

Fluid Links for Informed and Incremental Link Transitions

Polle T. Zellweger, Bay-Wei Chang, Jock D. Mackinlay

Xerox PARC

3333 Coyote Hill Road

Palo Alto, CA 94304, USA

E-mail: {zellweger, bchang, mackinlay}@parc.xerox.com

ABSTRACT

We have developed a novel user interface technique for hypertext, called *fluid links*, that has several advantages over current methods. Fluid links provide additional information at a link source to support readers in choosing among links and understanding the structure of a hypertext. Fluid links present this information in a convenient location that does not obscure the content or layout of the source material. The technique uses perceptually-based animation to provide a natural and lightweight feeling to readers. In their richer forms, fluid links can provide a novel hypertext navigation paradigm that blurs the boundaries of hypertext nodes and can allow readers to fluidly control the focus on the material to support their current reading goals.

KEYWORDS: Fluid UI, fluid links, hypertext navigation paradigms, rhetoric of departure, scent, user interface, animation

INTRODUCTION

The cognitive overhead associated with choosing whether or not to follow a given link when reading hypertext has long been noted [7]. Landow has described the hypertext rhetoric of departure [14], calling for authors to structure the material at each link anchor so as to help readers discover what the relationship between the source and destination material is, and hence to help them decide whether the destination material is likely to be relevant to their current reading goals. Similarly, Furnas has recently discussed the need for a link source to show a *scent* or *residue* of content that can be reached at the destination and beyond, to help readers find the information they seek [9].

In addition, we note that the cognitive experience of travelling to and from the destination can be disruptive for any link traversal. It seems especially burdensome when the reader follows a link only to find that it is not relevant to her current reading goals. The reader must leave the source context, enter the destination context, determine that it is not

relevant, exit the destination context, re-enter the source context, and find her previous reading location.

Associating explanatory material with a link in order to inform the reader about the link and its destination can help alleviate some of the cognitive overhead of hypertext. This information helps the reader decide whether or not to follow the link, thereby reducing the likelihood of following irrelevant links and of missing relevant ones.

We call such explanatory material at the anchor site a *gloss*, appropriating the notion of a gloss being a brief explanation positioned in the margin or between the lines of a text. A gloss can consist of a description of the destination, an actual excerpt from the destination, the relationship of the destination to the source material, an annotation on the source material, meta-information such as author and creation date, link popularity or recommendation data, and combinations of these as well as other information.

Most hypertext systems provide no special location or mechanism to support glosses, leaving it to authors to rewrite their source materials to support hypertext readers. However, while a gloss at a link anchor provides useful information, placing it inline with the anchor text clutters up the source. Indeed, if added to an existing document, it may greatly interfere with the author's intended presentation. Furthermore, for original works such as literature, images, or other hypermedia content, it may not be appropriate to alter the source content.

A few early hypertext systems addressed some of these issues. The Hyperties system [16] provides glosses for its links dynamically at a separate location. In Hyperties, a brief description (5-35 words) of the destination appears at the bottom of the screen when a reader touches a highlighted topic word. The reader may then continue reading in the original article, or he may jump to the full entry about the selected topic. However, placing the gloss so far away from its associated anchor makes it difficult to compare the gloss with the source material. It also forces the reader to make large visual saccades in order to read the gloss and return to the anchor.

The Guide system [5] addresses the proximity issue with two kinds of anchors: replacement-buttons and note-buttons. Selecting a replacement-button causes it to be replaced

inline by the material at the destination of the link. Selecting a note-button causes the destination material to appear in a pop-up window that remains in view only until the mouse button is released. Both of these techniques minimize the distance the reader's eye must move to read the new material and then return to the source material. But following Guide's links reveals the full destination material: it does not have a notion of an intermediate gloss display before transitioning to a potentially large amount of destination material. The cognitive overhead of selecting links to follow still exists.

Recently, Microsoft's Internet Explorer web browser began making use of "link titles" [18, 20] an attribute of HTML anchors that dates back to at least HTML 2.0 [28]. Link titles are glosses that combine the proximal pop-up style of Guide with the lightweight interaction of Hyperties. Mousing over a link causes any associated link title to display as a small floating window next to the anchor text. As in Hyperties, readers can choose to follow the link or continue reading the source text. This technique relies solely on overlaying the gloss on the source material, potentially obscuring parts of the source to which the reader wishes to refer.

The technique we describe in this paper, *fluid links*, is both an answer to placing glosses as well as an avenue for new ways of experiencing hypertext. For example, interline expansion, one of several variations on the technique, addresses the problem of finding a place for glosses by dynamically creating room for the gloss. This kind of fluid link combines the desirable features of the previous methods: it provides lightweight access to the glosses; it places glosses in context with the source material; it does not obscure the source text. In addition, it uses animation to move glosses between the foreground and the background, smoothing the experience of viewing the glosses and allowing it to become a perceptual rather than a cognitive activity. Fluid links also enable several kinds of augmented interaction with glosses and destination material in context with the source material.

The next section describes the graphical technique of fluid links in detail. The following section examines the kinds of glosses and other content that may be used in fluid links. We conclude with additional related work and a summary.

FLUID LINKS

"Fluid links" actually describes a class of techniques for augmenting hypertext links with an animated intermediate display of a gloss, presented in the context of the source material. We have implemented several fluid link techniques in two prototype systems, one written using Pad++ [1] and Python running on Unix platforms, and another written in C++ running on Windows. This section begins with a general discussion that focuses primarily on the interline expansion technique, and then describes three other graphical variations we have explored: margin callout, textual overlay, and high-resolution annotation.

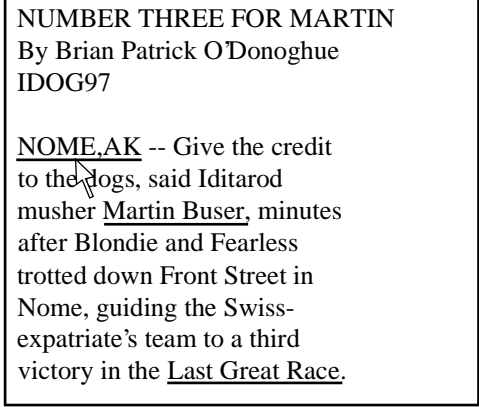


Figure 1. The hypertext source material. The "NOME, AK" fluid link activates when the mouse pointer moves over it.

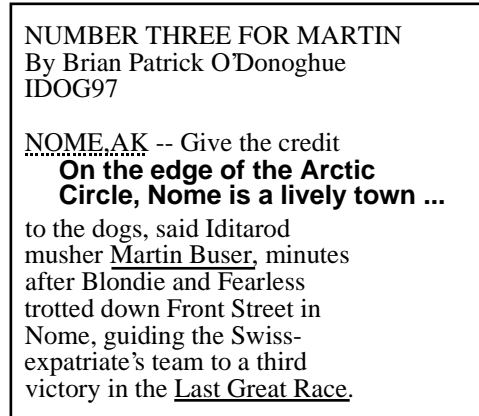


Figure 2. The gloss starts as the shape and size of the anchor underline. It grows smoothly to a readable size. At the same time the lines move apart to give it room.

As in most hypertext systems, the fluid link is indicated via a graphical mark that indicates the presence and (optionally) the extent of the link anchor. In our implementation, this is an underlining of the anchor text, but it could be bold-faced words, asterisks, or other marks. Like traditional hypertext, clicking on the anchor follows the link—the destination material replaces the current source material in the display. However, unlike traditional hypertext, simply mousing over the anchor causes a short phrase to expand from the mark, giving more information about the link and/or its destination. The presentation of this gloss is animated, allowing the reader's eye to follow the text growing smoothly from the underline to the fully sized text. The gloss expands to a position below the anchor, pushing the source lines apart temporarily to make room for the additional text (see Figures 1 and 2). We call this technique *interline expansion*. The space in between the lines at the anchor opens up to accommodate the expanding gloss, which pushes the rest of the source lines closer together.

To follow the link, the reader simply clicks on the expanded gloss. A traditional hypertext jump is then carried out, replacing the current view with the destination.

If instead the reader decides not to follow the link, she can simply move the mouse away (or on to the next anchor), backing out of the incremental step that she had taken. The gloss shrinks down again, and returns the reader to the anchored text (as in Figure 1). Viewing and dismissing the gloss is a lightweight operation: the reader need only indicate her interest by moving the pointer through the text, and the display smoothly changes to support that interest. Because the context of the rest of the document is maintained when glosses are expanded, there is less disruption for the reader.

The reader can freeze a gloss in view by clicking on it with an alternate mouse button. Clicking a frozen gloss again with the alternate button will cause the gloss to shrink to its original form. Selective gloss freezing can allow readers to control their view of the source and its glosses to support study and comparison. Frozen glosses can also act as a simple browsing "to do" list, allowing readers to delay following some or all of the seemingly relevant links until a more convenient time, without concern that such a link will be forgotten. The resulting decrease in interruption frequency may promote deeper engagement with the source material [15]. Useful auxiliary commands allow a reader to save the current arrangement of glosses for later use and to navigate through the current frozen glosses.

A final operation allows the reader to open all glosses for all visible links at once. This operation permits the reader to scan all glosses quickly.

The Fluid UI

We call our technique "fluid" because it provides lightweight, contextual, and animated access to the additional information, allowing the reader to fluidly shift attention from the source material to the gloss. Fluid links are a particular application of what we call Fluid User Interface techniques for accessing secondary information in the context of a primary information space. The Fluid UI approach has three parts:

1. *Visual cue.* A cue to the secondary material exists near its associated primary material. This cue can be textual or graphical. In the case of our hypertext application, this cue is the underline of the anchor.
2. *Animated transition.* The user indicates interest, triggering the secondary material to present itself more fully. For fluid links, we indicate interest simply by moving the mouse pointer over the anchor. When the pointer dwells on the anchor for more than a fraction of a second, the gloss begins to expand.
3. *Accommodation.* The primary material adjusts as needed to accommodate the expanded secondary material. This allows secondary material to be presented at full readable size, while still retaining the

context of the primary material. The interline expansion technique causes the source text to open up space in between its lines at the anchor site for the gloss.

Fluid UI techniques mediate the relative salience of primary and secondary material. In the beginning, the primary material is the focus and takes up most of the space and most of the user's attention. The secondary material is hinted at by graphical cues. When the user focuses on the secondary material, it will change itself graphically to become more salient: increase its size, move to a prominent location, display in boldface, and so on. The primary material may adjust itself to become less salient: decrease its size, move away from the center, fade out to a lighter shade. How exactly the primary and secondary material negotiate for space and salience is dependent on the particular application and is an area for design.

Although animation may seem unnecessary or even distracting for experienced users, it is an essential component of the Fluid UI approach. By supporting object constancy, animated transitions allow the human perceptual system to track changes [24]. In contrast, transitions without animation can be disruptive, requiring cognitive effort to reconcile their start and end states. We take care to ensure that our animations are short but perceptually appropriate, typically around one second in duration. Long animations can likewise disrupt the work process, whereas short animations occur during the time a user normally takes to react to a transition. Bernstein has noted the value of the related concept of visual enactments to show process in information tasks [3].

Fluid link techniques

We have explored a variety of fluid link techniques that use the principles of the Fluid UI, including the interline expansion technique described above, margin callout, textual overlay, and high-resolution annotation. Each has its particular strengths and weaknesses.

The interline expansion technique works well in both scrolled text and page-oriented text environments. To make space in a scrolling environment, interline expansion does not need to push other lines together as much, if at all. Nevertheless, it is desirable to keep as much of the original context as possible in place, so some squeezing may be warranted. However, in this case text can be pushed off the top and bottom of the display, remaining available via the usual scrolling mechanisms.

Page-oriented text displays maintain the relationship of the text content to the page divisions of the document: readers progress to more material by moving to a new page rather than scrolling the document. Page-oriented displays have a variety of advantages [21]. In fact, our initial investigation of Fluid UI techniques was motivated by considering interactions that would be useful on a page-oriented portable document reader, such as the XLibris active reading machine [23, 26].

However, a page-oriented display limits the amount that interline spacing can be reduced. We can then reduce the font size of the source text, but this technique again has a readability limit. For very large glosses—as might occur when following nested glosses—see the next section), we would like the more distant source text to begin to recede from view, either by moving off the page or by shrinking to less readable sizes, so that text closer to the anchor remains readable. A gradual fading out at the extreme ends would serve to smooth the transition.

The *margin callout* technique does not alter the source material in order to display the gloss. Instead, it places the gloss in existing white space on the page, typically the margins. When the reader indicates interest in an anchor, a line extends out from the underline to the nearby side margin, and the gloss expands at that location (see Figure 3). Because the callout line travels in between text lines, the source material is not altered. This technique can allow a good amount of gloss material to coexist with the original layout of the document. Although margin callouts sacrifice close proximity to the anchor, the animation of the callout line and the subsequent animated expansion of the gloss effectively draws the reader's eye to the gloss, and then back again when the callout line shrinks back to the anchor.

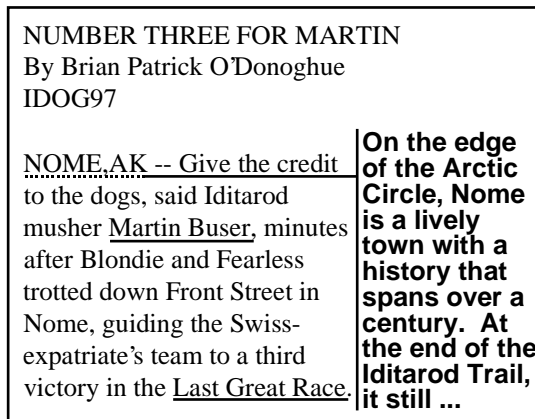


Figure 3. The margin callout animates a line from the anchor underline to the margin, where the gloss grows from a point to the fully readable size.

Another fluid link technique that preserves the layout of the source document is the *textual overlay*. A textual overlay link expands into the space below the anchor, as in the interline expansion technique. However, the source text in that space does not move away; instead, it fades in color. The difference in brightness of the two layers of text, as well as other differences like size and font, can make the gloss readable while still preserving some degree of readability of the underlying source text.

In all of the fluid link techniques, the positional and other graphical changes occur smoothly and quickly. Smoothly animating the changes is important for preserving the viewer's perception of the displayed objects. As the material

(text or graphics) grows, shrinks, and moves, the animated effect helps the eye to follow the object. For example, in the interlinear expansion technique, the gloss begins as a tiny string of text, and grows to a larger, readable size. The animation brings the eye from the anchor mark to the expanded gloss, and then back to the anchor mark and the source text as the gloss shrinks down again. Meanwhile, the surrounding source text moves apart to make room for the expanded gloss—in this case, the viewer's eye is not drawn to the animation, but the constancy that it affords in the viewer's periphery allows him to concentrate on the newly presented gloss. Animation serves both to direct the reader and to minimize surprising changes on the screen.

High-resolution fluid interfaces

Higher resolution displays are becoming more common. For example, dpiX, a Xerox company, is selling an evaluation prototype grayscale liquid crystal display with a resolution of 282 dpi [17]. Such high resolutions permit an extra level of gloss: a short description of the gloss itself can be placed at the anchor site at a very small font size, perhaps 3 or 4 points (see Figure 4). The fidelity of the display makes this tiny cue readable, and allows the reader to simply glance (or peer, if it is very tiny) at this augmented anchor to determine whether it is interesting enough to view the full gloss using the normal fluid link technique. In this *high-resolution annotation* technique, we have chosen to place the tiny gloss under its associated anchor to give an impression at a distance of a conventional anchor underline. On a display with lower resolution, it may in fact appear as a greeked line.

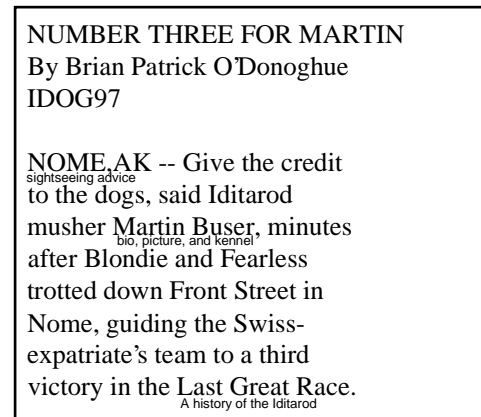


Figure 4. High resolution supports small fonts and textual link markers.

HYPertext CONTENT ENABLED BY FLUID LINKS

The previous section described how fluid links present glosses to facilitate hypertext viewing and decision-making. This section examines the content of those glosses, including several other novel techniques that fluid links enable.

Destination description

The most straightforward gloss is one describing the destination of the link. Many anchors do not have enough

contextual information from the surrounding source text to unambiguously convey the nature of the link destination. An anchor at the word "Iditarod" may lead to a short item defining it as the name of an annual dogsled race in Alaska, or to information about the town of Iditarod, or a listing of the competitors and their standings in the current race. A fluid link would supply a gloss so that the reader could learn something about the destination document before committing to leaving the current source material to go there. A useful variation would be to simply excerpt a small amount of the destination (its title, its first sentence) and use that as the gloss.

This kind of gloss can be explicitly authored for this link site, like the glosses used by Hyperties and HTML link titles.

Link relationship

Glosses also provide an opportunity for the author to explain the relationship between the source and the destination, an important piece of hypertext rhetoric that is frequently absent. Landow has observed that its absence can confuse and annoy readers [14].

Author-created link-specific glosses require explicit, manual author effort, but they are uniquely able to explain the intended relationship between the source and the destination. Such glosses may well be best for readers because of this capability. In addition, because they are local to the source document, accessing them is guaranteed to be fast enough to support the lightweight interaction required. However, they are subject to the traditional problems associated with distributed hypertexts, including a potential for becoming outdated if the destination changes without the source author's knowledge.

Dynamically-derived from the destination

It may be the case that it is impractical or undesirable for the author of the hypertext to supply her own glosses for the fluid links in the document and to maintain their accuracy over time. Dynamically-derived glosses for fluid links provide glosses with minimal author intervention. Such glosses could consist of the title of the destination, the first line, the first paragraph, a small image, a summary provided by the destination document's author, or even an automatic summarization of the destination text.

Dynamically-derived fluid links allow material from the other end of the hyperlink to appear inline and in context with the source document. This permits easier comparison of the source text and the destination text. In contrast with the traditional "go to" nature of link traversal, dynamically-derived fluid links provide a "bring from" approach that may better support readers and permit them to arrange the source and destination material to meet their needs.

Glosses that are dynamically derived from the destination do not require specific author action and are guaranteed to reflect the current state of the hypertext network. However, they cannot describe the relationship between the source and the destination. In addition, for distributed hypertext like the

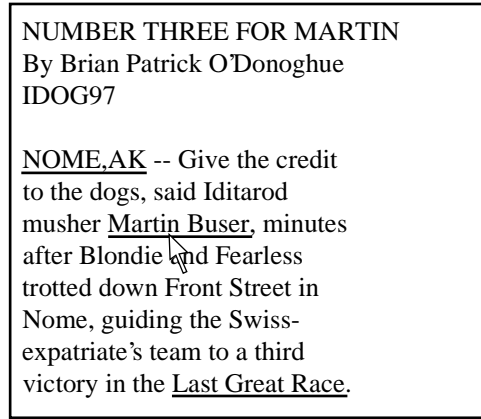


Figure 5. The reader is moving the mouse pointer over "Martin Buser" to activate the fluid link.

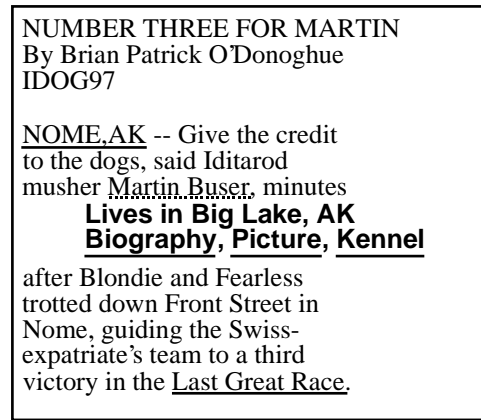


Figure 6. Multi-way links in a gloss. Clicking on an underlined anchor in the gloss will follow that link.

World Wide Web, accessing the destination may in general be too slow for lightweight interaction. Fortunately, the gloss material is usually small and thus quick to retrieve. The performance could also be enhanced by including specification and system support for partial node fetches, pre-fetching all dynamically derived link gloss contents as the final step of retrieving a node's source content. Finally, even if the gloss step is occasionally slow, it is guaranteed to be no slower than following the link in a standard way, and it has the cognitive advantage of appearing within the current context rather than requiring a full cognitive shift to the destination.

Meta-information

Glosses can also contain meta-information about the link and/or the destination, such as author, creation date, size, etc. Link popularity and recommendation information could also be valuable to readers as they consider whether or not to follow a link.

Multi-way links

Fluid links also provide a uniform mechanism and a convenient location to present multi-way links (see Figures 5 and 6). These may appear either from explicitly authored or dynamically derived situations. By embedding multi-way

links into a fluid link gloss, wasteful intermediate pages can often be avoided.

Simple fluid links allow the reader to click anywhere in the gloss to follow the link. In a gloss for a multi-way link, the reader clicks on the appropriate link within the gloss.

Nested glosses

Instead of presenting an additional level of standard links as in the multi-way link case, glosses may themselves contain fluid links. In this case, mousing over the fluid links in the gloss opens up a new gloss in the context of the original gloss (see Figures 7 and 8). Clicking on a sub-gloss takes the reader to that destination.

Nested glosses allow the reader to drill down through several nodes of hypertext in order to jump to a deeply buried destination, without taking on the overhead of changing context at each stage. The fluid interaction afforded by the fast, smooth animation of each level of gloss allows the reader to focus on the relevant new information (each new nested gloss) rather than on large state changes.

Complete destination

Finally, if the destination material is brief enough to fit on the source page in a way that retains readability of the source and the destination, the gloss can subsume the entire destination. In this case, the gloss/destination may be a terminal node of the hypertext (such as a definition or annotation) that contains no further links. Alternatively, it may contain links of its own that the reader would navigate as in multi-way links, or by fluidly expanding them, as in nested glosses.

Toward progressive disclosure of the destination

Even if the destination is lengthy, fluid links have the potential to support a progressive disclosure approach to its presentation. In this case, the gloss could begin as a dynamically-derived initial subset of the destination. As the reader continued to dwell on the gloss, more and more of the destination would arrive and be displayed. As the size of the destination grew, the source would be squashed more and more to the edges of the display, yielding a tunnel-like effect. Such novel navigation interactions are promising avenues for future exploration.

RELATED WORK

Hypertext

Although fluid links share many characteristics with the hypertext systems discussed in the introduction, they extend those capabilities by using animation to smoothly transition from the anchor to the gloss.

Hyperties has the notion of two-stage link-following, but unlike fluid links, the glosses are physically distant from the user's reading focus at the selected word, and authors are required to construct a single brief description for each topic manually.

Guide's anchors are replaced by their destination text. In contrast, fluid link anchors do not disappear when their

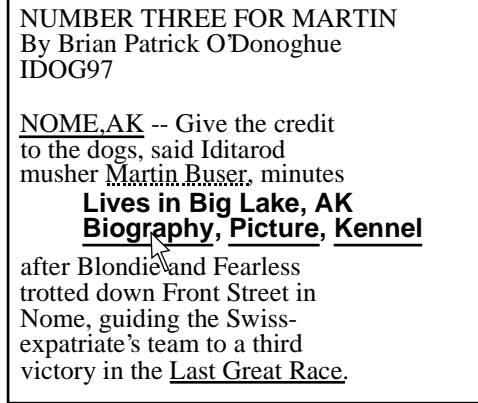


Figure 7. Moving the mouse pointer over an anchor in the gloss causes that fluid link to activate.

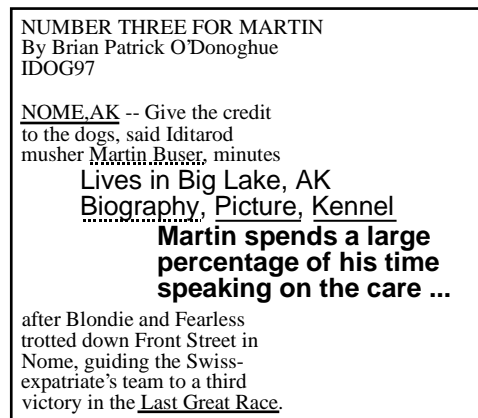


Figure 8. A nested gloss. Moving the mouse pointer away from the nested gloss will cause it to shrink back down.

destination material appears nearby, allowing for comparison. Guide also has no notion of a gloss for the link; all of the destination is used.

Internet Explorer 4 employs lightweight access (mouseovers) and proximal glosses via its pop-up window. For large glosses, however, its technique would obscure a great deal of the source text. In addition, there is no way to interact with the glosses; fluid links have the ability to implement multi-way links and nested glosses.

Many hypertext systems permit simultaneous viewing of source and destination content. A powerful example is the Elastic Windows Web browser, which also uses hierarchy and scaling to show multi-level context [12]. While simultaneous viewing can be useful either before or after following a link, it does not provide as contextualized a link previewing capability as the fluid links technique [10]. Fluid links bring information from or about the destination directly into the source view, reducing visual saccades and maintaining the reader's focus in the source context.

The concepts of hot and warm linking have previously been explored in hypertext systems. Hot linking refers to the notion that the contents of a source anchor will

automatically be updated to the current value of its destination anchor [27]. Warm linking is a related concept, introduced in the InterNotes extension to the Intermedia system [6], that allowed users to explicitly request such replacements: users could “pull” the contents of a destination anchor to replace the contents of the source anchor, or could “push” anchor contents in the opposite direction. In contrast with these concepts, fluid links do not replace the contents of the source anchor; rather, they allow for viewing portions of the destination in the context of viewing the source.

Overlapping

The NiF (News in the Future) Elastic Catalog [19] is a dense display of overlapping text that forms a rich hierarchical table of contents. Selecting a textual unit causes it to highlight and animate to a prominent visual space; at the same time, other textual units themselves highlight and move to form an organized layout based on the selected unit. Rosenberg has explored similar concepts of simultaneity in his Intergrams hypertext [25].

Like our work, the NiF Elastic Catalog uses animation and graphical changes to create visual relationships among elements of the information space. In addition, it also uses cropping to only show a few lines of text for items that are not being focused on. However, the Fluid UI and fluid links techniques emphasize scaling for handling secondary material, rather than relying only on overlapping.

Zooming

The Pad and Pad++ systems use interactive pan and zoom to place supporting material as small elements close to the corresponding primary material [1, 22]. A Pad++ technique called *semantic zooming* shows more of the supporting material as the zoom increases the size of the secondary elements. The Pad++ web browser [2] zooms each successive page to fill most of the screen, while surrounding pages are very small or zoomed out large behind the reader's viewpoint.

A problem with zooming as a technique for accessing small supporting elements is that it affects the entire display. When the reader zooms into a supporting element, the primary elements zoom out of view “behind” the reader. (The result is similar to hypertext techniques where the reader must remember the relationship between the destination material and the invisible source material.) The Fluid UI techniques maintain the context of the primary material by inverting the Pad++ zooming paradigm: they scale the material rather than zooming the view.

Focus+context

Distortion-oriented focus+context techniques such as the Perspective Wall and the Hyperbolic Tree [24, 13] scale up parts of the material on the display in order to focus on them. Non-focused areas of the display are spatially distorted to take up less space, but remain in spatial context. Furnas' generalized fisheye views [8] is a similar idea but without the notion of geometric scaling. Kaltenbach's rubber sheets scale space without scaling objects to permit

new items to be added to an existing layout [11]. Magic Lenses [4] use spatial lenses to change graphical material. In contrast, fluid links scale the gloss, but the scaling applies directly to the gloss text rather than as a result of geometric distortion or localized effect of the space.

SUMMARY

The technique we describe in this paper, *fluid links*, is both an answer to placing glosses as well as an avenue for new ways of experiencing hypertext. We have implemented several versions of the fluid link technique, including interline expansion, margin callout, textual overlay, and high-resolution annotation.

The fluid links technique combines the desirable features of previous approaches: it provides lightweight access to the glosses; it places glosses in context with the source material; it does not obscure the source text. In addition, it uses animation to move glosses between the foreground and the background, smoothing the experience of viewing the glosses themselves.

Levy has recently argued that hypertext is among several present-day forces combining to create shallower, more fragmented reading behavior [15]. Fluid links can serve to counteract this trend in several ways, promoting greater attention to a source document. First, link glosses appear in the visual context of the source, allowing readers to maintain the source as the primary focus. Second, their appearance is perceptually animated to reduce its load on cognitive processes. Third, by providing more information about the link and/or its destination, link glosses permit readers to avoid clearly irrelevant links without departing from the source context. Finally, the ability to freeze glosses allows readers to delay following some or all seemingly relevant links, thus permitting a longer uninterrupted engagement with the source document, without concern that such a link will be forgotten.

Fluid links also enable several kinds of augmented interaction with glosses and destination material in context with the source material, including multi-way links and nested glosses. In contrast with the traditional “go to” nature of link traversal, fluid links can provide a “bring from” approach that may better support readers and permit them to control their view of the source material to better meet their needs.

ACKNOWLEDGMENTS

Ken Fishkin was a great contributor to many early conversations on Fluid UI techniques, including their application to hypertext. Takeo Igarashi implemented the first version of interline expansion and the initial testbed for Fluid UI documents. Jim Pitkow suggested the use of meta-information in glosses. We thank the other members of the PDR team for their inspiration and enthusiasm.

REFERENCES

1. Bederson, Benjamin B. and James D. Hollan. Pad++: a zooming graphical interface for exploring alternate interface physics. *Proc. ACM UIST'94*, 1994.

2. Bederson, Benjamin B., James D. Hollan, Jason Stewart, David Rogers, Allison Druin, and David Vick. A zooming web browser. *Proc. SPIE Multimedia Computing and Networking'96*, Volume 2667, 1996, 260-271.
3. Bernstein, Mark. Enactment in information farming. *Proc. ACM Hypertext '93*, 1993, 242-249.
4. Bier, Eric A., Maureen C. Stone, Ken Pier, William Buxton, and Tony D. DeRose. Toolglass and Magic Lenses: the see-through interface. *Proc. SIGGRAPH'93*, 1993.
5. Brown, Peter J. Turning ideas into products: the Guide system. *Proc. ACM Hypertext'87*, 1987, 33-40.
6. Catlin, Timothy, Paulette Bush and Nicole Yankelovich. InterNote: Extending a hypermedia framework to support annotative collaboration. *Proc. ACM Hypertext'89*, 1989, 365-378.
7. Conklin, Jeff. Hypertext: an introduction and survey. *IEEE Computer*, September 1987, 17-41.
8. Furnas, George W. Generalized fisheye views. *Proc. ACM CHI'86*, 1986, 16-23.
9. Furnas, George W. Effective view navigation. *Proc. ACM CHI'97*, 1997, 367-374.
10. Kahn, Paul, Ronnie Peters, George P. Landow. Three fundamental elements of visual rhetoric in hypertext. In *Designing User Interfaces for Hypermedia*, W. Schuler, J. Hanneman, N. Streitz (Eds). ESPRIT Research Report Project 6532, HIFI, Vol. 1, Springer-Verlag, Berlin, 1995.
11. Kaltenbach, Marc. Screen management in hypertext systems with rubber sheet layouts. *Proc. ACM Hypertext'91*, 1991, 91-105.
12. Kandogan, Eser and Ben Shneiderman. Elastic Windows: a hierarchical multi-window World-Wide Web browser. *Proc. ACM UIST'97*, 1997, 169-177.
13. Lamping, John, Ramana Rao, and Peter Pirolli. A focus+context technique based on hyperbolic geometry for visualizing large hierarchies. *Proc. ACM CHI'95*, 1995.
14. Landow, George P. Relationally encoded links and the rhetoric of hypertext. *Proc. ACM Hypertext'87*, 1987, 331-338.
15. Levy, David M. I read the news today, oh boy: reading and attention in digital libraries. *Proc. ACM Digital Libraries'97*, 1997, 202-211.
16. Marchionini, Gary and Ben Shneiderman. Finding facts vs. browsing knowledge in hypertext systems. *IEEE Computer*, January 1988, 70-80.
17. Martin, Russel A. *et al.* An electronic document display: a 6.3-million-pixel AMLCD. *Journal of the SID* 4/2, 1996, 65-73. See also <http://www.dpix.com/>
18. Microsoft, Microsoft Internet Explorer 4.0, <http://www.microsoft.com/ie/ie40>.
19. Murtaugh, Michael. NiF Elastic Catalog. <http://table-saw.media.mit.edu/ecat/>.
20. Nielsen, Jakob. Jakob Nielsen's Alertbox for January 11, 1998. <http://www.useit.com/alertbox/980111.html>.
21. O'Hara, Kenton and Abigail Sellen. A comparison of reading paper and on-line documents. *Proc. ACM CHI'97*, 1997, 335-342.
22. Perlin, Ken and David Fox. Pad - an alternate approach to the computer interface. *Proc. SIGGRAPH'93*, 1993.
23. Price, Morgan N., Gene Golovchinsky, Bill N. Schilit. Linking by inking: trailblazing in a paper-like hypertext. *Proc. ACM Hypertext'98*, 1998.
24. Robertson, George G., Stuart K. Card and Jock D. Mackinlay. Information visualization using 3D interactive animation. *CACM*, 36, 4 (April 1993), 57-71.
25. Rosenberg, Jim. Navigating nowhere / hypertext infrawhere. *ACM SIGLINK Newsletter* Vol. III, No. 3 (December 1994), 16-19.
26. Schilit, Bill N., Gene Golovchinsky, and Morgan N. Price. Beyond paper: supporting active reading with free-form digital ink annotations. *Proc. ACM CHI'98*, 1998.
27. Towner, George. Auto-updating as a technical documentation tool. *Proc. ACM Conf. on Document Processing Systems*, 1988, 31-36.
28. World Wide Web Consortium. HTML 2.0 anchor element, http://www.w3.org/MarkUp/html-spec/html-spec_5.html#SEC5.7.3.