Beyond Location: Hypertext Workspaces and Non-Linear Views

Frank M. Shipman, III
Department of Computer Science
Texas A&M University
College Station, TX 77843-3112
shipman@cs.tamu.edu

Catherine C. Marshall
Xerox Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304
marshall@parc.xerox.com

Mark LeMere
Guidant Corporation
1360 O'Brian Drive
Menlo Park, CA 94025
markl@endovascular.com

ABSTRACT

With the growth of the Web as a public information resource, users need workspaces to support the collection, evaluation, organization, and annotation of the materials they retrieve. These analytic workspaces should be designed for both the casual and professional analyst, keeping in mind that different environments may be appropriate for each type of use. In this paper, we derive a set of requirements from observations and reports on the use of information workspaces, coupled with observations of people performing analytical tasks. These workspace requirements include: (1) support for performing multiple simultaneous tasks; (2) a variety of activity-based connections to information resources; (3) tailorable and manipulable reduced document representations; and (4) visualizations to help users manage screen space. We explore the trade-offs implied by these requirements using our implementation of multiple focus fisheye views as we have integrated them into the VIKI workspace.

KEYWORDS: information workspaces, analysis, interpretation, VIKI, fisheye views, reduced document representations, visualization, visual languages

INTRODUCTION

Hypertext, especially the World Wide Web, has taken an everyday role in the lives of many people. The Web is now used as a resource for both work and play. The home and office computer has changed from being a device for keeping track of finances, working with educational and specialized software, and playing games, to being one of the primary information sources available. The vision of the computer as an information appliance has come to pass.

Accordingly, there has been much work in the commercial world to help people locate desired information on the Web. Some sites provide categorization of Web content (e.g. Yahoo), while others provide full-text search facilities (e.g. Alta Vista and HotBot). Thus, while it still can be difficult to locate information, there is substantial support for this process.

But using the Web to perform real tasks is more than retrieving a list of documents. People need information to answer questions, solve problems, write reports, and as background for design, management, and general decision-making. In many cases, users have only begun their real task when a query has been refined to return a set of reasonable-looking documents. Once found, the documents may be scanned, read, interpreted, evaluated for validity, evaluated for utility, annotated, and filed for later use. While work on storage, transport, and retrieval of documents is necessary, there is also a need for interactive environments to use the information retrieved.

The Web and associated user environments must provide support for an ever expanding number and variety of users. Today, when a middle school student needs information on plants growing in her area for a homework assignment, she is likely to go to the home computer first. Of course, she may have difficulties locating information about the topic on the Internet. But, more likely, either by using a search engine like HotBot or Alta Vista or a topical index like Yahoo, she will uncover a large amount of potentially useful information. It is at this point that her task changes from one of location to one of interpretation and analysis.

As the student goes about collecting information for her school report, she may collect information from one or more digital libraries or information sources. She will create an "information space" one way or another through tools provided on the home computer (e.g. file system, folders, bookmarks). We (and others) have been working on visual information workspaces for analysis for a number of years [7]. Workspaces such as these provide an environment for the more casual user as well as the professional analyst for whom these tools were traditionally designed.

This paper discusses current practices surrounding the use of the Internet as an information resource and the concept of computational information workspaces. The next section recounts observations of information analysis in a variety of settings and how people use space. This discussion leads to a description of some of the requirements for developing a computational workspace to help both casual and professional analysts. We then explore one of these requirements in some depth – the need for tractable visualization techniques – by discussing our implementation of a multiple focus fisheye view for VIKI and how it has been integrated into the existing tool. Finally, we illustrate the results of applying the new view to existing VIKI workspaces.

INFORMATION TASKS AND SPATIAL HYPERTEXT

To develop an understanding of the current work practice surrounding the use of information, we have looked at how people use current computational visual workspaces, how people use the Web in their daily activities, and how visualizations emerge over time.

For some tasks, such as finding a phone number, the results of a football game, the price of a stock, or the capital of Kentucky, users usually perform a search (or go to some previously visited site directly) and find an answer to their questions. Such tasks are fairly common and on average take place relatively quickly. The standard bookmark list found in many Web browsers supports such tasks.

By contrast, research tasks such as writing a term paper, determining the state-of-the-art in some technology, or predicting which direction a stock (or a whole economy) will go, require in-depth analysis – collecting information from a variety of sources, skimming documents for the most useful information, and then developing an interpretation of those documents. Spatial hypertext goes well beyond Web browsers to support such tasks.

We have been developing VIKI, a spatial hypertext system to aid the analysis of information [6]. Spatial hypertext relies on proximity and other patterns of visual attributes to show relationships and linkages; most hypertext mechanisms, like those supported on the Web, rely on point-and-click traversal. For example, in the simplest case, two adjacent nodes in a spatial hypertext may be said to be linked. Other affordances of space are used to define different kinds of hypertextual structure (like hierarchies and composite nodes). VIKI allows users to construct a hierarchy of two-dimensional spaces, called collections, to organize symbols that represent information objects. Symbols may point to information objects that are stored in VIKI's own database, or they may point to external information like Web pages. Figure 1 shows a VIKI workspace with a combination of visual symbols and collections.

In VIKI, users collect, organize, and interpret information objects by using visual and spatial properties to signify classification schemes and relationships between objects. By organizing source materials into spatial structures, analysts create a visual interpretation that can easily evolve as their

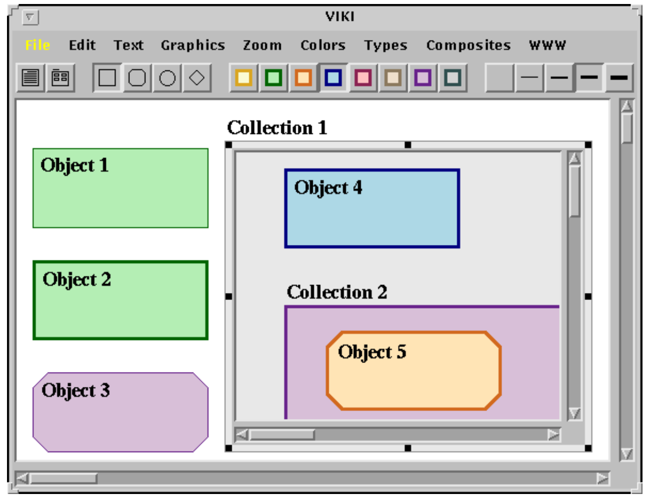

Figure 1. A VIKI workspace containing five visual symbols and two collections. Collections act as scrollable clipping windows onto the underlying information plane.

own understanding changes. For example, objects can quickly be moved from one category to another by changing their location or visual appearance.

Additionally, the categories themselves may change. The color of objects might initially be used to indicate a document's source while the object border width indicates some evaluation of relevance or importance. As the analysis progresses, the visual encoding might change so that object color identifies the topic of the information rather than the source. By enabling transient and partial interpretations and by aiding in the manipulation and expression of these interpretations, spatial analysis environments like VIKI support the evolutionary nature of the analysis process. The visual and spatial representation allows the user to non-verbally express the ambiguity and indecision that is natural during the analysis process.

Besides supporting evolving non-verbal interpretations, there are other benefits offered by the visual nature of the VIKI workspace. For long-term and infrequently accessed information spaces, the application of users' visual and spatial orientation abilities for relocating documents and remembering the state of a prior task is crucial. For collaborative tasks, a shared visual workspace can be used to coordinate and communicate. Other than the flexible visual representations of VIKI, the benefits described above apply to all visual workspaces. The following section will describe some observations of VIKI use and how they have led to new functionality to better support common work practices.

VIKI and the Web

While VIKI was not originally designed to work with the Web, we have been improving the communication between VIKI and Web applications, such as Netscape's Navigator, so VIKI may be used to collect, organize, and annotate Web-based materials.

There are three main connections between VIKI and the Web. First, as we have already mentioned, information objects in VIKI may point to information on the Web. When a user double-clicks on a visual symbol to open such an object, a Web browser displays the remote materials. This allows VIKI to work as a Web workspace where information of all media types can be organized within VIKI's hierarchy of two-dimensional spaces. Figure 2 shows a VIKI workspace containing symbols pointing to information about Macedonia that is available on the Web.

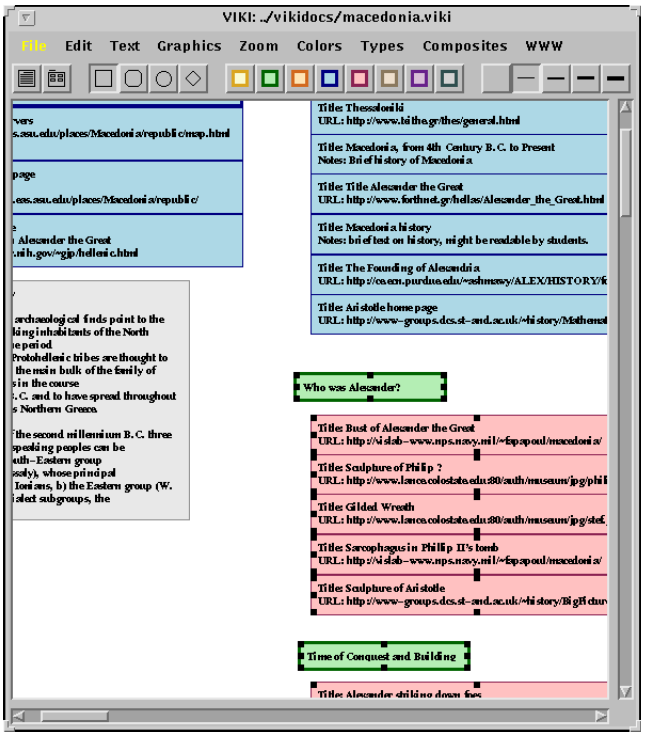

Figure 2. Information from the Web being collected and organized into paths in VIKI.

The second connection between VIKI and the Web is the interface between VIKI and specific Web search engines. Users may submit a query to be evaluated on a remote search engine. Each document that is returned by the engine generates a new visual symbol within the current VIKI information space. Users then may color code and spatially organize the results of their queries as they are evaluating the utility of the documents returned by the search engine.

The final VIKI-Web connection enables users to generate presentations of Web-based information within VIKI. The collection, organization, and annotation of information can eventually lead to a meta-document (a document composed of other documents) that presents one's findings. The internal presentation viewer in VIKI includes the ability to create annotated linear paths combining materials stored in VIKI with materials from the Web [16]. The selected items in the lower right of Figure 2 include the title, URL, and annotation of each element to be in the path.

INFORMATION WORK AND REQUIREMENTS

Most tasks are complex; people seldom progress through a particular activity in a straightforward linear fashion – consulting one document may bring to mind several others, which may in turn prompt a retrieval task. We have observed people using multiple Web browser windows as they search for information that they are collecting in the service of very different activities, a mode of action which may be brought on by exogenous factors like Internet response delays and network latencies. Given these exogenous factors, switching among concurrent tasks makes people feel more productive. Furthermore, people self-interrupt; they hit a stopping point on one task, and simply turn to another, hoping that when they return to the first task, it will be either more tractable or less boring (or both).

Do these observations necessarily imply that people are working with large numbers of documents at once? A recent survey of bookmark use shows that users amass an almost linearly increasing number of bookmarks over time, but only a few of the bookmarks are actually being used [17]. Hence we can envision hypertext workspaces that at any given time contain large numbers of inactive objects, and much smaller numbers of active ones.

When we design hypertext workspaces, we must consider just this types of task variability. People collect or amass materials they need for a particular activity; these materials may be related only through their function in the situation at hand. They work on many things at once, and mix active and inactive materials. In VIKI's data model, collections allow a user to gather materials in a "place", so that they can be used together. We consider this localization an essential characteristic of workspaces, and a good basis for the non-linear visualizations that we describe later in this paper. In the following sections we discuss requirements for how individual documents or lexia are represented in workspaces, and how these workspaces may then be viewed in a manner which supports existing practice.

Reduced Document Representations

In the physical world, documents don't simply function as information containers; in different situations, the same document may act in many roles [4]. For example, documents remind people of tasks they have yet to finish; passed hand-to-hand, they are tokens of a job that has been completed; and marked up and tossed in piles on a person's desk, they are records of partial interpretations of content (for example, that the document is important, or that the document's topic is signal analysis). It is important for workspaces to give people similar access to representations of individual documents, coupled with the ability to manipulate the representations in a fairly flexible manner so the document in the workspace can continue to fulfill multiple roles.

How might documents be represented in spatial hypertext? It is important to remember that the representations must be reduced, so that people can display and use multiple documents at a time, and that they must be in a form that suits the task. Reduced document representations include metadata-based displays, image thumbnails, and computed content derivatives (like summaries). Each is effective, and suggests a different way of working with documents.

VIKI represents individual text objects as tailorable views onto the object's metadata, since some tasks may be best served by working from, for example, an article's title, and others may be better served by working from its publication date or author (or both) [6]; this kind of reduced representation takes advantage of the idea that many information sources are pre-structured. Other systems, like the Integrator [3] and XLibris [15], take their cue from the way people recognize document genres (and individual documents or document parts) from their visual structure; these systems use thumbnails to make a particular reduced representation easily distinguishable from others. Systems may also use object summaries, or other computed content derivatives, as the basis for document representations. A few systems, like Pad++ [10] represent text objects in their entirety, and give the user the ability to zoom so they can control the level of reduction. In short, the reduced representations provide a variety of affordances for working with the documents in the workspace.

In any case, what we need to keep in mind about this aspect of workspaces is that the reduced representation must maintain some sort of integrity over various kinds of systematic transformations. In VIKI's case, it is particularly important that the focal objects' metadata maintain legibility – that the text does not become too small to read.

Nonlinear Workspace Visualizations

Our analyses of work practice have shown us that VIKI's users often attend to multiple documents at once; the tasks they engage in are characterized by extensive reading (see [5]), and may require that text objects be compared or integrated. When these text objects are proximate, the original spatial metaphor works well. However, when they are distributed in different portions of the workspace, it is difficult to maintain a multiple node focus without losing the overall spatial context for the items. A simple zoom may allow the user to see the desired objects simultaneously, but they often become so tiny that they are indistinguishable from their neighbors, and desirable affordances of the reduced representation – in this case, the metadata values – are swept from view. Likewise, duplicate references to objects may allow them to be placed in multiple spatial contexts; yet, for a variety of reasons, that is not always the strategy that people use [8].

We have been exploring ways to allow users to maintain attention on several documents at once, keeping in mind the affordances of VIKI's reduced document representation, as well as the kind of hypertext structures people build using VIKI's spatial metaphor. In particular, we have been exploring variations on a multiple focus fisheye transformation, starting from the transformation described by Sarkar and Brown [13]. In the following subsection, we describe our basic implementation of a multiple focus fisheye transformation for VIKI.

AN IMPLEMENTATION OF MULTIPLE FOCUS FISHEYE

We have extended Sarkar and Brown's basic fisheye transformation so users can choose multiple focus points for constrained-space fisheye views. The text objects at the focus points retain the largest possible proportions, and the active workspace containing them acts as the constraining area. The remaining objects in the workspace are all rendered (and made visible) in this constrained area. This strategy is in marked contrast to the fisheye view described in [14], in which collections are the focus of the fisheye transformation; collection size is changed by the transformation.

To illustrate how the algorithm works, we use a simple example. Let's say that the user selects focal objects at points A, B, and C (for this example, we won't draw the other objects in the space). The algorithm first determines linear equations for the perpendicular bisectors of the lines connecting each pair of focal points. In our example, the line segments connecting the foci are AB, BC, and AC. The intersection of the bisectors with the edge of the collection (which forms the bounding rectangle in Figure 3) is used to calculate the bisectors' end points. In our example, this gives us line segments AB', BC', and AC', as shown in Figure 3.

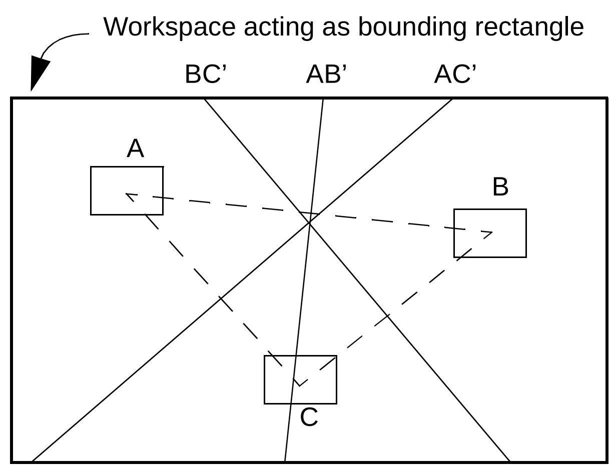

Figure 3. Calculating perpendicular bisectors AB', AC', and BC' of line segments AB, AC, and BC. A, B, and C are VIKI objects. Non-focal objects have been omitted for clarity.

The intersections of the perpendicular bisectors can now be used to produce "regions" around each of the focal points. Figure 4 shows the regions that arise in our example. The size and position of a given object in the fisheye view is computed according to the closest focal object and the corresponding region's boundary using the radial Sarkar and Brown transformation function.

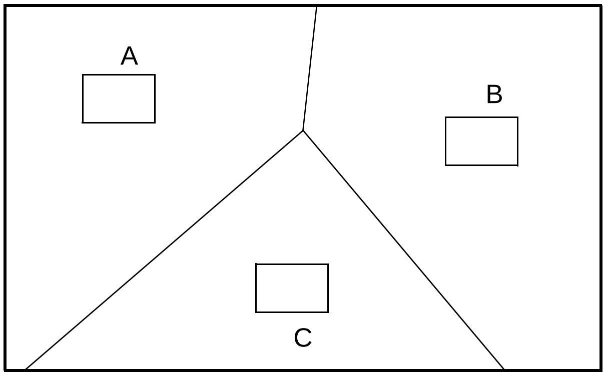

Figure 4. The regions that result from the application of the multi-focus fisheye algorithm

How does this transformation look with objects in addition to the focal objects? Figure 5 shows a three focus fisheye on a regular array of rectangular objects. Again, collection size is used to constrain the display region, and selected focal points determine which objects remain close to their original dimensions (and hence, greater readability of the textual metadata displayed on them).

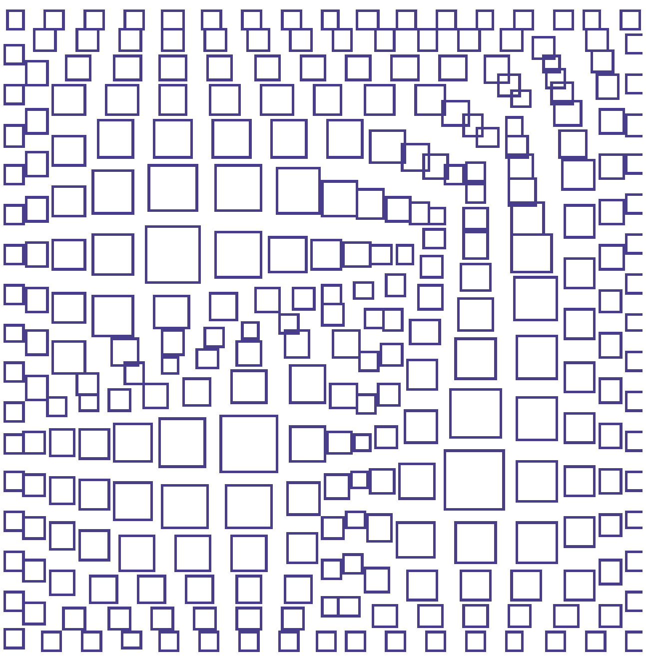

Figure 5. The basic transformation applied to a regular array of square objects, given three focus points.

In the development of a basic fisheye capability, we kept several constraints in mind. The first was that the user may be using more than one collection at a time; thus it was important to keep collection size constant. Secondly, there may be multiple focal objects, but the number would probably be kept within reasonable limits, since people work with but so many documents at a time. Finally, the focal objects must maintain their readability. These three constraints were satisfied.

But as we have noted before, much of the expressive power of spatial hypertext is derived from the ability to informally define and use spatial structures. What does this transformation do to the spatial structures that the user has laid out? These structures are presumably central to the intelligibility of the hypertext.

We considered two competing implementations of the fisheye algorithm, one that uses polar coordinates, and one that uses rectangular coordinates. They suggested an important trade-off in designing nonlinear views: preserving a workspace's visual structure v. optimizing display space. Figures 6a and 6b illustrate this tension by showing a polar coordinate transformation with three focal objects (6a), and a rectangular coordinate transformation with three focal objects (6b). Note that the first, polar, makes better use of display space and the second, rectangular, preserves the sense of the spatial structures. Optimizing display space is, of course, the fundamental motivation behind non-linear views. The importance of human-produced visual structure is a central premise behind spatial hypertext. How can this tension be resolved? In the following section, we explore ways to handle this tension and ways of allowing the user to resolve it.

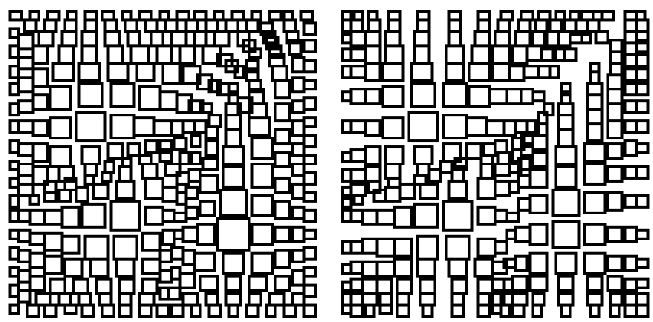

(a) (b)

Figure 6. (a) shows polar coordinates, three foci; (b) shows rectangular coordinates, three foci.

NONLINEAR VIEWS IN VIKI

As we described earlier, VIKI provides a hierarchy of two-dimensional spaces called collections for gathering and arranging information symbols. Observations that Web users frequently engage in multiple simultaneous information tasks and that each individual task may rely on multiple parts of the information space motivated us to add nonlinear views to VIKI. To integrate this capability, in addition to understanding common work practices, we also needed to evaluate the trade-offs between different potential interfaces to the non-linear view mechanism. In this section we discuss three main trade-offs we addressed in the integration: (1) how much control over the view transformation to offer users; (2) which coordinate system to use for the transformation, space-optimizing polar coordinates or structure-preserving rectangular coordinates; and (3) how to let users choose an acceptable distortion level.

Automation vs. User Control

One central trade-off in integrating non-linear views is how to balance ease of use with user control over the particular transformations. The most flexible interface would be one that allows a user to modify all of the coefficients in the mathematical transforms independently. At the opposite extreme, the system could generate one fisheye view based on a user's current situation.

While it is appealing to give users maximum control, this either requires that they learn the mathematics of the transform (so they can edit the numerical values in a dialog window) or that we develop a direct manipulation interface for altering the transform. On the other hand, having the system select the transform based on the current information space, foci, and screen space streamlines the interface into a single button.

As we will show, the final design ended up being a hybrid, retaining the interface simplicity of the automated approach – using only two buttons – while giving the user an option of six non-linear views for the current information space, foci, and screen space. The two buttons allow the selection of the type of transform (rectangular vs. polar) and one of four levels of distortion associated with the transform (none, low, medium, high.)

To enter a nonlinear view, users select the distortion-level icon and (if no symbols or collections have been selected) the current center of the visible space becomes the single focus for the transform. Also, VIKI users can select one or more symbols or collections before applying the new view, each of which will become a focus for the non-linear view.

Rectangular vs. Polar Nonlinear Views

A comparison between the rectangular and polar nonlinear transformations included in VIKI identifies some of the advantages and disadvantages of each. The following images show the differences between the two transforms. Figures 5 and 6 show the difference between the rectangular and polar transformation of a simple layout of visual symbols on a grid so as to provide the best view of how the transforms compare.

The polar and rectangular transforms have a single point of focus near the middle of the layout. Figure 7 shows the results of the polar transformation. Notice that the alignment of symbols is warped around the center of focus. Symbols which are vertically or horizontally aligned in a linear view are not aligned in this view. By contrast, Figure 8 shows the rectangular transform. In this case the symbol alignment is retained, making implied relationships between symbols easier to perceive. The disadvantage of the rectangular view is that there is more white space, reducing the amount of space devoted to showing symbols. In this coordinate system, the white space is used to keep symbols in vertical and horizontal alignment. The large horizontal gaps between symbols near the bottom in Figure 8 are the result of this use of extra space to retain alignment between symbols.

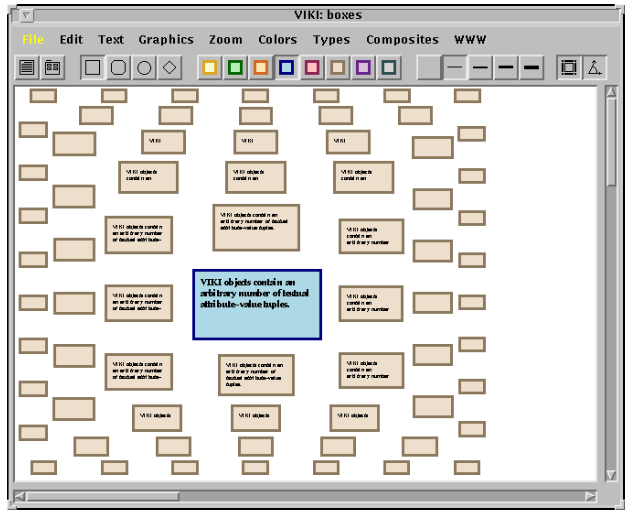

Figure 7. The polar transformation with one focus point applied to a regular array of visual symbols in VIKI.

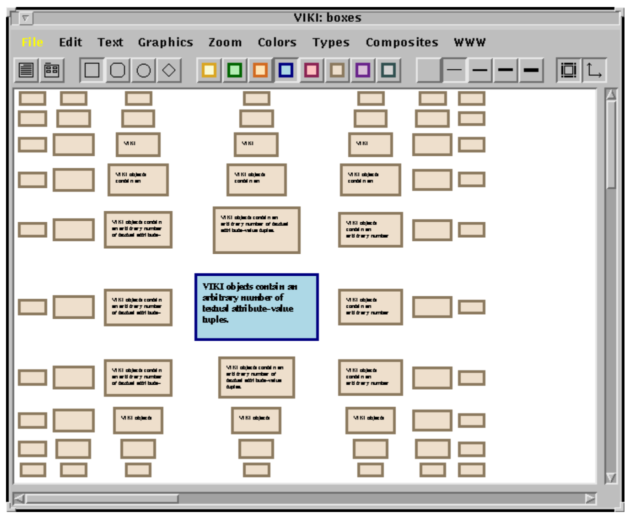

Figure 8. The rectangular transformation applied to a regular array of visual symbols in VIKI.

This might indicate that when perceiving alignment relationships is important, rectangular transforms are best and when having larger visual symbols is important (larger visual symbols can display more metadata about the information object), polar transforms are best. This is true for a single focus, but the issue is complicated when there are multiple foci.

Because the layout is divided into separate transform areas surrounding each focus point, there are discontinuities in the alignment of symbols at the borders between these areas. In the case of the polar transformation, the division can be seen by the diagonal regions of small symbols in Figure 6a. In the case of the rectangular transformation the borders include overlapping symbols and more extreme discontinuities in alignment (as seen in Figure 6b.) Alignment relationships

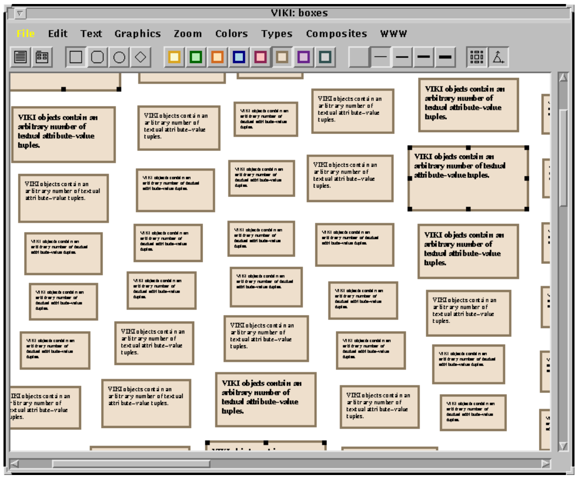

(a)

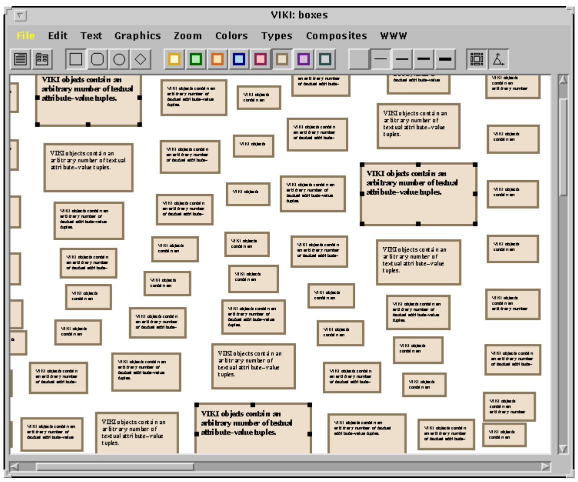

(b)

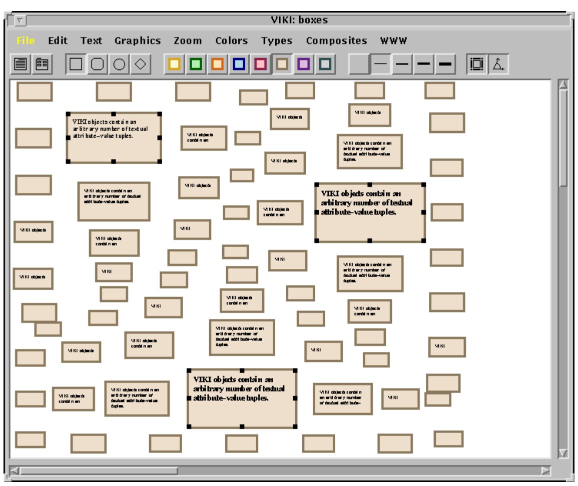

(c)

Figure 9. These three screendumps show results of increasing distortion levels using the polar transformation. This time there are three objects selected as focus points.

within a transform region are easier to perceive with the rectangular transform but relationships across transform regions are more visible with the polar transformation.

Because it may be difficult to determine whether a polar or rectangular transform is best suited for the current situation, the user can easily switch between the two tranforms.

Selecting a Distortion Level

In the prior examples, there was only one distortion level shown and it was computed to bring the entire display into the visible area. This approach allows users to perceive the high-level context of the selected or centered symbols, but the local context, such as the text in symbols near the selected ones, frequently is illegible due to the high degree of distortion required to fit all the visual symbols into the available screen space.

By selecting the distortion icon, the second icon from the right on the icon bar, the display cycles through the normal view and three levels of distortion. Figure 9 shows the gridded layout of symbols from Figures 7 and 8 in these three increasingly distorted views with three information symbols selected as foci.

NONLINEAR VIEWS OF HIERARCHIC COLLECTIONS

Besides investigating a variety of nonlinear views, we have also worked toward the use of different views in each of the collections within a VIKI space. As an example, we will use the layout of symbols from a collaborative VIKI space used by a group of researchers to collect information relevant to their group's research. Figure 10 shows the top-level collection for this information space.

Each of the five rectangular areas seen in Figure 10 is a subcollection that the user can navigate into to see more. These subcollections have labels to document how they are used. The subcollection in the lower-left shows that there are

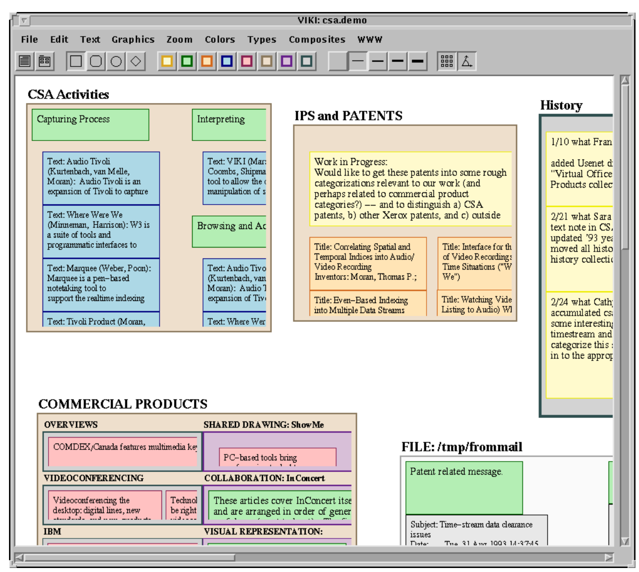

Figure 10. Top-level collection in VIKI workspace.

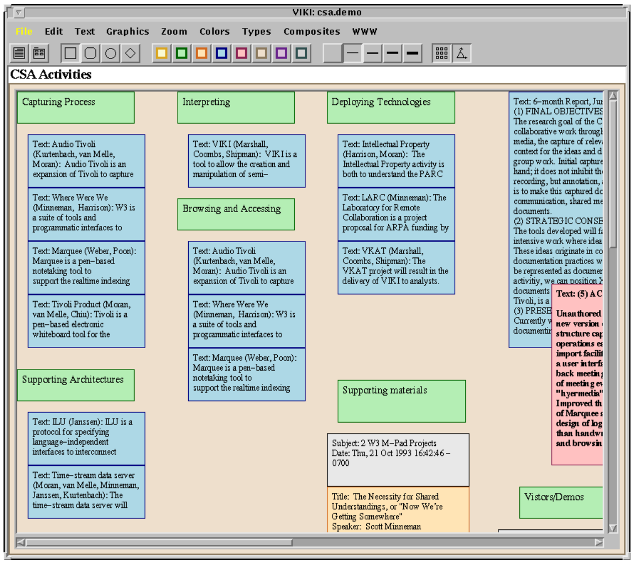

Figure 11. View after user navigates into "CSA Activities" subcollection.

further levels of subcollections within it. Figure 11 shows the user's view after navigating into the "CSA Activities" space in the upper-left corner of Figure 10.

Once in this subcollection the user may change to a non-linear view within this collection as is done with the top-level collection. Figure 12 shows such a view – the rectangles at the bottom of the figure are subcollections within this collection.

The view of the information within this collection will remain non-linear when the user navigates out of the subcollection. Figure 13 shows the view of the top-level collection after this change. Figure 14 shows the screen after the "IPs and Patents" subcollection is also non-linear and after the visible are of this subcollection has been adjusted to display the area near the point of focus.

The work on nonlinear views in VIKI provides users detailed control over the application and characteristics of the transformation while retaining ease-of-use.

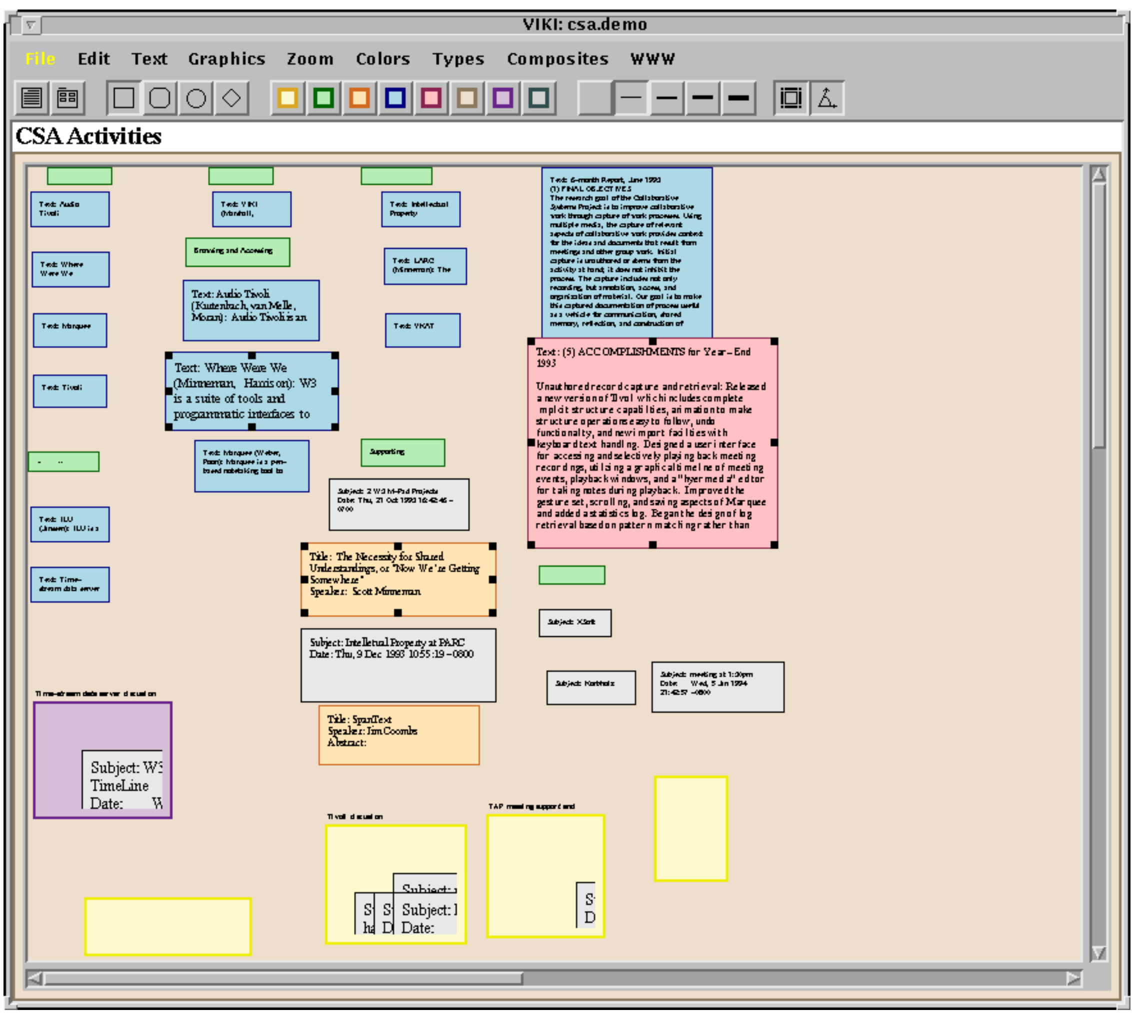

Figure 12. Non-linear view of is applied to subcollection with three foci.

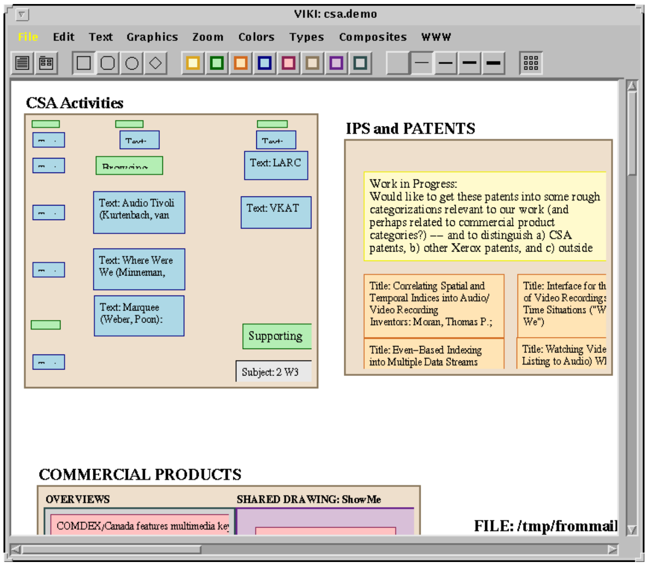

Figure 13. "CSA Activities" subcollection view is non-linear but others are all linear.

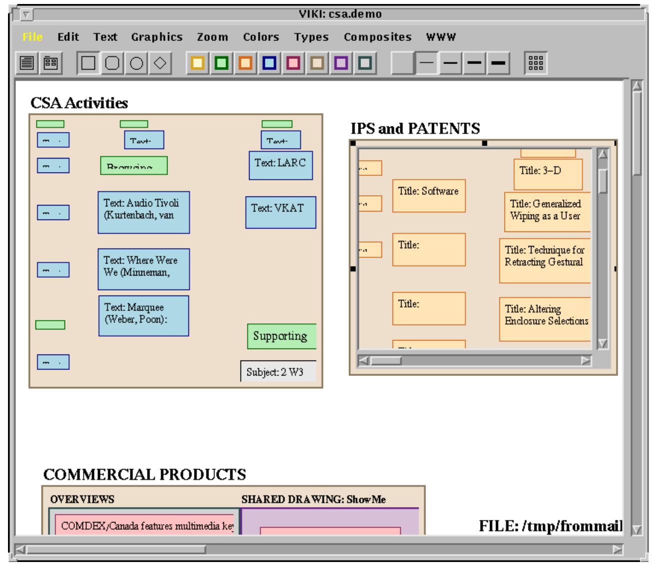

Figure 14. Two collections are now non-linear.

CONCLUSIONS

Gathering on-line materials and making sense of them is becoming more of an everyday activity; readers are becoming gatherers [12]. Information workspaces can provide the appropriate affordances for the range of gathering-related activities – searching, triage, interpretation, and presentation. We have identified basic requirements for making workspaces usable in this manner: the ability to collect materials directly from a source (like the Web); the ability to manipulate an appropriate form of reduced document representation; and the ability to focus on multiple portions of the workspace at once.

We have explored these requirements using VIKI with a live retrieval connection to the Web. In particular, we have implemented multiple focus fisheye views that display metadata for selected information objects. The integration of multiple focus fisheye views into VIKI illustrates a set of general trade-offs. First, from a user perspective, these views are unfamiliar. How much interaction should a user have to guide the algorithm parameter selection? Second, the choice of coordinate systems has an important bearing on which aspects of the layout are optimized or preserved. Is it more import to use the space well, or is it more important to preserve the spatial structures? Finally, fisheye views introduce distortion. Are there appropriate ways for the user to control the distortion while still rendering a large part of the area of interest? We examine how this implementation plays out in existing VIKI spaces.

The effectiveness of information workspaces is uncontroversial – the concept is embodied in products (for example, Web Squirrel [2]) and in new research (for example, see [11]). Previous work (see [1] or [9] for example), as well as our own, has explored how fisheye views might be used to make the kinds of large information spaces we expect to encounter in these workspaces more tractable. What remains for future work, however, is significant use experience to better understand the types of trade-offs we discuss in this paper.

REFERENCES

[1] L. Bartram, A. Ho, J. Dill, and F. Henigman, "The continuous zoom: a constrained fisheye technique for viewing and navigating large information spaces", *Proceedings of UIST'95*, New York: ACM (November 1995) pp. 207-215.

[2] M. Bernstein, "Visualize Cyberspace", http://www.eastgate.com/squirrel/FAQ.html.

[3] J. Blomberg, L. Suchman, and R. Trigg, "Reflections on a Work-Oriented Design Project." *Human-Computer Interaction*, 11(3) (1996), pp. 237-265.

[4] R. Harper, and A. Sellen, "Collaborative tools and the practicalities of professional work at the IMF", *CHI '95 Conference Proceedings*, New York: ACM (May 1995), pp. 122-129.

[5] D. Levy, "I read the news today, oh boy: Reading and attention in digital libraries", *Proceedings of Digital Libraries '97*, New York: ACM (1997), 202-211.

[6] C. Marshall, F. Shipman, and J. Coombs, "VIKI: Spatial Hypertext Supporting Emergent Structure", *Proc. of European Conference of Hypermedia Technologies (ECHT '94)*, New York: ACM (1994), pp. 13-23.

[7] C. Marshall and F. Shipman, "Spatial Hypertext: Designing for Change", *Communications of the ACM*, 38 (8) (August 1995), pp. 88-97.

[8] C. Marshall and F. Shipman, “Effects of Hypertext Technology on the Practice of Information Triage”, *Proceedings of Hypertext ‘97*, New York: ACM (1997), pp. 167-176.

[9] M. Noik, “Exploring large hyperdocuments: fisheye views of nested networks”, *Proceedings of Hypertext ‘93*, New York: ACM (November 1993), pp. 192-201.

[10] K. Perlin, and D. Fox, “An Alternative Approach to the Computer Interface”, *Proceedings of Multimedia ‘93*, New York: ACM (August 1993), pp. 57-64.

[11] G. Robertson, M. Czerwinski, K. Larson, D. Robbins, D. Thiel, and M. van Dantzich, “Data Mountain: Using Spatial Memory for Document Management”, *Proceedings of UIST ‘98*, New York: ACM (November 1998), pp. 153-162.

[12] J. Rosenberg, “The Structure of Hypertext Activity”, *Proceedings of HT’96*, New York: ACM (March, 1996), pp. 22-30.

[13] M. Sarkar, and M. Brown, “Graphical fisheye views.” *Communications of the ACM*, 37 (12) (December 1994), pp. 73-83.

[14] D. Schaffer, Z. Zuo, S. Greenberg, L. Bartram, J. Dill, S. Dubs, and M. Roseman, “Navigating hierarchically clustered networks through fisheye and full-zoom methods”, *ACM Transactions on Computer-Human Interaction*, 3 (2) (June 1996), pp. 162-188.

[15] B. Schilit, G. Golovchinsky, and M. Price, “Beyond Paper: Supporting Active Reading with Free Form Digital Ink Annotations”, *CHI ‘98 Conference Proceedings*, New York: ACM (1998), pp. 249-256.

[16] F. Shipman, R. Furuta, and C. Marshall, “Generating Web-Based Presentations in Spatial Hypertext”, *Proceedings of Intelligent User Interfaces ‘97*, New York: ACM (1997), pp. 71-78.

[17] C. Thomas, and G. Fischer, “Using Agents to Personalize the Web”, *Proceedings of Intelligent User Interfaces ‘97*, New York: ACM (1997), pp. 53-60.