

Reading and Writing Fluid Hypertext Narratives

Polle T. Zellweger
Xerox PARC
3333 Coyote Hill Road
Palo Alto, CA 94304 USA
+1.650.320.9360
pzellweger@acm.org

Anne Mangen
Volda University College
P.B. 500
N-6101 Volda, Norway
+47.70.07.51.66
anne.mangen@hivolda.no

Paula Newman
Xerox PARC
3333 Coyote Hill Road
Palo Alto, CA 94304 USA
+1.650.812.4239
pnewman@parc.xerox.com

ABSTRACT

We describe a new way to present and author hypertext narratives. The Fluid Reader constructs a unified interactive text from the content of multiple nodes and allows a reader to explore alternative paths within it. The Fluid Reader has been available as a hands-on museum exhibit for nearly a year to date, where it has been enjoyed by readers of all ages. Its success has prompted further interest and development in Fluid hypertexts. We have designed and implemented an authoring tool called the Fluid Writer that uses a new treetable visualization to help authors construct and manage alternative paths in a Fluid hypertext. Finally, an exploration of the narrative implications of Fluid hypertext suggests that it may be more suitable than conventional hypertext for formulaic fictions such as mystery stories.

Categories and Subject Descriptors

H.5.4. [Hypertext/Hypermedia]: hypertext

General Terms: Design, Human Factors.

Keywords: Fluid Documents, Fluid Reader, Fluid Writer, Fluid hypertext, hypertext narrative, authoring, visualization, treetable, stretchtext

1. INTRODUCTION

Conventional hypertext narratives are often experienced as disconnected and disorienting [1][7][8][11][26]. Many present-day hypertext authors embrace and exploit this very point. However, other ways to present hypertexts may allow an interactive unfolding of alternative, hypertextual narratives that offer a different experience for the reader.

In this paper we describe a new way to present and author hypertext narratives. The Fluid Reader constructs a unified interactive text from the content of multiple nodes and allows the reader to explore alternative paths within it. Interactive animation is used to adjust the content and typography of the document dynamically as the reader selects alternate segments to view.

An important difference between Fluid hypertexts and more conventional hypertext narratives lies in its continuously-visible

context. Because the unified view always shows a complete path through the structure, Fluid hypertext is suitable for very fine-grained alternatives, at the level of phrases or even words. In contrast, traversing from node to node within a single sentence in a traditional hypertext would in general be overly distracting and disorienting for readers.

The process of authoring these fine-grained hypertexts poses significant challenges to authors. One basic problem is to ensure that ideas flow sensibly and sentence mechanics are properly observed across multiple fine-grained nodes. Other problems include gaining an appropriate view of the entire narrative and its structure, and comparing two or more paths in detail.

To address these problems, we have designed an authoring tool called the Fluid Writer that includes a custom treetable visualization of the Fluid hypertext structure. The treetable visualization facilitates simultaneous reading of multiple alternative paths through the structure. Interactive operations upon the treetable permit adding and removing content as well as focusing on two or more paths for comparison purposes.

Fluid hypertext offers new opportunities for readers and authors. We conclude by reflecting on some of these possibilities by comparing Fluid hypertext narratives with a more conventional hypertext environment, namely Storyspace [27].

2. BACKGROUND

The Fluid Reader explores the possibilities of using Fluid Documents concepts and technologies [30][32] for storytelling. The Fluid Reader was initially designed as one exhibit within a cutting-edge museum exhibition entitled “eXperiments in the Future of Reading” (XFR), which incorporated digital technologies from a range of research projects underway at Xerox PARC. This section summarizes the relevant aspects of Fluid Documents and describes the XFR museum exhibition as a whole.

2.1 Fluid Documents

Fluid Documents seek to improve the online reading experience via contextual views and animated typography. Contextual views combine primary and supporting information into a single unified view. Animated typography is used to guide the reader’s attention calmly to the revealed supporting material and later guide it back to the primary material. Animation also clarifies the changes as supplemental material arrives and departs from the page. User studies of Fluid Documents support the validity of these two basic tenets [32].

An annotated fluid document looks initially much like a hypertext page. When an underlined word or phrase is selected, the text on the screen smoothly opens up, providing space for the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

HT’02, June 11-15, 2002, College Park, Maryland, USA.

Copyright 2002 ACM 1-58113-477-0/02/0006...\$5.00.

presentation of an annotation (termed a *gloss*). Although the gloss is typically visually distinct from the primary material, showing them together on the same page reduces distraction and enhances understanding. Fluid Links uses animated glosses at link anchors to help readers choose among links and understand the structure of a hypertext [30].

2.2 eXperiments in the Future of Reading

Some have feared that video and multimedia, among other current technological trends, spell the death-knell of reading. The XFR museum exhibition was designed by a group of researchers at Xerox PARC to demonstrate how digital technologies can actually make reading richer, more varied, and more pervasive [2]. The show, which was in daily use at the San Jose Tech Museum of Innovation for six months in 2000, consisted of 12 interactive, hands-on exhibits about the history and potential futures of reading. It attracted more than 250,000 visitors. Its success has led to a three-year touring contract to other museums across the United States.

3. THE FLUID READER

The goal of the Fluid Reader was to design an interactive, hands-on, engaging, and imaginative look at a possible future style of hypertext narrative.

The Fluid Reader shows alternative versions of a dynamic online hypertextual document by adjusting its content and typography using interactive animation. It is based upon the common storytelling devices of asides and elaborations: readers can choose areas of interest to explore more deeply.

One complete path through the narrative structure is presented initially, with embedded interactive “knobs” that the reader can explore. Selecting a knob presents an alternative ending to a section of the story (such as a sentence or paragraph); these alternatives may also contain further knobs. Alternative section endings can add additional information, explanations, details, asides, jokes, or even lies. Although the reader can explore other alternatives within one section, the other sections remain visible. Thus, in contrast with conventional hypertext narratives, where the reader may determine which of several broader directions are followed, a central core of the narrative (including some version of the ending) is always visible.

A single story, entitled “Harry the Ape,” was written for the Fluid Reader and the XFR exhibit. The story also experiments with the connection between content (a story about animals that nest in size) and medium (sentences that nest inside each other).

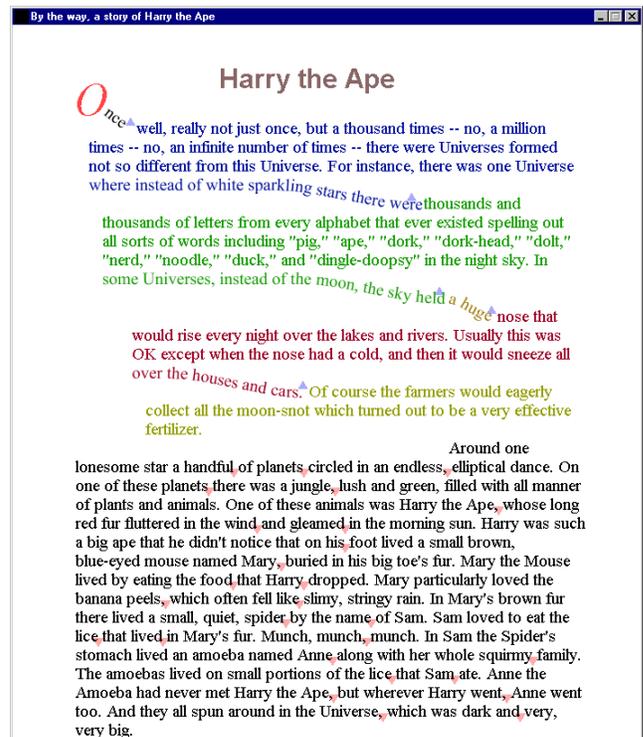
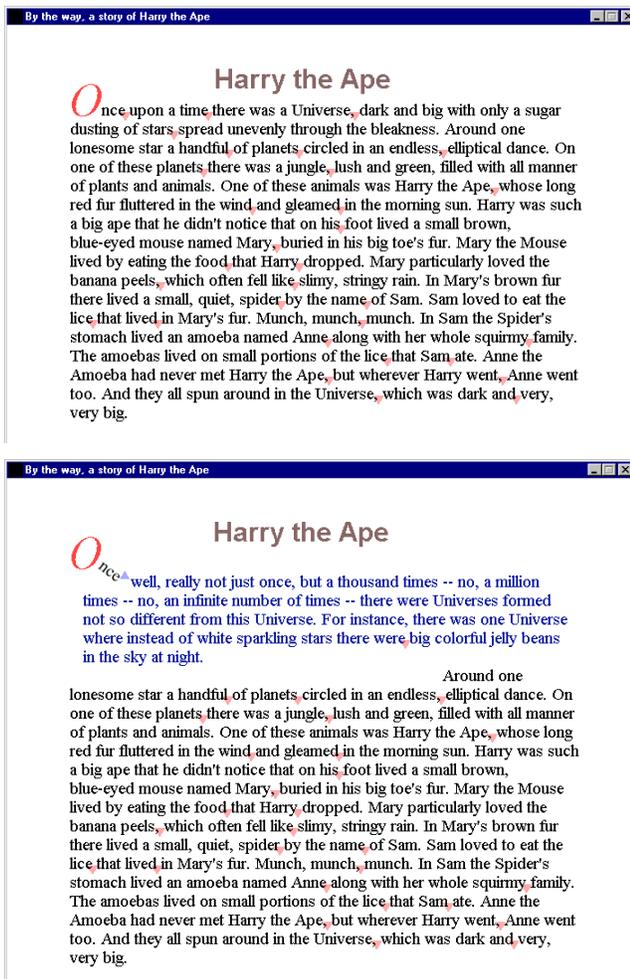


Figure 1. The Fluid Reader displaying the Harry the Ape story. The upper left screen image shows the reader's initial view of the Fluid Reader. The lower left image shows the result of opening the first red triangle in the story, after “Once”: an alternative sentence ending appears. Note that it also contains a red triangle. The upper right image shows the result of opening the first red triangle in each successive alternative ending.

3.1 Animated user interface

Figure 1 depicts a sequence of interactions with the Fluid Reader. Sprinkled throughout the narrative are small red triangles that mark choice points for alternative versions of the narrative. When the user selects (or “opens”) a red triangle, animated typography draws the reader further into the narrative. The downward-pointing red triangle gradually rotates and becomes an upward-pointing blue triangle. Meanwhile, the remainder of the sentence following the triangle gradually fades away, the segment of text preceding the triangle bends downward, and a new ending to the sentence appears. If the top-level narrative contains following sentences, these sentences are gradually pushed downward on the page to make room for the new ending. These changes are smoothly animated over an adjustable period of approximately one-half second.

Smooth animation is used for this process because only some of the material on the page is changing. Smooth animation draws upon basic characteristics of human visual perception. It helps the reader distinguish those parts that are changing, and must thus be read anew, from parts that have not changed.

The new material is also marked statically: it begins with a segment of bent text, its margins are wider and its text color is different. The new ending may contain further choice points, as shown within the indented colored text in Figure 1.

Selecting (or “closing”) a blue triangle reverses the animation to return to its original sentence ending. If a higher-level blue triangle is closed, animations for all deeper blue triangles are successively triggered as part of its closing animation.

3.2 Design for readability

Note that a major change from Fluid Documents to the Fluid Reader was the shift from glosses (annotations) to alternative endings. Glosses and alternative endings are both secondary, in the sense that neither appear at the surface level. However, glosses are intended to be viewed *together* with the related primary material, whereas alternative endings form a new text that *replaces* some part of the primary for a time. Helping the reader to follow the correct reading order effortlessly was a significant design problem. Note also the shift from link anchor (with extent) to choice point (without extent). For better or worse, the new mechanism thus far offers readers little advice or incentive to choose one direction over another.

The animations and the story appearance both before and after the selection of a choice point were carefully designed to reduce distractions, guide the correct reading order, and clarify the effects.

The “knobs” were instantiated as small filled pastel triangles placed along the baseline of a line, to make them visible while limiting distraction. The triangles appear in the background, so text can cross them if needed without sacrificing legibility.

The bending or “swooping” of the text allows existing text to be split and new text to be added without re-flowing text. Our experiences with re-flowing text indicate that its irregular movement is visually disruptive, and it does not support spatial memory. Instead, the swooping text guides the eye to the new sentence ending during the animation, emphasizes the change and has a statically meaningful result.

To make the exhibit inviting and successful, it had to be usable with minimal instructions. A placard at the exhibit said only: touch the triangles to explore the story more deeply. The exhibit used a large color display with a touchscreen rather than a more traditional mouse to reduce its “computerized” appearance and to avoid a sense of indirectness and intermediation for the visitor.

In order to simplify the reading experience for the museum setting, the museum version was limited to a single open alternative at a time (although additional choices can be made within that alternative). Thus opening a triangle in one sentence automatically triggers closing animations for any open triangles in other sentences.

3.3 Reactions of museum visitors

Our primary experience with readers to date has been in the museum context. The Fluid Reader has also been made available to a few universities, hypertext authors and students.

In the museum setting, readers often engaged deeply with the Harry story. Some spent 15 minutes or more exploring its various pathways, enjoying the discovery of surprises or jokes hidden inside deeply-nested alternatives. This is a considerable length of time, given that most museum designers expect a single exhibit to capture a visitor’s attention for about 30 seconds [2]. The most common reader was a young visitor between 8 and 12 years old, and the Harry story was written with this age group in mind. However, readers of all ages used and enjoyed the story.

3.4 Fluid hypertext data structure

In a conventional hypertext, nodes are arranged in an unrestricted graph structure. Each node may have an arbitrary number of in or out links, connected in arbitrary ways. Readers typically view the contents of one node at a time and may choose among several successors to each node. Alternative readings are constructed over time by selecting different nodes and/or a different ordering of nodes. A node’s content can vary widely in size or even media, but typically ranges from a paragraph to a few pages of material.

A Fluid hypertext has a more restrictive structure. At the top level, it consists of a sequence of trees: a backbone, if you will, with alternative tendrils sprouting from each segment. Within each tree, node boundaries define choice points for the reader.

Harry the Ape consists of a sequence of 14 trees in all; it was written so that each initial sentence corresponds directly to a single tree. However, this need not be the case in general: the initial content corresponding to a single tree could be either larger than or smaller than a sentence. Figure 2 shows the Fluid hypertext structure for Harry the Ape. In total, Harry the Ape consists of 14 trees, 196 nodes, and 2919 words. The largest two trees are sentences 1 and 2, with 35 nodes and 790 words¹, and 45 nodes and 633 words, respectively. The smallest tree is sentence 4 with 7 nodes and 82 words. In contrast, the initial top-level view of the story contains 221 words.

¹ The careful reader will find only 23 nodes for sentence 1 depicted in the figures for this paper. The structure has been simplified slightly to conserve space and maintain legibility. However, the screen images of Harry are actual images of the museum version.

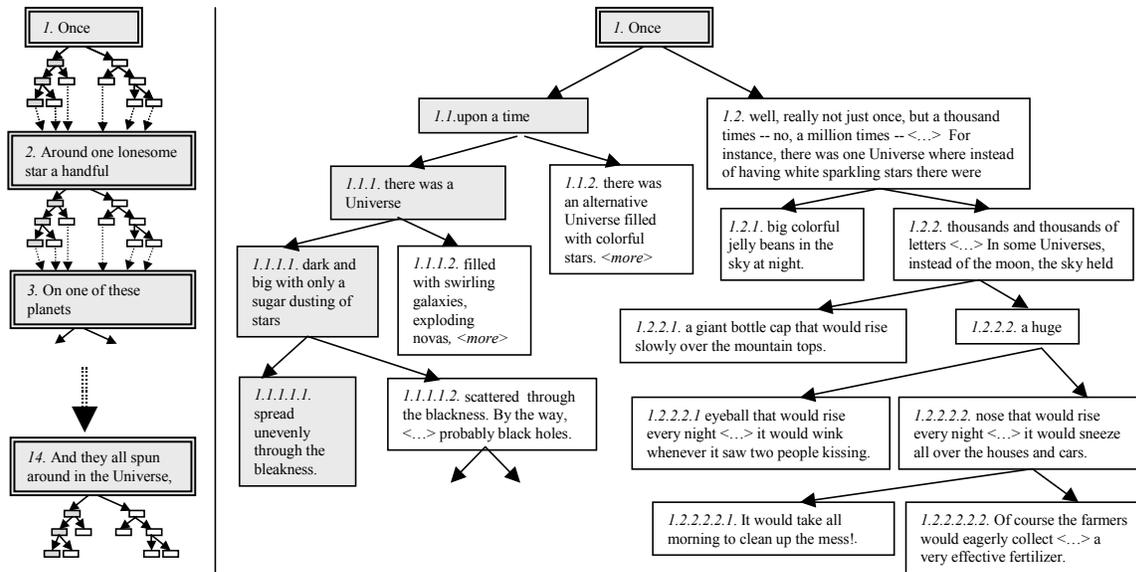


Figure 2. The Fluid hypertext structure for the Harry the Ape narrative. The left side of the figure shows an overview of the entire story as a sequence of trees. The right side focuses on the tree structure for the first sentence. The shaded nodes are visible in the initial view of the story. The double-outlined nodes are visible in every view of the story.

As stated earlier, the Fluid Reader shows a complete path through each tree at all times, from the root to a leaf node, followed by a path through the following tree, and so on. When a new alternative is displayed, all choice points occurring within that alternative’s subtree initially display their leftmost (or .1) node.

The current Fluid Reader is restricted to displaying binary trees (that is, trees with two children for each interior node). However, the Fluid Writer described in the next section can display trees having a varied number of children for interior nodes. The Fluid Writer can also display more general acyclic graph structures that permit two alternatives to merge into a single ending, thus providing a way to insert additional detail into the middle of a path while leaving the ending the same. We plan to extend the Fluid Reader to accommodate these structures in the future.

4. THE FLUID WRITER

The success of the Fluid Reader has prompted further interest and development in Fluid hypertexts. We have designed and implemented an authoring tool called the Fluid Writer, which uses a custom treetable visualization to help authors construct and manage alternative paths in a Fluid hypertext.

4.1 Previous makeshift authoring method

Because only a single, comparatively short, story instance was written for the museum, the primary focus of effort was placed on the visitor’s reading experience in the Fluid Reader.

The original authoring process was therefore decidedly makeshift and could not be recommended to additional authors. The author, Rich Gold, used Microsoft Word in outline mode to write the many branching endings that comprise the story. This approach posed significant challenges. Most importantly, it made it difficult to see paths and read along them easily without laborious transitions to and navigation within the Fluid Reader itself.

The ability to see paths and read along them easily is central to viewing a complete idea in a fine-grained hypertext. It is thus needed to support the construction of grammatically-correct sentences across fine-grained node boundaries, as well as proofreading. This ability is also central to following a particular narrative progression amid the branching alternatives.

Unfortunately, conventional representations for tree-structured material—namely, node-link diagrams (as in a typical hypertext structure view) and linear, indented displays (as in an outline)—obscure path information. Node-link diagrams with significant content are typically either quite large and sparse, to keep unrelated nodes distant, or their physical layout intersperses unrelated paths, to keep the diagram compact. Either way, reading along a path requires potentially distracting wayfinding effort between nodes. Another style of node-link diagram uses small nodes with little or no content to show the structure, and requires users to view node content in a separate window. Separating structure and content is particularly harmful for fine-grained hypertexts. In a linear, indented display, a given node (shown as one or more lines of text) may appear at a considerable distance from its parent or children, so that an individual path is difficult to follow.

4.2 The unaligned treetable visualization

The limitations of the previous representations, experienced firsthand by the first author as editor of the Harry the Ape story, led us to design a new visualization. The resulting unaligned treetable visualization is a variation on and extension of treetables recently developed for visualizing email threads [22].

Figure 3 shows an unaligned treetable representation of the first sentence of Harry the Ape. Each column represents a single path from the root to a leaf, and each cell exactly spans the cells representing its children in the tree. Unlike a conventional table with well-defined rows, cells representing different children may

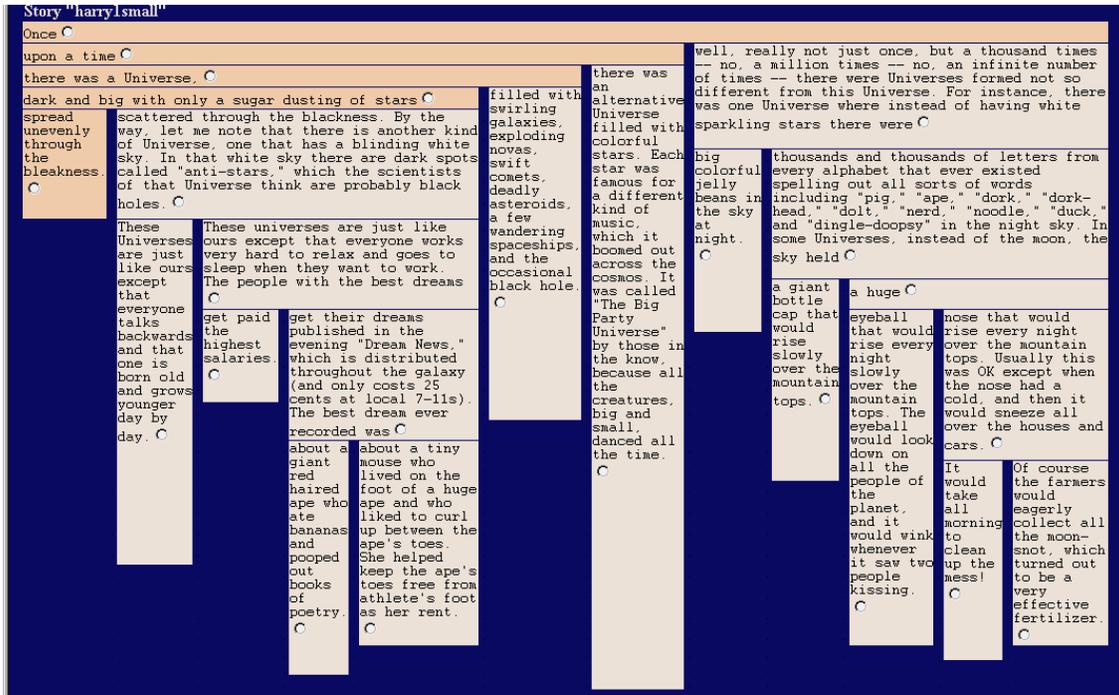


Figure 3. An unaligned treetable visualization for the first sentence of Harry the Ape.

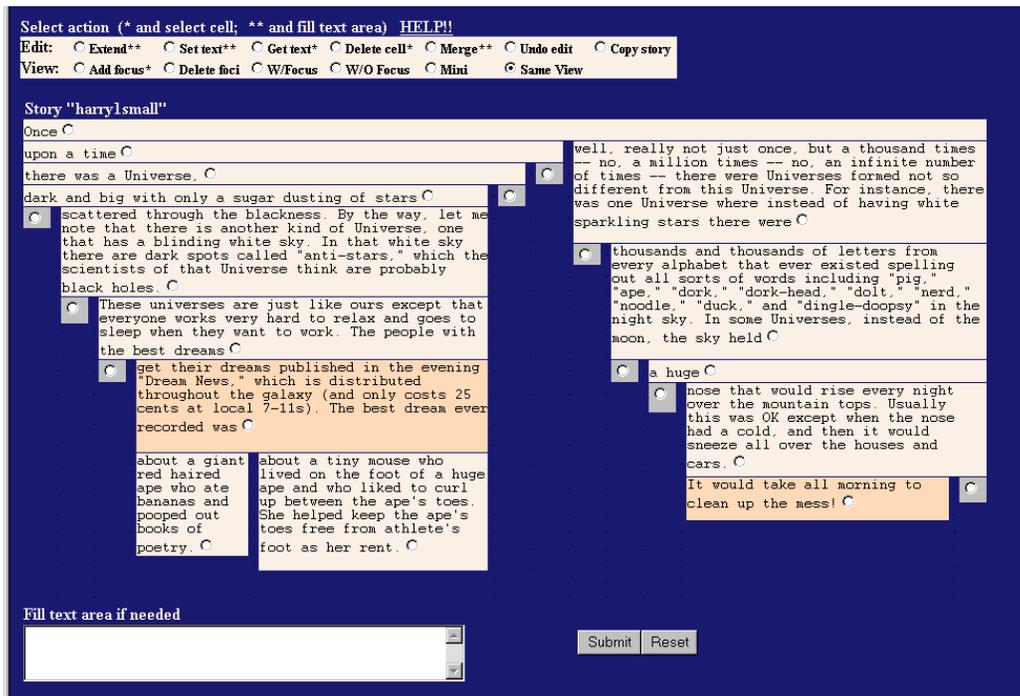


Figure 4. Treetable focus operations within the prototype Fluid Writer window, showing also the control menu at the top of the screen. The author has selected two nodes as focus nodes: the interior node “get their dreams published...” on the left and the leaf node “It would take all morning...” on the right, both shaded. All paths that pass through the focus nodes are displayed, and all other nodes are shrunk to small stubs. In this initial web-based version of the Fluid Writer, asterisks in the control panel indicate how many parameters each operation requires: 0, 1 (a cell), or 2 (a cell and the text area).

vary in height.² Also, columns may have different widths to minimize wasted space and limit the height of the whole.

Each vertical path through the treetable represents a reader-selectable version of the narrative, where each successive level of the tree represents a choice point among several alternatives. This new visualization allows authors both to see the entire structure of the narrative and to read each alternative path conveniently (by reading vertically through the treetable). This layout is particularly useful for the proofreading phase of the authoring task, because the author can check punctuation, subject-verb agreement, etc. across the parent-child boundary for each of the children of a node. In addition, this visualization allows the author to compare different paths easily and simultaneously. In contrast, access from the dynamic Fluid Reader would require successive views, separated by one or more alternative selections to transition from the view of one path to the other.

Any tree visualization, including cone trees [25] or hyperbolic browsers [14], could potentially be used to lay out textual content of hierarchical nodes. However, treetables provide a particularly effective way to present text in a way that permits reading smoothly from one node to the next, across node boundaries.

Treetables offer the ability to read text in vertical columns while simultaneously viewing hierarchical relationships. In fact, unaligned treetables are the most compact static way to lay out hierarchical text. They draw on the common conventions of text layout in lines and vertical columns (as in newspapers), and squeeze out the wasted space that a more traditional node and links view of tree-structured data consumes for the links and spacing (see again Figure 2). Relaxing a conventional table's row alignment enhances the ability to read columns downward across node boundaries because there is neither extra whitespace nor an abrupt clip of the content to fit within a fixed height.

4.3 The Fluid Writer prototype

The Fluid Writer prototype has been implemented as a WWW application. The layout algorithm is written in Java, with a Perl front end. Although this simple user interface reduces our ability to offer a direct manipulation editing style, it does permit easy long-distance collaboration between developers and authors.

A basic set of interaction controls for unaligned treetables is shown at the top of Figure 4. The controls are divided into edit controls and view controls. The current edit controls permit the addition of child cells, editing of cell content by extraction of text and replacement, addition of merge cells, deletion of subtrees, and an "undo" to successively revert to earlier states. The current view controls are intended primarily to obtain views emphasizing particular paths. As in the Table Lens facility described by Rao and Card [24] for true tables, multiple foci are of interest, here to show several alternative versions of an authored paragraph or section.

² This difference in cell height across a row gives rise to the "unaligned" in the name "unaligned treetables." Aligned treetables have a related ability to facilitate reading hierarchical information, but they also make the level structure more prominent, which is not necessary here. They are also taller for a given set of text to display.

The view controls distinguish between those used to select focus cells (currently "add focus" and "delete foci") and those indicating the kind of display desired. Currently, given a set of focus cells, the standard display ("w/focus") shows only cells on paths passing through those cells. For example, Figure 4 shows a view of the treetable of Figure 3 with two different focus cells and associated focus areas.

Other displays distinguish focus paths only by highlighting ("w/o focus"), to assist in shifting foci. "Mini" displays, showing only initial fragments of cell content, are used for the same purpose for larger tables. These alternatives could be augmented by intermediate focus expansions showing full text for in-focus paths and indicative fragments for others.

It should also be noted that animation techniques would be useful in helping to maintain orientation during shifts in focus and in displaying layout changes due to editing operations. Wittenburg and Sigman [30] provide a detailed discussion of ways to use animation to maintain orientation in the context of multiple foci in a tree representation.

5. FLUID NARRATIVES

In parallel with the development of the Fluid environment by the first and third authors of this paper, the second author began exploring Fluid narratives from the perspective of literary theory, intent on studying how different kinds of hypertext structures affect how we write, and read, various kinds of narratives. In order to explore these issues, she is currently writing a mystery story in Storyspace and the Fluid Writer, focusing in particular on the narratological aspects of plot and closure. These two hypertext systems were chosen because they provide quite different options for an author, both at the level of detailed writing, and at the level of structuring and visualizing larger amounts of text.

5.1 Plot and closure in hypertext narratives

Hypertext is claimed to create a shallow, fragmented, reading behavior, its alluring links prompting the reader to scan the page and click on links rather than focusing on the (deep) reading of the text [15]. Interesting, and still only partially answered, questions are how this pertains to the writing, and reading, of hypertext narratives.

The point of departure for this part of the paper is the observation that a major part of existing hyperfiction might be said to belong to a category/categories that can be described as experimental, spatial, avant-garde – in other words, literary genres and aesthetic conventions that are often associated with a so-called post-modern aesthetics.

As a tool for creative writing, hypertext is considered to facilitate a spatial mode of aesthetics and narration, most obviously illustrated in many of the hyperfictions of the Eastgate school. Hyperfictions like *afternoon*, *Patchwork Girl*, *Victory Garden*, and *Figurski at Findhorn on Acid*, all written in Storyspace, are characterized by multiple narrative levels and points of view, shifting contexts and perspectives, digressions and other kinds of narrative and aesthetic devices working against a straightforward telling of a story. Their composition principles and narrative strategies are dominated by other mechanisms than a clear-cut plot starting out from a clearly stated exposition and an initiation of

the chain of events, continuing through a rising curve towards a climax, followed by closure and narrative completion.

Such plot-reliance and (more or less) straightforward unfolding of the action and events is prevalent in formulaic narratives – such as detective or mystery stories, romances, and fairy tales. Formulaic narratives place a heavy load on the reader's cognitive capacity, making for the modes of reading associated with escapism and so-called “ludic reading,” where we become fully immersed – “lost in the book.” [20][6][7][26].

An intriguing issue is whether, why, and to what extent, hypertext writing tools (or even hypertext itself) are incompatible with the telling of typically immersive narratives, as for instance, detective or mystery stories. Several have pointed to the lack of hypertext mysteries and detective stories [3][7][29],³ and some claim that an important aim for future hypertext authors should be to avoid the theorizing way of writing hyperfiction, and instead opt for more purely entertaining and straightforward narratives [9].

The transition from traditional narrative form to the new narratives supported by and made possible by digital media can be characterized as a move from linearity to multi-linearity – from the confines of print to the flexibility of the hypertext network [17]. It may seem that the very core of hypertext, its generic features so to speak (multi-linearity, network structure, the reader choosing her paths) collide with core features of narratives. Whereas Liestøl proposes some genres of computer games as examples of possible solutions, the issue remains largely unresolved in hypertext narratives.

In addition to a clearly discernible plot, and character development, the aspect of closure is crucial in our reading for pleasure. As Peter Brooks claims in his pivotal work *Reading for the Plot*, we read “in anticipation of retrospection,” [4] and this urge for the end, which will put the text as a whole in the “proper,” final, perspective, is our very motivation for reading [7][1][26][4]. However, in hyperfictions like *afternoon*, closure becomes largely a matter for the reader to decide – whenever she feels that main narrative tensions are resolved and puzzles explained [7][26].

Fluid hypertext provides several intriguing options for shedding light on these issues. In order to highlight specifics of Fluid hypertext for narrative purposes in general, and for plot-reliant narrative fictions in particular, comparisons with Storyspace (the authoring tool itself, as well as hyperfictions written in Storyspace) will prove valuable.

5.2 Reading Storyspace and Fluid narratives

George Stuart Joyce notes that a conventional hypertext network structure, such as that employed in Storyspace, poses two important problems for narrative: the problem of reader confusion, and the problem for the narrative to emerge from the unconstrained graph structure, which he calls a *rhizome*. Because the rhizome structure entails a lack of narrative progression, it provides too little support for developing character and plot – arguably two of the most important ingredients in narratives that typically provide pleasure/escapist reading [12].

Studies of hyperfiction reading confirm this assessment. There are by now quite a number of studies covering individual readings and analyses of hyperfictions, focusing on narrative and/or aesthetic issues. Empirical research on hypertext reading, however, has primarily concentrated on informational/non-fiction hypertexts. Although research on actual readers reading hyperfiction is still relatively scant, some studies have addressed readers reading hyperfictions written in Storyspace [18] [7][8][19]. Most empirical studies conclude that readers require considerable cognitive work when reading hyperfiction, often to such an extent that the confusion and frustration from lack of coherence and structure overshadows whatever pleasure is gained from the narrative (or even the ability to “find” the narrative – the plot – in the first place).

The Fluid Reader counteracts such reader confusion. The always-visible backbone provides a consistent and continuous (con)textual framework making the construction of the conceptual map, the graphical model of the fictional universe, much easier than in a hyperfiction like *afternoon*, in which the reader has to (re)construct both the fictional universe (in narratological terms, the *story*), the textual rendering of this universe (the *discourse*), and what can be termed the “discourse-as-discoursed,” i.e.,

the actual use and reading of the digitally stored text [...],
the creation of a path based on the selection and
combination of elements existing in a spatial and nonlinear
arrangement of nodes and links [16].

In the Fluid Reader, a coherent narrative is present and visible from the beginning, so that whatever exploration the reader chooses to indulge in (whatever alternative endings she chooses to activate), does not impede an overall narrative experience.

As a result, and in contrast with hyperfictions presented in Storyspace, the reader of a Fluid hyperfiction doesn't have to concentrate as much on merely understanding what is going on, and continuously establishing a coherence and context so as to obtain a good-enough understanding in order to continue her reading. This leaves the reader more capable of paying attention to detail, to concentrate more on subtle narrative and aesthetic devices that may enhance the total reading experience. And as dedicated readers know, details are of utmost importance for a successful mystery or detective story. As a result, the author can take advantage of the reader's cognitive surplus by providing paths to explore within a stable narrative context. Fluid hypertext thus combines the requirements for a closed narrative, a set plot and outcome, with the possibilities of explorations. Such a combination is clearly advantageous, because it gives the reader “a feeling of free movement and yet even with that free movement provide[s] a seamless story experience [...]. Moreover, it makes the reader want to repeat the same story to explore the other paths.” [23], cited in [12]

As Willerton claims, the experienced mystery reader is not primarily interested in speeding toward the solution, but takes pleasure in the repeated postponement of a desired end [29]. The mystery author can exploit the hypertext structure by creating “delicious delays,” for example by having hyperlinks lead to background information on characters, milieu, or clues to the mystery – or, indeed, red herrings or delectable distraction, experienced by the reader as aesthetic pleasures rather than annoying disruptions. As such, a Fluid mystery might provide what Marie-Laure Ryan calls “spatial immersion,” inviting us to

³ Eastgate has just announced a forthcoming mystery written in Storyspace: Chris Willerton's *Londale Hotel*.

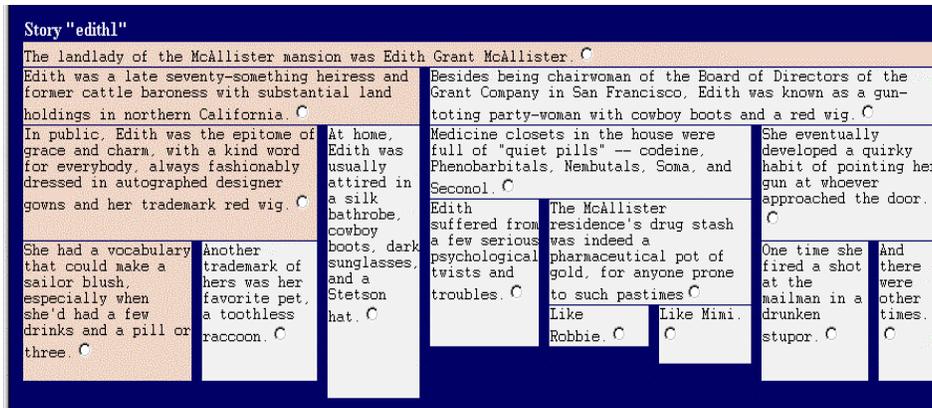


Figure 5. Character development in a Fluid narrative. The shaded nodes on the left edge of the treetable are visible initially, with choice points to access the second, third, and fourth paths (labeled preceding rightward). Choice points for the fifth and seventh paths appear with the fourth path; for the sixth path with the fifth; and for the eighth path with the seventh.

slow down the pace of reading and occasionally to linger on particularly interesting places [26].

5.3 Writing Storyspace and Fluid narratives

As a writing tool, Storyspace facilitates exploiting intricate and complex network structures on a large scale, easy manipulation of links and nodes, and possibilities of extensive linking (uni-directional and bi-/multi-directional). The different options for visualizing the structure also allow authors to obtain an overview and to keep track of relations between nodes at various levels. However, possible advantages for authors at a macro-structural level do not necessarily help authors write narratives that require local coherence and narrative development, or help readers in search of signs of character or plot developments.

Fluid hypertext, on the other hand, provides totally different prospects than Storyspace, for authors and readers. The axial structure of Fluid hypertext allows for a simultaneous evolving of a narrative along two diverging, but closely connected, axes: the “forward” axis of the backbone forming a top-level narrative, and the “downward/inward” axes of the alternatives (the branching nodes). In a Fluid narrative, the backbone, or top level of the text, can represent the forward movement of the narrative, whereas the branching nodes, the alternatives, can be used for elaborations of the narrative in different ways.

In the case of detective or mystery stories, character and plot development are crucial issues for which the Fluid Writer may provide better local support than Storyspace.

For instance, the author can use the top-level narrative for a first, general, presentation of each character involved in the mystery to be solved – a minimal description of the cast, alongside the presentation of the main events forming the mystery. In the Fluid Writer, it is easy to keep track of the relations between the top level narrative, and the branching nodes, at all times, because the text sections forming the top-level narrative span all branching nodes within. The alternative section endings, or branching nodes, can be exploited for presenting multiple focalizations and traits of a character, thus rendering each of the characters more nuanced – more finely grained – as the nodes branch. The alternative section endings can also be used to provide multiple points of view on a character, or on events and possible solutions, thus complicating the detection of the mystery, and maintaining and even

heightening the suspense so crucial in mystery/detective stories. See Figure 5 for an example of character development in a Fluid narrative.

At the same time, it is imperative to keep track of the relations (semantic, syntactic, narrative) between current and preceding nodes, and the relation to the backbone. This is made easy in the Fluid Writer, in that the author can see the preceding nodes in immediate proximity all the time. In Storyspace, by contrast, the author has to consider how to make nodes functional according to different and shifting contexts, depending on what path the reader has followed, what node was visited prior to the current, etc. (although the option of “guard fields” may to some extent ease this challenge). The Fluid Writer maintains the context, allowing the author to explore multiple writing at a detailed level within a coherent structure, without losing the overview.

Similarly, for plot development, the backbone may present the “shallow” version of the mystery, containing only the main events and presenting core information and the main mystery/-ies to be solved, without any background information or solution. At the subsequent level(s), in the branching nodes/alternative endings, the author can gradually elaborate on the events, providing background material that sheds light on the mystery and triggers the reader to explore deeper down in the narrative for additional clues.

The Fluid Writer is an authoring tool that supports a combination of narration (in the backbone) and description and additive material (in the alternatives/branching nodes); thus, Fluid narratives can provide a solution to the seeming lack of dynamic, plot-based narratives with a forward-moving story combined with occasional delay in the form of descriptions, additional details about characters, setting, plot twists, or events.

6. DISCUSSION

6.1 Animated form of generalized stretchtext

Fluid hypertext is related to Nelson’s stretchtext proposal [21], which allowed readers to control how much material is presented in a document via an extra “level of detail” throttle. Increasing the level of detail would cause gaps to appear between the phrases; additional words and phrases would pop into the gaps. The stretchtext concept continues in popularity, with recent

implementations, such as `pop`text [28], continuing to appear. However, unlike `stretch`text, Fluid hypertext also allows the reader to choose among different paths within the expanding and contracting text.

The Fluid Reader represents a successful implementation of an extension to `stretch`text that incorporates animation to smooth the reading experience and to help readers distinguish the added details from the previous summaries. Although the fine-grained, inline additions inherent in `stretch`text call out for animation (in fact, Nelson's description seems to suggest it), there are many details involved in creating a readable result (see section 3.2). For example, animated re-flowing of text typically distracts readers from the reading task. The Fluid Reader, with its animation and bending text, offers a concrete implementation that many readers have found usable, engaging, and enjoyable.

6.2 Extending the Fluid hypertext structure

The structure of Fluid hypertext most closely resembles that of the `replace`-buttons in the early Guide system [5]. Again, `replace`-buttons provide the opportunity for fine-grained, inline expansion of text. Unlike the `stretch`text concept, `replace`-buttons explicitly acknowledge that material at a higher level may disappear as new material is added. However, Fluid hypertext's choice points, which replace the content of several nodes with a new path (with possibly additional embedded choice points) provide yet another wrinkle on the replacement concept.

Guide also included several more traditional hypertext constructs: cross-reference links (which could either insert their contents inline or jump the reader to another location in the hypertext) and note links, which would pop up a short note.

The current version of the Fluid Reader operates on what Phelps calls a braided multi-linear structure [23]: one that has a linear sequence of central nodes that the reader experiences, with alternate ways of traveling from one of these central nodes to another. (Note that such a structure can refer either to the reader's possible paths through a traditional hypertext network or the structure itself.) Because the Fluid hypertext structure is actually a sequence of trees, we will call it a braided multi-tree. As we have noted earlier, this structure is not as general as a traditional hypertext network.

However, for greater generality several extensions could be made. First, links could be permitted within nodes, thereby allowing readers to jump from one multi-linear structure to another. This addition would permit authors to escape into a more conventional hypertext structure if the multi-tree structure becomes too confining. Also, following Guide, another form of link could allow another subtree to be expanded in place in the current tree.

6.3 Appearance controls for the Fluid Reader

The current Fluid Reader represents an initial step toward an environment that supports successful reading and allows some latitude for authors. We made certain choices in the Fluid Reader museum exhibit to provide a sense of playfulness that we thought would appeal to visitors. Naturally, other situations and stories may suggest other choices.

Authors can currently specify a variety of document characteristics, including fonts, text colors, margin sizes, and animation speeds. The images used for some or all of the choice points can be changed: distinct images could indicate possible

different kinds of content in that direction. Authors can also select whether a previous alternative disappears completely (as in the museum version) or remains, fading to a less-prominent color if desired. Although we are quite pleased with the text swoop as a way to guide the reader's eye to a new alternative, other line shapes, or even the simple method of inserting line breaks and placing the new text at the beginning of a new line, as in Guide's `replace`-buttons, could be placed under author control. We expect that the set of desirable author controls will become clearer as more authors write stories for this new medium.

6.4 Navigation aids in the Fluid Reader

Although the engagement of the museum visitors clearly demonstrates readability (even for young readers), ease of use, enjoyment, and possibly even an immersive quality, we expect that additional navigational aids analogous to those available in standard hypertext environments will be needed to support more serious reading of potentially longer narratives. The additional aids described here are being developed for the Fluid Reader.

Saved views. It can be quite difficult to remember how to find a particular passage encountered within the narrative. That is, which choice points must be opened to bring the passage into view again? Readers must be able to save the current view to permit returning to it later. In effect, this action saves the states of all choice points on the screen. Because a saved view is a large piece of text encompassing the content of many nodes, the reader may select a text region to highlight when it is displayed.

Breadcrumbs. Readers may also grow confused in a tangle of triangles: have I come this way before or not? Like colored text and underlines that show previously-traversed links, triangles can be recolored to show whether or not this choice point has been opened previously.

Completeness. A significant and unsolved problem in conventional hypertexts, the restricted nature of Fluid hypertext structures makes it possible to further recolor choice triangles to show that *all* nodes reachable from that choice point onward have been viewed. Reachability is easily computed for Fluid hypertexts because links (i.e., choice points) are not permitted to point to arbitrary destinations.

7. CONCLUSIONS

The Fluid Reader provides a new hypertext experience for readers. Its animated typography presents a new style of fine-grained hypertext with a continuously-visible backbone.

Our recent development effort in the area of Fluid Narrative hypertexts has resulted in improvements for both authors and readers. The Fluid Writer provides a new unaligned treetable visualization to allow authors to read easily along one or more paths in a Fluid Narrative hyperstructure. Navigational aids in the Fluid Reader support more serious reading activity.

The state of the art of hypertext narratives show the difficulty of authoring in hypertext while at the same time maintaining the coherence, plot line, and closure so pertinent to popular narratives like, for instance, mystery and detective stories. Fluid Narrative employs features that might expand the canon of hypertext narratives beyond the current state of more or less experimental, disorienting narratives, making it possible to write narratives that rely on the coherence and closure we savor in formulaic fiction, while at the same time exploiting the hypertext features for

narrative and aesthetic purposes, thus expanding the narrative experience for both writer and reader alike.

8. ACKNOWLEDGMENTS

The Fluid Reader was a collaboration between Rich Gold, Bay-Wei Chang, Jock Mackinlay, and the first author.

We thank David F. Vasquez for constructive discussions and valuable feedback on mystery writing, and for being an inspirational sounding board, and Marjorie Luesebrink and Maribeth Back for very interesting conversations. Jock Mackinlay and the anonymous reviewers provided thoughtful comments on earlier drafts.

Note: Some ideas in this paper have been presented in a less elaborated version at the Digital Arts and Culture Conference, Brown University, April 26-28, 2001 (Maribeth Back & Anne Mangen: "Lost in Fictional Space: Immersion by Animation and Sound. Experiments in Dynamic Fiction at Xerox PARC.")

9. REFERENCES

- [1] Aarseth, E. *Cybertext*. Johns Hopkins Univ. Press, 1997.
- [2] Back, M., M. Chow, R. Gold, A. Balsamo, M. Gorbet, S. Harrison, D. MacDonald, and S. Minneman. Designing innovative reading experiences for a museum exhibition. *IEEE Computer*, Vol. 34, No. 1, January 2001, 80-87.
- [3] Bernstein, M. No Mystery. <http://www.eastgate.com/HypertextNow/archives/Mystery.html> , 1997.
- [4] Brooks, P. *Reading for the Plot*. Harvard Univ. Press, 1984.
- [5] Brown, P.J. Do we need maps to navigate round hypertext documents? *Electronic Publishing*, Vol. 2(2), Dec 1989.
- [6] Douglas, J.Y., A. Hargadon. The Pleasure Principle: immersion, engagement, flow. *Proc. ACM Hypertext 2000*.
- [7] Douglas, J.Y. *The End of Books – Or Books without End?* Univ. of Michigan Press, 2000.
- [8] Gee, K. The ergonomics of hypertext narrative: usability testing as a tool for evaluation and redesign, *ACM J. Computer Documentation*, Vol. 25, No. 1, 2001, p 3-16.
- [9] Jensen, N.F. *Internet Hyperfiction*. Thesis, Department of English, Univ. of Copenhagen, Denmark, Feb. 2001.
- [10] Johnson, B. and B. Schneiderman. Treemaps: a space-filling approach to the visualization of hierarchical information structures. *Proc. 2nd IEEE Visualization Conference*, 1991.
- [11] Johnson-Eilola, J. *Nostalgic Angels: Rearticulating Hypertext Writing*. Ablex Publishing, 1997.
- [12] Joyce, G.S. *Authoring as Architecture. Toward a Hyperfiction Poetics*. Dissertation, 1999. <http://skyscraper.fortunecity.com/dns/689/hyperf.html>
- [13] Joyce, M. *afternoon – a story*. Eastgate, 1987.
- [14] Lamping, J., R. Rao. Laying out and visualizing large trees using a hyperbolic space. *Proc. ACM UIST'94*.
- [15] Levy, D. I read the news today, oh boy: reading and attention in digital libraries. *Proc. ACM Digital Libraries '97*.
- [16] Liestøl, G. Wittgenstein, Genette, and the reader's narrative in hypertext. In G. Landow (ed.), *Hyper/Text/Theory*. Johns Hopkins Univ. Press, 1994.
- [17] Liestøl, G. *Essays in Rhetorics of Hypermedia Design*. Univ. of Oslo, Dept. of Media and Communication (doctoral diss.), 1999.
- [18] Mishra, P., K. Nguyen-Jahiel. Reading print and hypertext fiction: Reader stance and its impact on meaning making. Paper presented at the *American Educational Research Association Conference*, April 1998. <http://punya.educ.msu.edu/PunyaWeb/pubs/print/hypertext/hypertext.html>
- [19] Moulthrop, S., N. Kaplan. Something to imagine: Literature, composition and interactive fiction. *Computers and Composition* 9, 1, 1991, p 7-24.
- [20] Nell, V. *Lost in a Book. The Psychology of Reading for Pleasure*. Yale Univ. Press, 1988.
- [21] Nelson, T. *Dream Machines / Computer Lib*. Microsoft Press, 1987.
- [22] Newman, P. Treetables and other visualizations for email threads. Xerox PARC Technical Report, Sept 2001. http://www.parc.xerox.com/istl/groups/hdi/papers/psn_emailvis01.pdf
- [23] Phelps, K. Story Shapes for Digital Media. <http://www.glasswings.com.au/modern/shapes/>
- [24] Rao, R., S. Card. The Table Lens: merging graphical and symbolic representations in an interactive focus+context visualization for tabular information. *Proc. ACM CHI'94*.
- [25] Robertson, G., J. Mackinlay, S. Card. Cone trees: Animated 3D visualizations of hierarchical information. *Proc ACM CHI'91*.
- [26] Ryan, M-L. *Narrative as Virtual Reality*. Johns Hopkins Univ. Press, 2001.
- [27] Storyspace. <http://www.eastgate.com/Storyspace.html>
- [28] Wadge, W., m.c. schraefel. Putting the hyper back in hypertext. *Intensional Programming II: Papers based on the 1999 Int'l Symposium on Lucid and Intensional Programming*. M. Gergatsoulis, P. Rondogiannis (eds). Singapore, New Jersey, London: World Scientific, 31-39.
- [29] Willerton, C. Structure problems in hypertext mysteries. *Proc. ACM Hypertext 2000*.
- [30] Wittenburg, K., E. Sigman. Visual focusing and transition techniques in a Treeviewer for Web information access. *Proc. IEEE Visual Languages '97*.
- [31] Zellweger, P., B. Chang, J. Mackinlay. Fluid links for informed and incremental link transitions. *Proc. ACM Hypertext '98*, 50-57.
- [32] Zellweger, P., S.H. Regli, J. Mackinlay and B. Chang. The impact of Fluid Documents on reading and browsing: An observational study. *Proc. ACM CHI 2000*