appeared in the 1995 *Communications of the ACM* special issue on Hypermedia Applications

Spatial Hypertext: Designing for Change

Catherine C. Marshall and Frank M. Shipman III

Introduction

Hypertext[1], in its most general sense, allows *content* to appear in different *contexts*. The immediate setting in which readers encounter a specific segment of material then changes from reading to reading or from reader to reader. Authors collect and structure materials to reflect their own understanding or in anticipation of readers' possible interests, needs, or ability to comprehend the substrate of interrelated content.

This powerful underlying concept is usually realized in both research and practical efforts within a node-link model: nodes are the holders of content, and links are the means by which the content is given context. In this model, links are closely associated with navigation and mechanisms for traversal; they are a way to move from node to node, to keep readers focused on the current node or document, until they decide to move on to the next. We refer to this style of hypertext as *document-centered hypertext*.

As systems and applications designers have gained more practical experience with hypertext, models of structure have grown more sophisticated, more expressive, and in many cases, more flexible than conventional node-link models. To address the needs of specialized applications, some models diverge entirely from standard notions of links; for example, Parunak introduces set-based hypertext for applications that involve taxonomic reasoning [21]. The expressiveness of hypertext models has been extended by adding types to nodes and links [5][20] and by using structures in which n-place relations replace typed links [15]. Finally, for the sake of flexibility, in some applications links are left implicit, computed dynamically by using text analysis methods or other heuristics for determining regular interconnections [2][6]. Thus we see systems and applications designers broadening the field, in some cases by theory, but also through practical experience. Spatial hypertext has grown out of just such an intertwining of theory and experience-based extensions to hypertext concepts.

Spatial hypertext has arisen through our experiences with applications that explore

1. We use the term hypertext broadly, to cover both textual and multimedia content.

alternative structures for content and applications in which the domain structure is not well understood at the outset or changes during the course of a task. Many of these applications involve the collection, comprehension, and interpretation of diverse materials; they are information-intensive activities, like analysis, design, or evaluation, and are often collaborative efforts. *Spatial hypertext is most appropriate when there is no distinction between readers and writers, and more prescriptive design methods might hamper exploratory structuring.*

Spatial hypertext has its origins in *browser-based* approaches in which the emerging hypertext network is portrayed graphically, in an overview, to promote coherence (see Thüring et al., this issue); authors and readers interact with the hypertext using a map of its structure. In early conceptions of browser-based hypertext, boxes or other icons symbolized nodes and lines represented the links among them; hypertext browsers were oriented toward presenting a graph representation of a network. NoteCards [9] and gIBIS [5] are two examples of hypertext systems that used such a spatial map. But gIBIS differed from NoteCards in one crucial way: in NoteCards, the browser was used mainly as a means of *visualizing* existing network structure[1]; in gIBIS, the browser was the primary means of *interacting* with the emerging network as well as presenting it.

Extensions to the node-link model, such as types and complex structures, have found their way into browser-based approaches. By watching people use these extensions, our sense of "what matters" in spatial hypertext evolved. The ability to *create and move nodes freely* matters; the ability to *express relationships by spatial proximity and visual cues* matters. If we remove the explicit links from a browser, it can become a dynamic canvas for interaction. Nodes may appear in different contexts through multiple spatial references to the same underlying content; authors may use "any unit of text as a new element in an expanding vocabulary of signs" [4].

Thus, spatial hypertext is not only a means of presenting readers and authors with visualizations of existing structure; it is also a way to take advantage of human perceptual abilities in hypertext navigation, and to provide users with a fairly intuitive medium through which they may express new structure and manipulate existing structure.

1. Although NoteCards allowed users to construct hypertext through browser interaction, most users chose other, more document-centered means of making links.

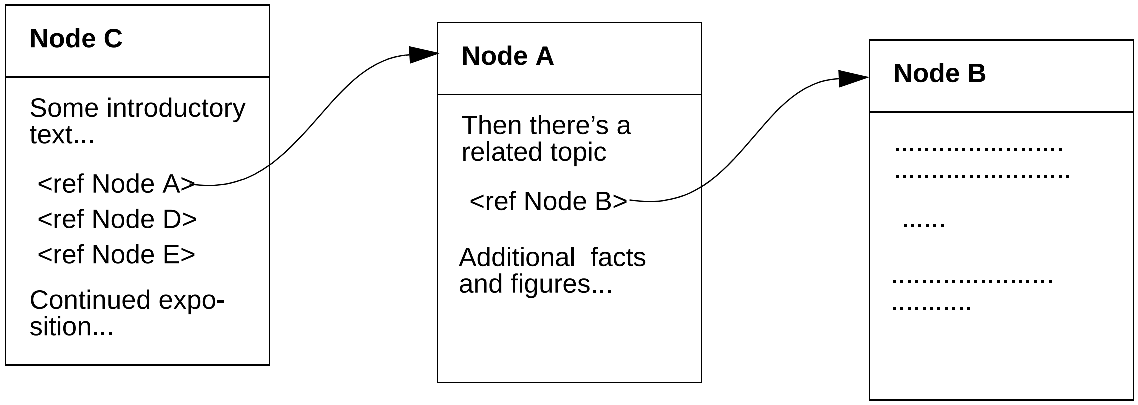

Figure 1 a. Document-centered hypertext

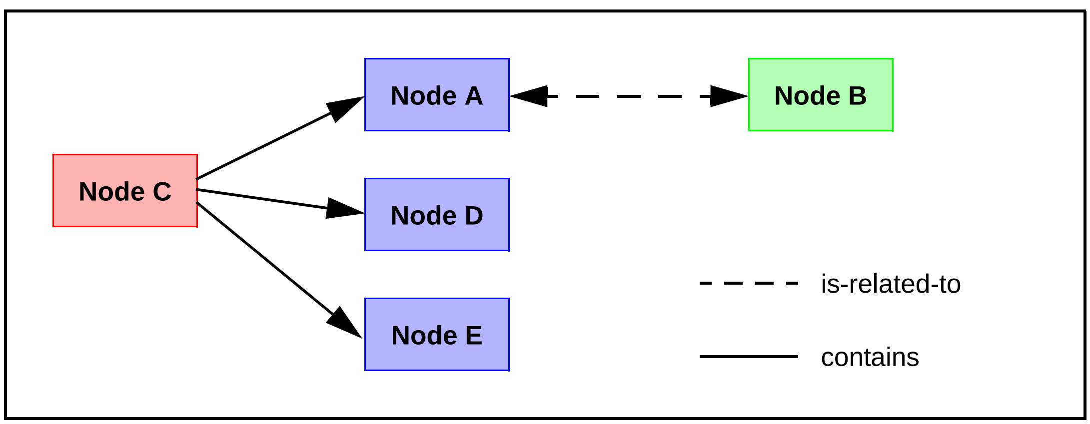

Figure 1 b. Browser-based hypertext

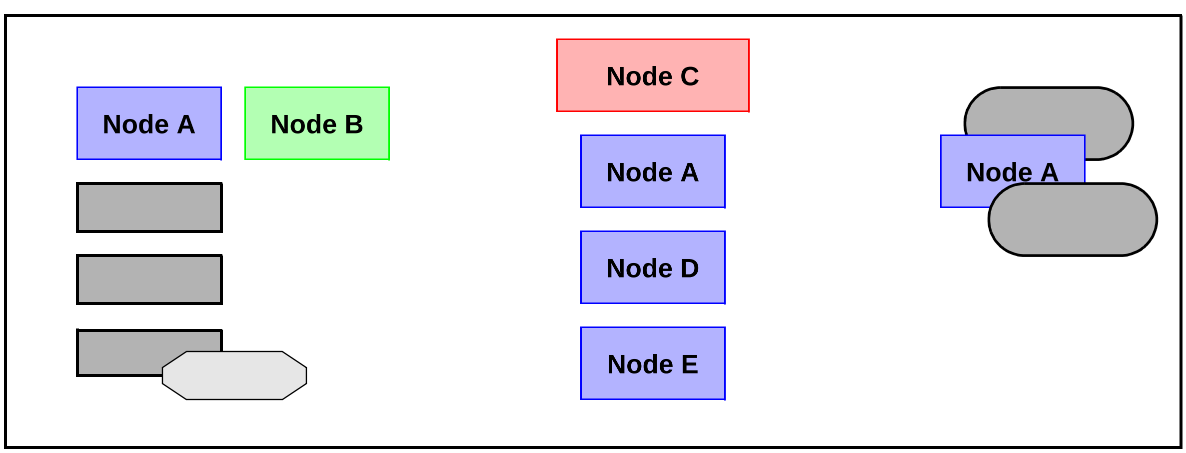

Figure 1c. Spatial hypertext

Figure 1. Transitions from document-centered hypertext to browser-based hypertext to spatial hypertext

Figure 1 illustrates how hypertextual structures can be realized in a spatial setting, by showing the transition from a document-centered hypertext (Figure 1a) to a browser-based hypertext (Figure 1b) to a spatial hypertext (Figure 1c). First, let us look at Figure 1a, and imagine Nodes C and A to be documents displayed as one might find in popular interfaces such as NCSA's Mosaic; clicking on the anchor (or, more precisely, the marker for the anchor) embedded in Node C causes the document corresponding to Node A to be displayed in the window. Node A may replace Node C in the viewer, as it does in Mosaic's default behavior, or open a new window, as it would in many multi-window hypertext systems, but in either case, the links are used primarily as a means of traversal.

Figure 1b shows the same hypertext portrayed in a spatial browser. Nodes A, B, C, D, and E are represented by *visual symbols* -- boxes, in this case -- that refer to the underlying documents. We have added types to Figure 1b; node types are represented by box colors and link types by line dashing. Arrowheads indicate link directionality. Nodes A, D, and E are all contained by Node C, and Node B is related to Node A. Figure 1b is fairly typical of the style of hypertext network display presented by browser-based hypertext (see Figure 2 in Thüring et al., this issue and Figure 4 in Nanard and Nanard, this issue).

Figure 1c extends and generalizes Figure 1b's browser to a simple spatial hypertext. Once again, we refer to content of underlying nodes by visual symbols, applying the color-coded types from Figure 1b. First we see that Node A's visual symbol is near Node B's visual symbol, so Node A is very probably related to Node B. From the characteristics of their visual symbols, Nodes A, D, and E are understood to belong to the same set, contained in Node C, whose symbol is at the head of the list; the set relationship is implicitly conveyed through proximity, alignment, and homogeneity of type. To illustrate one final important property of spatial hypertext, we have populated the space with some gray unlabeled nodes that show Node A as part of other relationships elsewhere in the space -- its visual symbol is heaped with some others on the right side of the example hypertext, as well as heading up a list on the left. Thus hypertext -- encounters with content in different contexts -- is realized through recurrent references that use visual symbols of underlying objects instead of through traversal-oriented links.[1]

1. Rosenberg contrasts the choices a reader must make in traversing links to Node A *or* D *or* E (see, for example, Figure 1a) with spatial simultaneities, in which a reader can at once apprehend "A *and* D *and* E"[22].

Figure 1 illustrates several important characteristics of spatial hypertext: the separation of symbol and underlying content; the use of these visual symbols to create hypertextual meaning; and the ability to leave structure *implicit* and *informal*. This final point -- the ability to leave structure implicit and informal -- is one of the crucial distinctions of spatial hypertext. Of course, in document-based hypertext, structure may be left implicit and reclaimed through analysis of document content, but the structure that is computed is regular; spatial hypertext allows people to express what DeRose refers to as “extensional structure” [6], idiosyncratic and dependent on the situation at hand, perceived through context, not computed from content. Ambiguity may be left unresolved without compromising the integrity of the structure[1]. For example, in Figure 1c, if we pushed Node E out of alignment with Nodes A and D, we could use the misalignment to express our uncertainty about whether Node E really belongs to the same set as A and D. The characteristics of spatial hypertext allow structure to emerge very gradually, as people work with the visual characteristics and spatial positions of symbols.

Spatial hypertext is thus inherently flexible, decidedly less formal than other models of hypertext, and readily supports volatility and change.

Many of these characteristics are derived from the basic affordances offered by visual and spatial modes of working. People may find it difficult to express how or why content is interconnected, but they are accustomed to arranging media -- either physical or electronic -- in space. Computational tools like text editors, structured graphics editors, outliners, some hypertext systems, and even multiple window displays (where different documents or different portions of the same document can be readily juxtaposed) all support spatialized content. Office workers frequently shuffle papers to make sense of them, and use the physical space around their offices as important adjuncts to their more organized file cabinets [13][14].

Sometimes links are noted physically or graphically in these media -- in computational media, lines are drawn from one content item to another, or boxes group multiple items; in non-computational environments, people fasten documents together with paper clips, staple them, put rubber bands around sets of videotapes, or otherwise physically connect related things -- but often

1. In explicit hypertextual linking, links are either there, or they are not. When links are given weights or may exist independently from nodes, the *existence* of the link still must be specified unambiguously.

interconnections are left implicit, to be resolved only if use demands it.

Looking at practice

Instances of spatial hypertext-like constructs appear in many computational and non-computational settings. How can we learn from practice to design appropriate components for and interface to a spatial hypertext system? It is crucial to watch users in action, to look at the results of their work in a variety of environments, and to talk to them about what they are trying to accomplish. Experiences with the Aquanet hypertext system, which uses a browser-based approach, coupled with a survey of the types of spatial structures people created in three different hypertext systems, NoteCards, the Virtual Notebook System, and Aquanet, provided us with a basis for designing a spatial hypertext system and reflecting on directions for spatial hypertext [16].

In our survey, we found that authors sometimes prefer to express relationships among nodes[1] by using geometric cues like proximity and alignment, and visual cues like graphical similarity. These geometric and visual cues correspond to Bertin's notion of planar and retinal variables [3]. By combining geometric and visual cues, authors may build up surprisingly complex hypertext-like structures.

Geometrically-based relationships arise from spatial configurations of nodes: nodes may be "close to" or "on top of" each other (be proximate); nodes may be "under," "over," "to the left of," or "to the right of" one another (show deliberate alignment in either the x or y dimension); or nodes may clustered or grouped (set apart from other elements of a space). People readily use these spatial configurations to portray relationships. For example, two adjacent or overlapping nodes might be related in a manner analogous to a simple binary link.

Visual characteristics also are used to signify relationships or clarify spatial cues. So nodes of the same color distributed throughout a space might belong to a common set. Or, nodes that are similar in appearance might elaborate the role of a geometrically-perceived link. For example, if we know Node Y is an annotation on its neighbor, Node Z, and Node Y is distinctive because it is

1. We use the term "node" here to mean the node itself or a visual reference to a node's full content, since both terminology and functionality varied among the systems and applications surveyed.

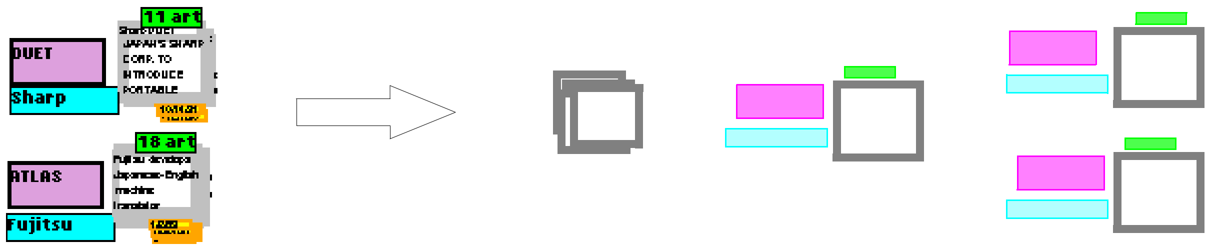

Excerpt from Figure 2 Sets Composites Set of composites

Figure 3. Finding structure in a portion of Figure 2.

yellow, we can perceive the same annotative relationship between Node X and its yellow neighbor, Node W. While we have observed that people develop complex visual codings to represent different properties of nodes, these codings are apt to be undeclared and inaccessible for processing by a spatial hypertext system. Thus, from our analysis, we found that the most important and basic visual characteristic to consider is “similar to” (or, equivalently, “different from”). Using homogeneity and analogy as guides, people can portray and perceive different kinds of hypertextual structures.

Taken in combination, visual and geometric characteristics also suggest hypertext *composites*. Composites are higher level structures that include other nodes [8]. By this definition, any group of visual symbols that appears to be deliberately organized -- specific symbols occupying specific relative positions -- is a composite. More interesting are node layouts that exhibit repeated patterns of visual and geometric characteristics. These we refer to as *composite types*. We found composite types in many of the layouts surveyed.

Figure 2 shows a complex spatial structure created in Aquanet as part of a long-term analysis of machine translation systems and technologies. The figure includes graphical objects that are instances of seven different object types; each instance refers to underlying content. What kinds of structure can we, as readers, perceive in this information space? First, coupling geometric and visual characteristics, we can see that each corner of the space seems to be inhabited by different distinct patterns of graphical objects. It is useful to take a closer look at a small portion of one of these regions to examine other common types of hypertextual structure.

Figure 3 shows several structures taken from the lower right corner of this space. First, we can see that there are overlapping stacks of visually-similar grey-bordered objects; we can

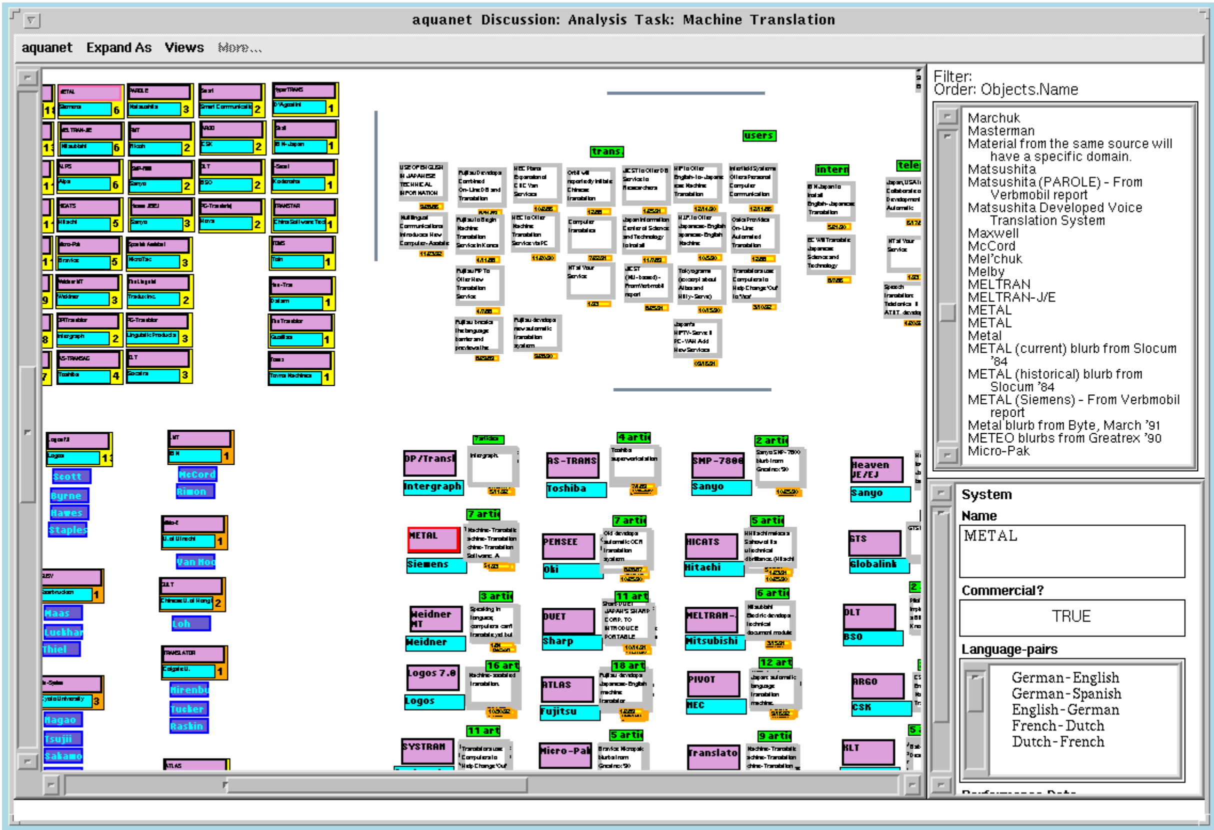

Figure 2. A multi-region space created in an Aquanet application

consider these to be sets. At a slightly coarser level of structure, we can see two instances of a composite type, the type that consists of four symbols: a purple rectangle above a cyan rectangle next to a green rectangle over a set of grey-bordered rectangles. Since the two instances of the composite type are aligned so regularly, they appear to be a set as well. Although little structure has been declared beyond node types, there is much hypertextual structure to be perceived in this example.

Given the results of our analysis, we can begin to anticipate common kinds of spatial structures and ways people would like to manipulate and interact with these structures. These results can then be used as a basis for designing an experimental spatial hypertext system.

An experimental spatial hypertext system

We have developed an experimental system, VIKI, to explore spatial hypertext as a geometric and visual structuring paradigm [17]. The system is designed to take advantage of the human perceptual system, spatial and geographic memory, and more generally, spatial intelligence; its emphasis is on flexibility, informality, and change. For readers, the system provides an opportunity to read in context, with awareness of the related, nearby nodes. For writers, the system supports the development of a visual language through informal interaction with a user interface similar to a graphics editor. A hierarchy of spaces helps keep complexity tractable for both readers and writers; users traverse through levels of spaces, each containing references to a particular set of nodes and subspaces -- an organizational strategy similar to that of Storyspace [2].

The interpretive task

In designing VIKI, we focused our efforts on supporting the interpretive process. Interpretation is that part of writing, collaborating, or thinking in which people collect the materials of interest and make sense of them in the light of their task and the background they bring to it. Of course, groups and individuals are neither monolithic nor consistent in their beliefs; to support interpretation, we considered ways of allowing people to record varying interpretations of a body of materials, using a multiplicity of frameworks or structuring schemes that make very different use of the same content.

Interpretive use of spatial hypertext is an important avenue of research to pursue. Through the increasing availability of on-line information sources (including other people, in addition to databases and corpora), people have access to vast quantities of materials relevant to their work. Yet there are relatively few tools that allow people to gather diverse materials and actively explore the various ways the materials may be interpreted.

People need such tools in many settings. For example, a successful product development organization will collect business-related material relevant to the positioning and viability of a new product. This business-related material may include electronic announcements and reviews of related products; competitors' home pages from the World-Wide Web; relevant patent documents retrieved from an information provider's extensive database; internal project-related materials (like engineering documents) created during the course of product development; electronic mail exchanges held within the organization; and, perhaps, information mined from Usenet newsgroup discussions. Once these materials have been gathered, members of the organization must make sense of them -- they need to decide which material is important, extract the salient information, evaluate its veracity, and organize it for their purposes. It is the understanding that a group develops and shares over time that makes a collection of materials valuable; the materials themselves are of limited utility without this shared understanding. The difficulty lies in expressing and recording this mutually-defined, ever-evolving interpretive structure.

Thus, interpretive structure is not inherent in document content, but evolves opportunistically through work. In practice, interpretive structure may arise through filing, in conversation, in juxtapositions of documents on one's desktop, in notes and annotations, in markings and marginalia [12]. All of these methods of expressing interpretations are necessarily partial, since they are not fully articulated; they readily tolerate ambiguity and fuzziness. Relationships among documents, assessments of their content, and interpretive abstractions can remain tacit through the entire sense-making process.

But externalizing these relationships, assessments, and interpretive abstractions is a crucial part of reflecting on one's own understanding of a problem and communicating one's understanding to others during the course of collaborative work. Our experiences demonstrate that such structure is difficult to express, even in hypertext; a tension arises between the tacit knowledge of the expert practitioner (the reader/author), and the means that hypertext tools provide to coax out interpretations as partial, emerging forms [23]. Spatial hypertext promises to be a good medium to do just that. VIKI forms the basis of our experimentation with the use of spatial hypertext in service of interpretive tasks.

VIKI's spatial hypertext model

VIKI's spatial hypertext model is simple, and follows from our survey of practice[1]; it consists of *objects*, which are node-like holders of content; *visual symbols*, which are the references to objects that a user manipulates in an *information space*[2]; *collections*, which are user-defined subspaces; and *composites*, which are the groupings of objects that form higher-level structures. Extensible structured *types* provide users with an abstraction mechanism. We describe

1. Much of our survey centered around information-intensive tasks, which were by-and-large analytic in nature.
2. Because VIKI is a spatial hypertext system, users interact through an information space; all content and references to content are accessible through this space.

each element of the model very briefly to illustrate how spatial hypertext is realized in our prototype. To illustrate the model, we use Figures 4 and 5 which again refer to an analysis of machine translation software, in which retrieved materials and notes are organized according to the systems they cover. The analysis is shown in progress -- some portions of the spatial hypertext are more highly structured than others.

Objects. Objects are the content-holding entities of VIKI's spatial hypertext model. Objects are semi-structured: they may have no internal structure, or they may have an unlimited number of fields that can be added when users wish to highlight or create content that doesn't fit within the object's existing internal structure. Each object that is part of the information space is referred to by one or more visual symbols. Double-clicking on a visual symbol causes the corresponding object's full content to appear in a separate window.

Visual symbols. Visual symbols are manipulable references to an object that may be moved freely. Visual symbols each have a graphical appearance, determined either by direct manipulation or through the types mechanism described below. Users can manipulate the size, shape, color, line thickness, and font characteristics of individual visual symbols, apart from their content. Indeed, users have much control over what is visible on a visual symbol at a given time. They can use the symbol's size to limit the amount of content revealed; they can specify which fields' contents are shown; and they can scroll through content to focus attention on a desired segment.

In contrast to the tight correspondence between instances of object types and their graphical appearance that is common to most browser-based hypertext systems, VIKI's visual symbols decouple objects and particular references to them in the display space. This decoupling, along with the ability to manipulate the appearance of individual symbols, provides users with an informal way to record different properties of an object according to its spatial context. For example, Figure 4a contains two references to the same content, a lengthy survey article about machine translation software, "Babelware for the Desktop." One reference to the article has a wide, dark border; the other, a narrow, less conspicuous outline. The author has used this convention, line thickness, to indicate how important the information in the article is. Apparently, the article is a more important source of information about Intergraph's DP-Translator system than it is about Globalink’s GTS.

Collections. Collections contain arbitrary spatial arrangements of objects or other collections, enabling them to form a system-supported hierarchy of spaces. They too are represented by manipulable visual symbols that refer to large, scrollable planes, in which objects may overlap. Figure 4a includes five collections (each with a blue background) titled according to the systems investigated. The edges of the collections act as clipping regions -- collections appear as portals into subspaces. Users may resize or zoom collections to reveal more or less of their contents. In Figure 4a, the collections are sized to reveal the degree to which their contents have been organized; the collection at the lower left is clearly less well organized than the others.

Collections give users the ability to immerse themselves in local meaning without losing global context. A collection may be motivated by the use of a particular visual/spatial structuring scheme (for example, as exhibited by the regions in Figure 2); it may group materials that pertain to a specific subtask (as is the intent of Rooms [10]); or it may indicate semantic cohesion, as is the case in Figure 4.

Since collections are acyclic, visually-nested hierarchical structures, the system provides a navigational mechanism that allows people to move between focal collections. Like Boxer [7], VIKI uses a spatial metaphor for navigation.

Figure 4 shows an example of spatial navigation in VIKI. Figure 4a shows the top level of an information space. Traversing into the collection titled “ATLAS” (by double-clicking on its

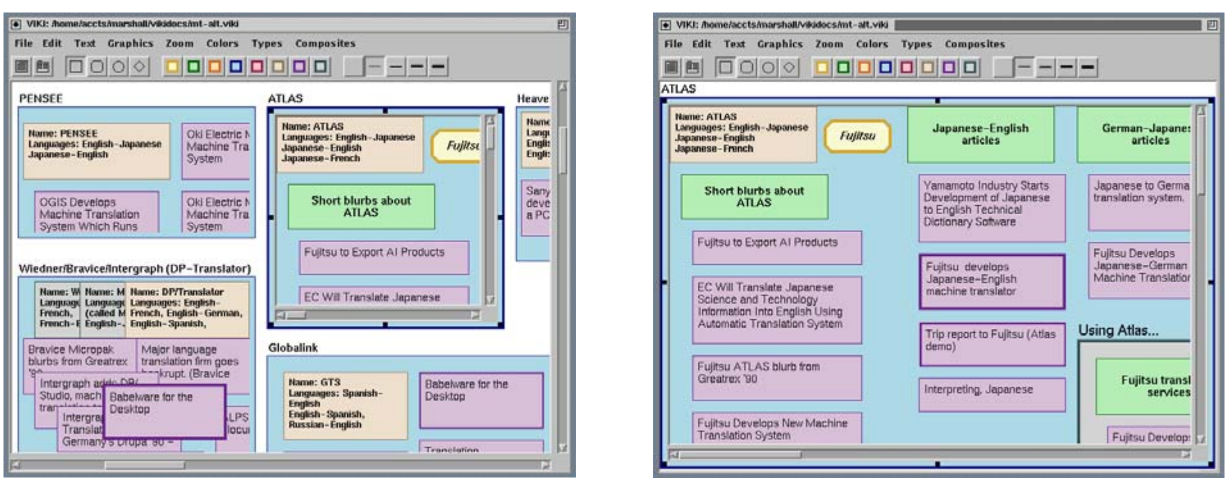

(a) (b)

Figure 4 - An example application showing objects, visual symbols, collections, and navigation.

border) causes it to become the focus. Figure 4b shows the result of this traversal. Now more of the contents of “ATLAS” are visible. Other objects have been revealed by the traversal, and we see that “ATLAS” contains another collection, “Using Atlas,” which is clipped. Multiple levels of the hierarchy may be traversed in a single step by clicking on any collection’s border.

Types. In our spatial hypertext model, types arise through use; we expect users to work first with graphical appearance and spatial layout of instances, and build up abstractions from these examples as generalities become apparent (see the description of prototypes and bottom-up design processes in Nanard and Nanard, this issue). A type encapsulates a particular appearance and internal structure, along with how this internal structure is displayed by visual symbols. Because content is often irregular, *ad hoc* fields may be added to the internal structure of any instance. For example, if one of the machine translation systems used a particular intermediate representation like Esperanto, we might want to add a field and value “Interlingua: Esperanto” to a single object.

Composites. Composites are structures that consist of combinations of two or more objects or collections in particular visual/spatial configurations; they closely model the structures we have observed in practice. Composite types are formed through recurrent patterns of constituent types; they may be defined directly or suggested to authors based on the results of recognition algorithms, as illustrated in Figure 5. We see composites (especially automatically perceived and maintained composites) as a viable way of building up relational structures in spatial hypertext.

Figure 4b shows an example of four composites that are instances of a composite type, which may be described roughly as, “a green rectangle over a set of purple rectangles.” Several instances of the purple rectangle (an object type) use a visual property -- border thickness -- to represent an attribute of the underlying objects.

Working within such a spatial hypertext model, authors can use objects, collections, composites, and the visual symbols that refer to them to gradually build up interpretations of the contents of a corpus of collected relevant material, eventually forming types, abstractions, and a sophisticated visual language for performing an analysis of this material.

Recognition

Instead of providing a mechanism for defining specific links between objects, VIKI focuses on the relationships that can be noted visually and spatially. These relationships may either remain implicit, or be recognized by structure-finding algorithms that automatically analyze spatial layout and visually salient properties of the constituents of a space [24]. Because structure is not enforced, but available on demand, expressive flexibility is not hampered by the definition of formal interconnections[1]. This approach addresses common problems associated with authoring hypertexts, such as the need to articulate tacit knowledge or commit to structure prematurely [25]. In short, structure is created easily, perceived by readers and writers, and, if requested, recognized by the system and used as the basis for further interaction.

Automated structure recognition, as we use it in our spatial hypertext prototype, assists authors in three different ways. VIKI may (1) use undeclared structure as a basis for interaction (for example, users may select and move the entirety of an undeclared structure); (2) help people use the object-collection-composite data model (for example, by locating potential collections and composites); and (3) provide inter-object abstraction (for example, by suggesting composite *types*) on demand.

Figure 5 shows an example of how the recognition facilities can help users locate and declare composite types, and more generally, incrementally formalize structure. In Figure 5, the structure recognition facilities have identified and suggested a possible composite type, a green rectangle (a type called “Label”) over a set of purple rectangles (a type called “Dialog”). The user has accepted the system’s suggestion, and has declared the new type.

Recognition is an accelerator, and enhances the usability of the basic facilities; it also usefully demonstrates that certain kinds of regular structures do tend to emerge in real use situations and can be made into formal hypertext constructs as needed.

Spatial hypertext in use

We have found that information-intensive applications refer to external materials; for

1. This approach is philosophically similar to Moran et al.’s notion of *freeform interaction [19]*.

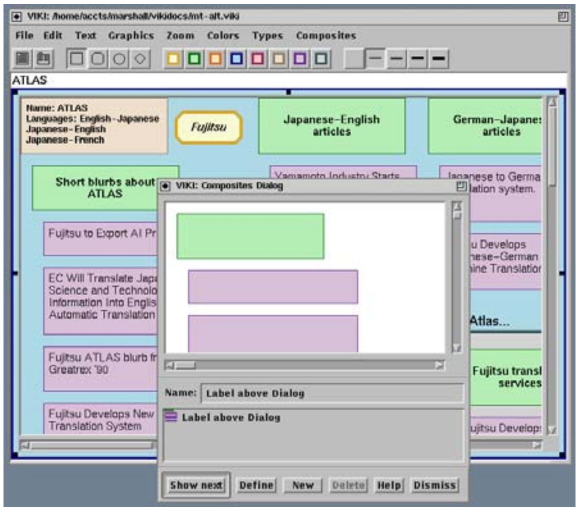

Figure 5 - Applying composite type recognition to the collection shown in Figure 4b.

these applications it is important to see spatial hypertext not as the primary holder of content, but rather as a *superstructure* for organizing, and more importantly, interpreting large-scale information resources. This perspective takes account of the mutability of any external materials, and the role of structuring paradigms in making large scale information resources serve communities of practice [18].

To demonstrate the use of spatial hypertext in this role, we have extended VIKI so that objects may refer (through automatic processing of user-supplied URLs) to World-Wide Web pages [1]; users can start Web viewers (such as Mosaic) by interacting with VIKI objects. This facility allows people to take advantage of the Web as an information resource, enabling them to interpret and situate Web materials in the context of their work; they are free to mix annotation and source materials, and add as much structure as need be. Thus VIKI provides an expressive superstructure -- a medium for customizing views of the Web.

Figure 6 shows an in-progress VIKI information space that refers to Web pages pertaining to research topics of interest to its author. The author has used collections to gather materials on different topics and has used color to indicate the types of Web pages the visual symbols refer to. For example, red symbols refer to journals, purple to ACM Special Interest Groups, green to conferences. The author has also annotated materials in several of the collections with his own notes and questions, such as his notes about how visual languages use spatial parsing. In this

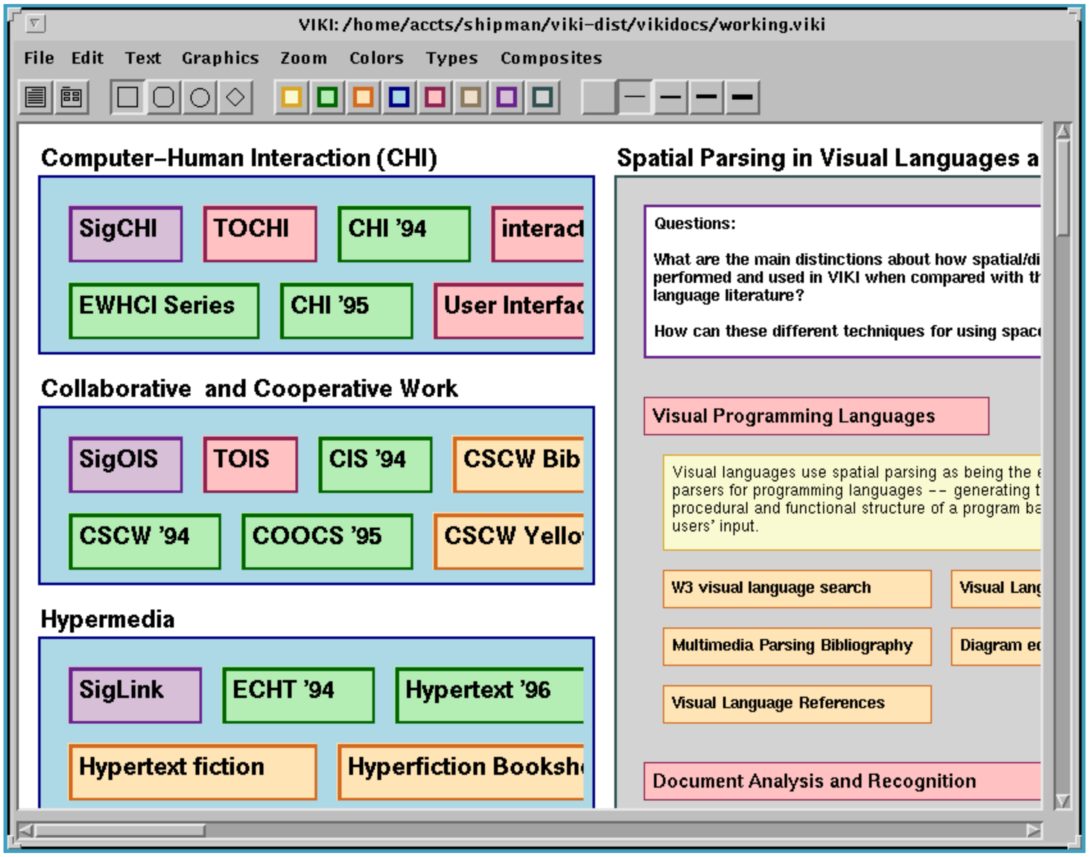

Figure 6. A VIKI information space used to organize topical World-Wide Web pages and notes

example, we see a mixture of structured and unstructured materials, and a mixture of external references and annotations. We see this as a potentially useful way for communities to cull materials from information resources and record their interpretations of these external materials.

Summary

Spatial hypertext can provide workgroups with a means of recording their interpretations of the large bodies of information relevant to their day-to-day activities; this information is quickly becoming available through a variety of networked electronic resources such as the World-Wide Web, on-line information services, and various digital library efforts.

By addressing the needs of applications in which structure is emerging, changing, or difficult to articulate, spatial hypertext promises to be a useful alternative to traditional node-link models of hypertext. Spatial hypertext uses perceptible characteristics of recurrent visual symbols -- references to underlying content -- to create structures such as annotative links, sets, and composites. We have observed that these structures arise in practice, through human interaction with computational and non-computational materials.

Spatial hypertext not only provides the means of visualizing complex structures; it also can provide perceptually-based feedback for navigation and a range of possibilities for expressing emerging relationships. As Kaplan and Moulthrop point out, spatial hypertext need not only reflect an author's picture of relationships among nodes; it may also be used to represent a reader's growing conceptualization of meaning [11]. Hyperspace may draw upon physical metaphors (see Dieberger and Bolter sidebar, this issue); it may also serve as fertile ground for exploring new kinds of abstract spaces that exhibit characteristics not found in the physical world.

We have begun such an exploration of abstract information spaces with an experimental system, VIKI, which uses a simple spatial hypertext model and a familiar user interface, coupled with mechanisms to support the gradual emergence of structure, such as types and recognition algorithms. VIKI is designed to support the volatility and change we have observed in information-intensive work. By investigating the use of spatial hypertext in such applications, we can extend hypertext concepts to realms in which prescriptive design methods break down.

Acknowledgments

This paper is dedicated to the memory of our colleague James H. Coombs, who contributed extensively to the development of the VIKI spatial hypertext system.

References

[1] Berners-Lee, T., Cailliau, R., Luotonen, A., Nielsen, H.F., Secret, A. The World-Wide Web. *Communications of the ACM* 37, 8 (August 1994), 76-82.

[2] Bernstein, M., Bolter, J.D., Joyce, M., and Mylonas, E. Architectures for Volatile Hypertext. In *Proceedings of Hypertext '91* (Dec.16-18, San Antonio, TX). ACM, New York, 1991, pp. 243-260.

[3] Bertin, J. *Semiology of Graphics*. University of Wisconsin Press, Madison, Wisconsin, 1983.

[4] Bolter, J.D. *Writing Space: The Computer, Hypertext, and the History of Writing*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1991.

[5] Conklin, J. and Begeman, M. gIBIS: A Hypertext Tool for Exploratory Policy Discussion. ACM Transactions on Office Information Systems 6, 4 (October, 1988), 303-331.

[6] DeRose, S.J. Expanding the Notion of Links. In *Proceedings of Hypertext '89* (Nov. 5-8, Pittsburgh, PA). ACM, New York, 1989, pp. 249-258.

[7] diSessa, A. and Abelson, H. Boxer: A Reconstructible Computational Medium. *Communications of the ACM* 29, 9, (1986), 859-868.

[8] Grønbæk, K. and Trigg, R.H. Design Issues for a Dexter-Based Hypermedia System. *Communications of the ACM* 37, 2 (Feb. 1994), 40-49.

[9] Halasz, F. Reflections on NoteCards: Seven Issues for the Next Generation of Hypermedia Systems. *Communications of the ACM* 31, 7 (July 1988), 836-852.

[10] Henderson, D.A., Jr. and Card, S.K. Rooms: The use of multiple virtual workspaces to reduce space contention in a window-based graphical user interface. *ACM Transactions on Graphics* 5, 3 (July 1986), 211-243.

[11] Kaplan, N. and Moulthrop S. Where No Mind Has Gone Before: Ontological Design for Virtual Spaces. In *Proceedings of the European Conference on Hypermedia Technologies* (Sept. 18-23, Edinburgh, Scotland). ACM, New York, 1994, pp. 206-216.

[12] Levy D. and Marshall, C.C. Going Digital: A Look at Assumptions Underlying Digital Libraries. *Communications of the ACM* 38, 4 (April 1995), 77-84.

[13] Malone, T.W. How do People Organize their Desks? Implications for the Design of Office Information Systems. *ACM Transactions on Office Information Systems* 1, 1 (Jan. 1983), 99-112.

[14] Mander, R., Salomon, G., and Wong, Y.Y. A 'Pile' Metaphor for Supporting Organization of Information. In *Proceedings of CHI '92* (May 3-7, Monterey, CA). ACM, New York, 1992, pp. 627-634.

[15] Marshall, C.C., Halasz, F.G., Rogers, R.A., and Janssen, W.C. Jr. Aquanet: a hypertext tool to hold your knowledge in place. In *Proceedings of Hypertext '91* (Dec. 16-18, San Antonio, TX). ACM, New York, 1991, pp. 261-275.

[16] Marshall, C.C. and Shipman, F.M. III Searching for the Missing Link: Discovering Implicit Structure in Spatial Hypertext. In *Proceedings of Hypertext '93* (Nov. 14-18, Seattle, Washington). ACM, New York, 1993, pp. 217-230.

[17] Marshall, C.C., Shipman, F.M. III, and Coombs, J.H. VIKI: Spatial Hypertext Supporting Emergent Structure. In *Proceedings of the European Conference on Hypermedia Technologies* (Sept. 18-23, Edinburgh, Scotland). ACM, New York, 1994, pp. 13-23.

[18] Marshall, C.C., Shipman, F.M. III, and McCall, R.J. Making Large-Scale Information Resources Serve Communities of Practice. To appear in *Journal of Management Information Systems* 11, 4 (Spring 1995).

[19] Moran, T., Chiu, P., van Melle, W., and Kurtenbach, G. Implicit Structures for Pen-Based Systems Within a Freeform Interaction Paradigm. To appear in *Proceedings of CHI '95* (May 7-11, Denver, Colorado). ACM, New York, 1995.

[20] Nanard J. and Nanard M. Should Anchors be Typed Too? An Experiment with MacWeb. In *Proceedings of Hypertext '93* (Nov. 14-18, Seattle, Washington). ACM, New York, 1993, pp. 51-62.

[21] Parunak, H.V. Don't Link Me In: Set Based Hypermedia for Taxonomic Reasoning. In *Proceedings of Hypertext '91* (Dec. 16-18, San Antonio, TX). ACM, New York, 1991, pp. 233-242.

[22] Rosenberg, J. Navigating Nowhere / Hypertext Infrawhere. *SIGLINK Newsletter* 3, 3 (Dec. 1994),16-19.

[23] Shipman, F.M. III and Marshall, C.C. Formality Considered Harmful: Experiences, Emerging Themes, and Directions. Xerox PARC Technical Report ISTL-CSA-94-08-02, 1994.

[24] Shipman, F.M. III, Marshall, C.C., and Moran, T.P. Finding and Using Implicit Structure in Human-Organized Spatial Layouts of Information. To appear in *Proceedings of CHI '95* (May 7-11, Denver, Colorado). ACM, New York, 1995.

[25] Shum, S. Cognitive Dimensions of Design Rationale. In *People and Computers* VI (eds.) D. Diaper and N.V. Hammond, Cambridge University Press, Cambridge, 1991.

About the Authors:

Catherine C. Marshall is an Associate Research Scientist in the Department of Computer Science at Texas A&M University. She is currently the US Vice Chair of ACM SIGLINK. Her research interests include hypermedia, digital libraries, and computer support for information-intensive intellectual work. Author's Present Address: Department of Computer Science, Texas A&M University, College Station, TX 77843-3112.

Electronic addresses: marshall@cs.tamu.edu, http://bush.cs.tamu.edu/~marshall

Frank M. Shipman III is an Assistant Research Scientist in the Department of Computer Science at Texas A&M University. He has been pursuing research in the areas of hypermedia, computer-supported cooperative work, and intelligent user interfaces since 1987. The focus of Frank's research is the combining of informal and formal representations in interfaces and methods for supporting incremental formalization. Author's Present Address: Department of Computer Science, Texas A&M University, College Station, TX, 77843-3112.

Electronic addresses: shipman@cs.tamu.edu, http://bush.cs.tamu.edu/~shipman