boundary between visual workspace and visual presentation become more blurred and uncertain. Established visual techniques such as concept maps may provide a source for new forms of spatial hypertext. Embracing this diversity raises the question as to what is, in fact, spatial hypertext rather than some other visual tool.

Previous visual interfaces to digital libraries have borrowed from spatial hypertext, but often have limited representation of individual documents and a similarly restricted range of controls over the appearance of documents. The issue of representation and expression seen in the small-screen context re-emerge as spatial hypertext features are adopted by other systems, and vice-versa. For example, in DLITE [3], there is no representation of individual documents, except as members of a search result set.

2.2.6 Summary

In integrating spatial hypertext and digital libraries, a number of long-standing issues in spatial hypertext can be explored. Initial results are encouraging, in that many anticipated problems have proven so far to be less of an impediment that might have been expected.

However, in moving from hypertexts used only by a single user to workspaces used collaboratively, many more problems emerge. These now need to be addressed systematically, alongside the further exploration of single-user use.

3. CONCLUSION

In experimenting with spatial hypertexts in integrated environments, I have found that a limited scope for expression is not a critical impediment to effective information structuring. However, interpreting the results and then using them as a basis for further work by the author is a demanding area of research.

In facing the problems of integration, considerations in spatial parsing need to be widened (e.g. should objects of different types be treated similarly or differently). Limiting the scope for human expression can assist usability, as with NoteTwig, but there is a tension between ease-of-use and function.

Users show surprising consistency in their own organization of material, but whether the organization of one is usable for another is an open question. Exploring the possibilities for interpreting spatial hypertext opens the door to supporting authors of spatial hypertexts in their own tasks, and also as a possible gateway to map between users.

4. ACKNOWLEDGMENTS

Many thanks go to Rob Akscyn for his comments on NoteTwig. Also, to Matt Jones in his collaboration on small-screen research.

5. REFERENCES

[1] Bederson, B. B., Hollan, J. D., Perlin, K., Meyer, J., Bacon, D., and Furnas, G. W. "Pad++: A Zoomable Graphical Sketchpad for Exploring Alternate Interface Physics", Journal of Visual Languages and Computing, 7, 1196, 3-31.

[2] Buchanan, G., Blandford, A., Thimbleby, H., Jones, M., "Integrating Information seeking and structuring: exploring the role of spatial hypertexts in a Digital Library ", Proceedings of ACM Hypertext, 2004, in press.

[3] Cousins, S.B., Paepcke, A., Winograd, T., Bier, E.A., Pier, K.A.: "The Digital Library Integrated Task Environment (DLITE)". ACM Conference on Digital Libraries, pp.142-151, 1997.

[4] Furnas, G.W., "Generalized fisheye views", Procs. SIGCHI conference on Human factors in computing systems, ACM Press, 1986, 16-23.

[5] Furnas, G. W., and Rauch, S.: "Considerations for information environments and the NaviQue workspace". Proceedings of the Third ACM Conference on Digital Libraries (DL '98), ACM Press, New York, pp. 79-88, 1998.

[6] Luchini, K., Quintana, C., Soloway, E. “Pocket PiCoMap: a case study in designing and assessing a handheld concept mapping tool for learners”, Proc. Conference on Human factors in computing systems, ACM Press, 2003, 321-328

[7] Hendry, D. G., Harper, D. J., “An informal information-seeking environment”. Journal of the American Society for Information Science, 48(11), pp. 1036-1048, 1997.

[8] Jones, M., Buchanan, G., Mohd-Nasir, N., “An Evaluation of WebTwig”, Proc. International Symposium on Handheld & Ubiquitous Computing, 1999, 343–348.

[9] Marshall, C., “The Haunting Question of Intelligibility”, First International Workshop on Spatial Hypertext, ACM Hypertext 2001, Aarhus, Denmark, 2001. Online at: http://www.csdl.tamu.edu/~shipman/SpatialHypertext/SH1/marshall.pdf

[10] 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), pp. 346-353, 1995.

[11] Shipman, F.: “Seven Directions for Spatial Hypertext Research”, First International Workshop on Spatial Hypertext, ACM Hypertext Conference 2001, Aarhus, Denmark, 2001. Online at: http://www.csdl.tamu.edu/~shipman/SpatialHypertext/SH1/shipman.pdf

[12] Shipman, F., Moore, J. M., Maloor, P., Hsieh, H.W., Akkapeddi, R.: “Semantics happen: knowledge building in spatial hypertext”. Proceedings of ACM Conference on Hypertext, pp. 25-34, 2002.

[13] 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, AAAI Press, Menlo Park, California. pp. 287-290, 1997.

Columns on Last Page Should Be Made As Close As Possible to Equal Length