Proceedings of UIST'98 (San Francisco, CA, November 1-4, 1998), p 123-132.

A Negotiation Architecture for Fluid Documents

Bay-Wei Chang, Jock D. Mackinlay, Polle T. Zellweger, Takeo Igarashi[*]
Xerox PARC
3333 Coyote Hill Road
Palo Alto, CA 94304 USA
{bchang, mackinlay, zellweger}@parc.xerox.com, takeo@mtl.t.u-tokyo.ac.jp

ABSTRACT
The information presented in a document often consists of primary content as well as supporting material such as explanatory notes, detailed derivations, illustrations, and the like. We introduce a class of user interface techniques for *fluid documents* that supports the reader's shift to supporting material while maintaining the context of the primary material. Our approach initially minimizes the intrusion of supporting material by presenting it as a small visual cue near the annotated primary material. When the user expresses interest in the annotation, it expands smoothly to a readable size. At the same time, the primary material makes space for the expanded annotation. The expanded supporting material must be given *space* to occupy, and it must be made *salient* with respect to the surrounding primary material. These two aspects, space and salience, are subject to a negotiation between the primary and supporting material. This paper presents the components of our fluid document techniques and describes the negotiation architecture for ensuring that the presentations of both primary and supporting material are honored.

KEYWORDS: fluid user interfaces, fluid documents, negotiation architecture, scaling, zooming, focus+context, annotation

INTRODUCTION
Documents often supplement their primary information with related supporting material. For example, Shakespeare plays typically include supporting material to explain the archaic English and the obscure cultural references for the modern reader. Annual reports from corporations often include footnotes describing the details behind a profit or loss. Schematics and architectural plans include references to technical specifications that must be satisfied to produce appropriate yields or structural integrity.

Various techniques have been developed for combining primary and supporting information. Paper documents use well-known conventions such as footnotes, wide margins, endnotes, cross-references, and the like. These methods disrupt readers because they must find distant supporting material and then return to their prior reading location. More recently, hypertext systems have introduced more direct ways of connecting primary and supporting material. In a typical hypertext application, underlines indicate links and mouse clicks permit automatic transitions to supporting material. However, these transitions can also be disruptive because they remove the reader from the context of the primary material: a new page containing the supporting material may replace the primary material, a pop-up window may partially occlude it, or a new window containing supporting material may appear at some distance from the link location.

In contrast, our *fluid documents* provide lightweight, contextual, animated access to supporting information, allowing the reader to fluidly shift attention from the primary material to the supporting material.

Consider Figure 1 and Figure 2 for an example. In Figure 1, underlines act as visual cues to signal the existence of supporting material. The reader is indicating interest in a cue by moving the mouse pointer over it. In Figure 2, the supporting material has expanded to a position below the cue, while the primary material has moved apart to make room for it. The presentation of the supporting material is animated, allowing the reader's eye to follow the text growing smoothly from the underline to its fully expanded form.

To generate the presentation in Figure 2, the primary text yields space near the annotated text ("her maid") to the annotation ("the moon…"). In addition, the annotation is displayed in a contrasting font and style, making it more salient. How the primary material makes space and how the supporting material presents itself saliently comprise the central choices in a fluid document. In our fluid documents, these choices are arrived at through a process of *negotiation* between the primary and supporting material.

Fluid documents have several advantages over prior approaches. They promote engagement with the primary

[*] University of Tokyo, H. Tanaka Lab, Bldg 13, Electrical Engineering, 7-3-1 Hongo, Bunkyo-ku, Tokyo, 113, Japan

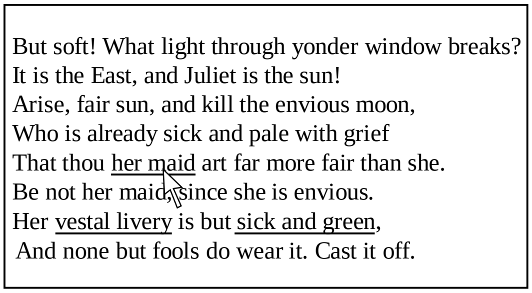

Figure 1. Excerpt of *Romeo and Juliet* with underline cues indicating annotations.

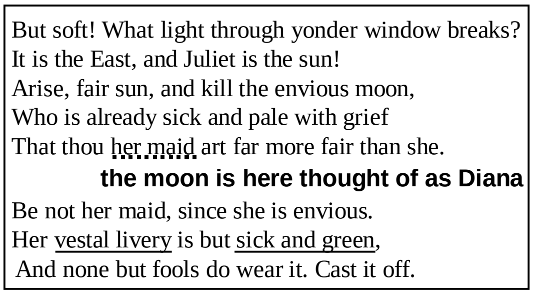

Figure 2. An animation expands the annotation fluidly, moving the surrounding text apart.

material by reducing the distraction of supporting material until it is of current interest and maintaining the context of the primary material while the supporting material is presented. They reduce visual saccades and promote comparison by placing supporting material close to its primary material. They reduce cognitive disruption by animating the presentation of the supporting material. Fluid documents can help readers interactively control their focus on the material to support their current reading goals.

We have implemented several prototype fluid document applications: a page-oriented document reading application that supports fluid annotations, both authored and user-created, a fluid hypertext browser [19], and a spreadsheet editor that uses fluid user interface techniques to show formulas and relationships among cells [9]. These prototypes have led to the implementation of a system that manages the fluid document negotiation process. After a discussion of related work, this paper describes fluid document interfaces and then explains the negotiation architecture that supports the interaction.

RELATED WORK

The development of documents that dynamically change to incorporate supporting information has been a history of increasing graphic power. Modern graphics hardware enables a rich set of graphical techniques such as scaling, shading, and the like that can be used to create such dynamic documents. Scaling and interspersing text [7, 10, 12, 15] and overlaying objects upon one another [5, 8, 18] serve to increase the amount of information a screen can display. A recent provocative example is the Java-based NiF (News in the Future) Elastic Catalog [14], which creates a dense display of overlapping text that provides visually rich, complex access to a hierarchical table of contents. Our goal here is to develop an architecture that can be used to marshal this emerging graphical power to create effective dynamic documents.

The seminal research relating to documents that dynamically change their layout is Furnas's generalized fisheye views [6], which use a degree of interest function that takes into account both spatial proximity and overall importance of a focal element to suppress other parts of the display. His work led to various focus+context techniques [11] that 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. Fluid documents draw heavily on this research tradition. However, focus+context systems deal primarily with uniform information spaces whose elements are related by spatial proximity, and thus can be effectively distorted in a geometric fashion. In contrast, our work deals with multi-layered information spaces whose elements may have both spatial and explicit semantic relationships. The negotiation architecture described in this paper has been expressly designed to support fluid access to many different types of supporting information and many different strategies for achieving space and salience.

Significant similarities to our work also exist in Zooming User Interfaces, which grew out of the Pad/Pad++ systems [2, 16]. Interactive pan and zoom can be used to place supporting material as small elements close to the corresponding primary material. A drawback 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. The result is similar to hypertext techniques where the reader must remember the relationship between the supporting material and the invisible primary material. In contrast, our goal is to maintain the context of the primary material by inverting the zooming paradigm: we scale individual objects rather than zoom the entire view. However, this inversion creates a problem not addressed by the Pad++ research, namely, how to distort the primary information to accommodate the supporting information. Our negotiation architecture embodies an explicit protocol that is used by the primary and supporting objects to negotiate both space and salience.

A non-document example of the object scaling approach is the Continuous Zoom system [1], which deals with a large network depicted as a collection of non-overlapping rectangles connected by links. When the user makes a rectangle grow, the other rectangles in the network shrink using a fisheye algorithm that takes into account the network proximity.

Finally, a fluid document must rely on some kind of algorithm for document layout. At one end of the spectrum are knowledge-based techniques [13]. However, they typically run slowly, which is unacceptable for interactive applications. Furthermore, the conflicting goals of the primary and supporting information can create cycles that might keep a general-purpose knowledge-based algorithm from terminating. At the other end of the spectrum are constraint algorithms that can describe and adjust the fine-grained layout of a dynamic document. Borning [4] and others have developed many techniques for solving constraints efficiently and avoiding problems with cycles. However, fluid documents require knowledge about space and salience that go beyond simple layout constraints. Sophisticated strategies for adjusting fluid documents that consider interaction history and the current state of the dynamic document are even more difficult to describe with a system of constraints. Therefore, we adopt here a middle ground between these two endpoints. Our negotiation incorporates detailed knowledge about space and salience as well as a negotiation protocol that is guaranteed to terminate.

FLUID DOCUMENTS

Fluid documents mediate the relative salience of their primary and supporting material, allowing the reader to quickly and smoothly change focus to the portions of the document that she is interested in at a given moment. Initially, the primary material is the focus, using most of the space and capturing most of the user's attention, with supporting material limited to graphical cues. When the user focuses on the supporting 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, or fade to a lighter shade. The primary and supporting material choose their presentations via a negotiation process, taking into consideration their graphical properties and capabilities.

This approach can be divided into four steps:

1. A visual cue is placed near annotated primary material (referred to as the *anchor*) to indicate the existence of supporting material.
2. Readers indicate interest in a cue or anchor, triggering the supporting material to present itself more fully. For lightweight use, readers indicate interest by simply moving the mouse pointer over the cue. When the pointer dwells on the cue for more than a fraction of a second, the supporting material expands.
3. Behind the scenes, the primary and supporting material negotiate their relative graphical appearance, including their salience and the space in which to display the supporting material. Our goal is to display the selected supporting material near its associated primary material while minimizing the visual disturbance to the primary material.
4. Finally, the primary and supporting material animate smoothly to their new states. The supporting material is now visible and readable, and typically close to the primary material that it annotates.

We now discuss each step in turn.

Visual cue to supporting material

Visual cues in fluid documents can build upon various conventions from paper and hypertext documents to indicate supporting material. They can be textual (e.g., asterisks or footnote numbers), graphical (e.g., underlines), or utilize properties of the primary material (e.g., making anchor text boldfaced). In addition, cues can indicate more than the simple presence of supporting material. For example, different kinds of cues can indicate different kinds of supporting material. Cues can even preview the actual content of the supporting material via small text or images, providing additional information to help the reader decide whether to expand it. We have found that on a 282 dpi display, text in a 3 or 4 point font is an effective visual cue [19].

Lightweight UI for indicating interest

A goal of the fluid document interface is to develop a lightweight user interface that has the feel of eye saccades on a paper document. We use the pointing device as a surrogate for the reader's focus of attention. When the mouse pointer is over a cue, the supporting material starts to grow to a predetermined size. When the pointer moves away, the supporting material shrinks back down to its cue. If the reader clicks, the supporting material remains expanded after the mouse pointer leaves, allowing the user to focus on multiple annotations at once.

Adjusting the primary and supporting material

In a fluid document, the primary and supporting material adjust their displays in concert to allow supporting material to be shown in a salient way while maintaining the primary material for context and comparison. Bertin's graphical vocabulary suggests a wide range of potential techniques for adjusting the primary material to make room for supporting material, including position, color/saturation, size, shape, orientation, texture, and marks [3].

Size is the main graphical dimension that fluid documents manipulate in order to show supporting material. In our text annotation and hypertext applications, for example, cues are always small and unobtrusive, and they expand to become the textual annotation. There is usually no existing space reserved for the supporting material to show itself at full size—if there were, the supporting material might as well be displayed there to begin with—so the primary material must arrange for some space to be given to the expanding supporting material. Because the primary and supporting material each have their own preferences for space and positioning, they engage in a negotiation to determine how much space is allocated and where it is located.

The space allocation may have implications for the presentations of the primary and supporting material. For example, the two may need to contrast in order to emphasize one or the other, or to more clearly distinguish them from each other. Since each party also has its own preference for how to display itself, these graphical changes are also the result of negotiation. We discuss the negotiation architecture in detail later in the paper.

The following discussion describes some of the fluid document techniques that we have explored for text documents with textual annotations. Each technique is a set of design decisions for displaying annotations in the context of the main text that they support. The manipulated graphical dimensions that are illustrated in these examples are position, color, and marks.

Position: Block movement and interline compression

The simplest positional technique for adjusting primary material moves it away as blocks to make space, as shown earlier in Figure 2. One advantage is that the supporting material stays close to the primary material and the primary material is not distorted. However, a fixed display size with no margins would result in cropping the primary material, losing some of the context.

A related technique for use in fluid text documents moves the position of each line of text individually, compressing the interline space and allocating the savings to the annotation. This approach still maintains the annotation near the primary information but at the expense of somewhat distorting the primary text body. Our trials suggest that compressing the interline spacing of all primary text lines equally is less obtrusive than selectively compressing lines either near to or far from the annotation.

When the supporting material is the same type of object as the primary material (primary text and textual annotation, for example), it may be difficult to distinguish between the two when the supporting material has expanded. A fluid document principle is to enhance the salience of the supporting material by altering some of its graphical characteristics. In a textual annotation situation, we change the font face, font style, and/or text color.

Marks: Margin callout

Adding graphical marks helps to show the relationship between the primary and supporting material when they cannot be adjacent. The *margin callout* technique avoids altering the primary material in order to display supporting material. Instead, it places the supporting material in existing white space on the page, for example, the margins. When the reader indicates interest in a cue, a line extends out from the cue to the nearby side margin, and the supporting material expands at that location (see Figure 3). Because the callout line travels in between text lines, the primary material is neither moved nor altered. This technique can allow a good amount of supporting 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 supporting material effectively draw the reader's eye to the supporting material, and then back again when the callout line shrinks back to the cue.

But soft! What light through yonder window breaks?
It is the East, and Juliet is the sun!
Arise, fair sun, and kill the envious moon,
Who is already sick and pale with grief
That thou her maid art far more fair than she.
Be not her maid, since she is envious.
Her vestal livery is but sick and green,
And none but fools do wear it. Cast it off.

The moon is here thought of as Diana

Figure 3. Margins can also be used for supporting material.

Color: Overlay

Another technique that preserves the layout of the source document is the *overlay* (see Figure 4). An overlay expands into the space below the anchor, as in the interline compression technique. However, the source material in that space does not move away; rather, it fades in color. The difference in darkness of the two layers of text, as well as other differences like size and font, can make the supporting material readable while still preserving some degree of readability of the underlying source text.

But soft! What light through yonder window breaks?
It is the East, and Juliet is the sun!
Arise, fair sun, and kill the envious moon,
Who is already sick and pale with grief
That thou her maid art far more fair than she.
Be not her maid, since she is envious.
the moon is here thought of as Diana
Her vestal livery is but sick and green,
And none but fools do wear it. Cast it off.

Figure 4. When space is at a premium, supporting material can overlay the primary material.

Combining techniques

These techniques can be combined, as shown in Figure 5. In this example, the interline compression technique did not yield enough space for the annotation, so the overlay technique was used for the remainder of the second line. Combining techniques gives the annotation as much empty space as possible and still allows it to be readable where it must overlap the primary text.

But soft! What light through yonder window breaks?
It is the East, and Juliet is the sun!
Arise, fair sun, and kill the envious moon,
Who is already sick and pale with grief
That thou her maid art far more fair than she.
the moon is here thought of as Diana,
the goddess and patroness of virgins
Be not her maid, since she is envious.
Her vestal livery is but sick and green,
And none but fools do wear it. Cast it off.

Figure 5. This example combines interline compression and overlay.

Interactions with existing expansions

Since fluid documents allow more than one annotation to be in an expanded state at the same time, newly expanding annotations may interact with previously expanded annotations in several ways. Figure 6(a) shows in schematic form how several expanded annotations may look on a page.

Similarly, Figure 6(b) illustrates a *nested annotation*. Supporting material may have its own supporting material, creating a situation of nested expansions. When creating space for nested annotations, negotiations between each pair of primary/supporting material may be required.

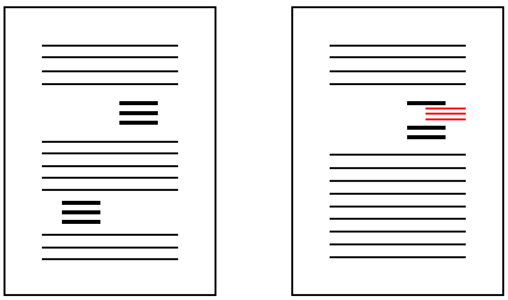

(a) Multiple (b) Nested

Figure 6. Multiple expansions and nested expansions.

In Figure 7(a), a central annotation is being expanded into the margin between two previously expanded callouts. To make room for the new callout, existing callouts must be moved.

Finally, if the primary object has already created some space for an existing expanded supporting object, it may have space left over in which to direct new expanding supporting material. In Figure 7(b), the primary grants the newly expanded annotation (on the left) some space in between lines that are already pushed apart. Since the annotation is no longer near its anchor, a connection line connects the annotation to its anchor.

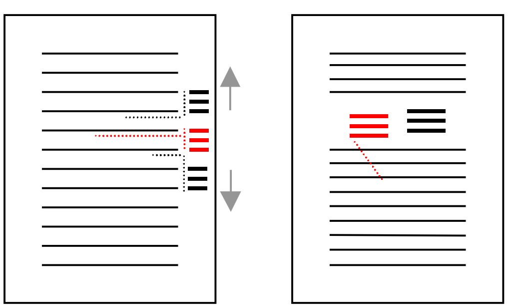

(a) Move other (b) Reuse

Figure 7. Moving annotations and reusing space created for previously expanded annotations.

Animated transitions

In our fluid document 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 textual or graphical material grows, shrinks, and moves, the animation helps the eye follow the action. For example, in the interline compression technique, the supporting material begins as a tiny object and grows to a larger, readable size. The animation brings the eye from the cue to the expanded supporting material, and then back to the cue and the primary material as the supporting material shrinks down again. Meanwhile, the surrounding primary text moves apart to make room for the expanded supporting material—in this case, the purpose of the animation is to provide constancy of the textual objects in the viewer's periphery to allow him to concentrate on the newly-presented supporting material. Animation serves both to direct the reader and to minimize surprising changes on the screen.

Our animations are designed to be short but perceptually appropriate, typically less than a half second in duration. Long animations can disrupt the work process, whereas short animations occur during the time a user normally takes to react to a transition [17].

A FLUID NEGOTIATION ARCHITECTURE

We have designed and implemented an architecture for supporting fluid documents. The architecture takes into account the spatial extent and graphical presentation needs of both the primary and supporting material. (For this discussion we use the terms "primary object" and "supporting object," or in general, "fluid objects.") When a supporting object is to be expanded, a structured negotiation session occurs between it and the primary object, resulting in a cooperating visual presentation of both.

The process is one of "negotiation" because neither participant in the exchange has complete control over the layout and presentation decision. The primary object cannot dictate how the supporting object is displayed, nor can the supporting object demand a particular region into which to expand itself. Instead, the participants make requests and proposals based on the information they have of their own layout and presentation needs. The resulting agreement allocates *space* for the supporting object, as well as ensuring that its presentation makes it *salient* with respect to the primary object.

We designed an architecture based on negotiation between fluid objects because the back-and-forth nature of negotiation keeps complex decision-making within the objects themselves. Consider that, in general, primary objects may have many different strategies for allocating space for supporting objects, and supporting objects may have many different strategies for presenting themselves, depending on how much space they receive or other graphical restrictions imposed by the primary. Each object has idiosyncratic algorithms for choosing the best strategy, based upon the graphical requests made by the other object. In contrast to a scheme in which both sets of strategies, the context, and the algorithms for selecting strategies are all exposed to a central decision-making engine, negotiation permits strategies and selection algorithms to be encapsulated within each object. At each stage of the negotiation, a relatively small amount of information is communicated from one participant to the other. The information is then used to generate a reply, which may be the result of an arbitrarily complex internal algorithm.

Thus, a negotiation architecture allows fluid objects to be flexible in the number and kinds of strategies they consider, and simplifies the system by hiding the strategies and selection within each object. This is especially important when cascading negotiations are taken into account. Because supporting material such as an annotation may have its own supporting material (and so on), expanding a nested annotation may result in a series of negotiations up a chain of objects if each primary cannot adequately satisfy its supporting material's space request without expanding itself. In this case the localized nature of each negotiation significantly simplifies the situation: rather than the system attempting to simultaneously resolve the selection of *n* mutually compatible strategies, each pairing of objects in the chain simply negotiate between themselves.

The following sections describe the negotiation architecture. First we introduce the framework for hierarchical graphical characteristics protocols which allows fluid objects to gracefully handle requests from other fluid objects that present different or more complex graphical characteristics. This enables a component model of fluid documents in which different fluid objects (components) can co-exist without being designed specifically to interact with one another.

Then we discuss the two intertwined subjects of negotiation, space and salient presentation. Finally we step through the negotiation sequence. The examples will be in the realm of text annotation, one of the areas we have explored most thoroughly. However, the architecture is designed to support non-textual fluid objects, including such things as embedded property sheets for text and figure editing, annotations on diagrams, map details, and others. We are currently implementing several of these other fluid applications.

Graphical protocols

Supporting material in a fluid document must be able to communicate its presentation capabilities to its primary material. For example, if a primary object plans to overlay a supporting object on itself, it must be able to request that the supporting object make itself relatively darker in order to provide enough contrast. Similarly, if a body of text plans to place a textual annotation in between its lines, it should be able to determine whether the annotation can display itself in a contrasting font face, size, or color, in order to be visually distinguishable.

Different kinds of objects (text, annotations, hypertext glosses, images, schematics, maps) have different graphical capabilities. The particular language for communicating a fluid object's graphical features, its *graphical characteristics protocol*, is specific to a class of fluid objects. For example, our textual annotations understand graphical characteristics that include position, size, font (face, style, size, and color), background color, and border. In addition to specific values, these characteristics can be expressed as quantitative ranges (font size: 10 to 14 point) or qualitative ranges (font face: sans serif). Specific values describe the actual or proposed state of the object. Ranges describe acceptable states of the object, either to describe a supporting object's capabilities, or to request the supporting object to conform to certain graphical conditions.

Some protocols allow the space to be specified as an area, with minimum width and minimum height. Textual annotations, for example, can reformat themselves by reflowing across multiple lines. Requesting an area rather than a specific width and height allows the primary more leeway in selecting an effective space.

The protocol must be known to both parties of the negotiation. In the simplest case, a fluid object acting as a supporting object in a negotiation need only know the protocol for its graphical type. However, any fluid object that may embed supporting material may be confronted with a foreign protocol.

To handle these situations, protocols are arranged hierarchically, where protocols closer to the root of the hierarchy are strict subsets of those descending from it. All protocols descend from a common root, which is the basic protocol that all objects will obey: size and, implicitly, location.

A fluid object confronted with a foreign protocol simply follows the protocol's parent chain until it finds a protocol it knows about. It then communicates using this simpler protocol. In the case where the root protocol must be used, only the basic characteristic of size will be negotiated.

The protocol hierarchy allows new types of fluid objects to be created and used in existing fluid documents without re-coding. A new fluid object may have graphical parameters that the existing document is not informed about, but the system is likely to still settle upon reasonable presentations.

Strategies for making space and for presentation

Fluid objects must be able to give space to a supporting object that wishes to expand. Each primary object provides its own strategies for coming up with space. The strategies are not described declaratively; rather, they are procedures that consider how to alter the layout of the fluid object in order to accommodate an expanding supporting object.

Space-making strategies can be grouped into four categories:

1. *Block movement.* The primary object divides into sections, which rigidly move apart. The supporting object then can expand into the space in between the sections.
2. *Deformation.* The primary object divides into sections, which deform in some manner to create space in between the sections. (Alternatively, the primary object simply deforms itself as a whole.) The sections do not increase the size of the bounding box of the primary object by moving apart. The supporting object expands into the space in between the sections, or into the vacated space.

   Among the many strategies that fall in this category are interline compression, shrinking the entire text body or image, curving lines into a "bulge" around the expanding supporting object, and squashing the characters of the text itself.
3. *Overlay*. The primary object does not move out of the way; instead, it permits the supporting object to use some region *on top of* itself.
4. *Outside allocation.* The primary object negotiates for space outside of its bounding box for the supporting object. Other objects must agree to make space for the expanding object by way of block movement, deformation and/or overlay. Margin callout is an example of this strategy.

The possible space-making strategies each fluid object has are ordered, from the most preferable strategies (because they result in the most aesthetic or understandable layouts) to the least preferable. Each strategy has limitations on the maximum size it can accommodate in a given situation. During the negotiation phase, the fluid object evaluates each strategy in turn until it finds one that will satisfy the space requirements requested by the expanding supporting object.

For example, a text block may choose to move apart to allow an annotation to expand. However, if the annotation is very large, there may not be enough top and bottom margin space for the text block to move into. In that case, the text block may choose another strategy, such as allowing the annotation to be overlaid upon it.

In addition, as shown in Figure 5, strategies across and within the categories may be composable. The system tries each of the strategies singly and if none of them are satisfactory, it will attempt to compose strategies. When composing strategies, the system keeps track of how much space each strategy would provide. Then it seeks the remaining space requested as it evaluates the next strategy. The ordering of strategies is also used to order the composition of multiple strategies.

In addition to determining space, each strategy also dictates the kind of presentation needed by the supporting object in order to ensure its salience (recall the example of a supporting object being required to be darker than the primary object when it is overlaid on the primary). In order to respond to these requirements, supporting objects have strategies for presenting themselves. For example, a textual annotation may prefer to display itself in black, but when requested to contrast with a black main text body (communicating via a graphical protocol), it can instead display itself in red.

The amount of space allocated to a supporting object will also affect its presentation choice. A textual annotation may be able to scale itself to fit the amount of space available, but if the space is particularly small, it may then fall back to showing just a summary phrase or its first few lines. This behavior is similar to semantic zooming [2]. Just as the primary object chooses an appropriate strategy for making space, given the supporting object's request and the current layout situation, the supporting object chooses an appropriate strategy for presenting itself, given the primary's proposed space allocation and presentation guidelines.

Currently, the rules for selecting salient presentations are embedded in the primary object's strategies; in the future it could be interesting to pursue a "salience oracle" approach in which such rules are centralized and customizable by both implementors and end users.

THE NEGOTIATION

Negotiations are encapsulated in a Negotiation object. This object carries all the information that transpires during the negotiation, including proposals for space and presentation, as well as the information required to construct the coordinated animation once the negotiation completes. Negotiation objects are retained after the negotiation for use in future negotiations (if the primary must modify the expanded supporting object) and for reverting the expansion when the supporting object is no longer of interest.

The central elements of a Negotiation are Proposal objects. These carry the space and presentation requests. Specifically, a Proposal bundles up two graphical characteristic objects, or GChars, which are instances conforming to a particular graphical protocol. (Although space is a special consideration of the negotiation, it happens to be represented along with other graphical characteristics in the GChar.)

One of the GChars in a Proposal is a SpecificGChar, which describes a specific graphical state of the object. The other is a RangedGChar, which describe the range of graphical presentations that the object can take on—for example, 8 to 14 point, bold or italic or plain, sans serif. A Proposal thus expresses both a preferred state and the permitted leeway to deviate from that preferred state.

Step 1: Initial proposal (supporting object)

When the user indicates interest in the supporting object, the negotiation is initiated. The supporting object contacts its primary, indicating that it would like to expand. It indicates its graphical protocol and submits a Proposal based on that protocol. The Proposal's SpecificGChar describes its *a priori* preferred amount of space and its preferred presentation. The RangedGChar describes the range of sizes and presentations it can display.

Step 2: Guidelines (primary object)

The primary object first determines which graphical protocol it will use to communicate with the supporting object (see the earlier section on graphical protocols). Then, using the space request, the primary object performs a preliminary determination of where the supporting object should expand and how it should present itself. To do this, the primary tries each of its strategies in turn, settling upon the one that can satisfy, or best satisfy, the supporting object's space request. This process results in a maximum amount of space that can be given to the expanding supporting object.

Since in general a supporting object does not know anything about its primary's presentation, it does not initially know if its expanded form will present an effective contrast to the surrounding primary information. The primary object is therefore in charge of communicating guidelines to the supporting object regarding desired graphical characteristics. The primary narrows the range provided by the supporting object to what it determines would be effectively salient.

The primary then returns to the supporting object a Proposal that includes these guidelines, including the maximum space available, as the RangedGChar. Our policy is to have the primary keep the guidelines as open as it can given its own layout constraints, in order to allow the supporting object the maximum amount of flexibility in choosing its own presentation. In addition to the guidelines, the primary creates a suggested presentation as the Proposal's SpecificGChar, taking into account the supporting object's preferred presentation.

Step 3: Presentation strategy choice (supporting object)

In the initial proposal, the supporting object offered its preferred presentation. At this stage, it may have to try alternate presentation strategies in order to conform to the guidelines received from the primary. Because these guidelines are a subset of the supporting object's initial proposal, it is likely that a conforming presentation can be found. Most importantly, the chosen presentation must fit within the maximum space allowance. If possible, it should also be close to the suggested presentation. This final graphical presentation, along with the precise amount of space needed by the expanded object, is communicated in a SpecificGChar to the primary object.

Step 4: Space-making strategy choice (primary object)

Using this information, the primary makes a final choice of space-making strategy.

Now that the negotiation is complete, the primary and supporting objects place animation specifications into an Animation object in the Negotiation object, and the animation to the newly expanded state is started. The single Animation object ensures that the various changes required to expand the supporting object and alter the primary object will be coordinated and smooth.

An example negotiation

The following example illustrates the negotiation process. The primary object is a body of text, and the secondary object is a textual annotation. Initially, the display simply shows the text body, with an underline to indicate the presence of an annotation:

Arise, fair sun, and kill the envious moon,
Who is already sick and pale with grief
That thou her maid art far more fair than she.
Be not her maid, since she is envious.

Step 1 (supporting). When the reader moves the cursor over the anchor, the negotiation begins with the annotation requesting space in which to expand. It calculates its space request based on its preferred presentation, 12pt Times. The annotation sends the following initial proposal to the text body, along with the ranges of font face, size, and color that it can display:

SpecificGChar

size	170x28 pixels
font face	Times
font size	12pt
color	black

RangedGChar

size	>= 115x10 pixels
font face	*<any>*
font size	>= 8pt
color	*<any>*

(For clarity, we only show a subset of the graphical characteristics within the annotation's graphical protocol.)

Step 2 (primary). In this example, the text body has two strategies for making space: moving its lines apart rigidly into the margins, or compressing the interline spacing. In both cases, the annotation is to be displayed under its anchor, in between the text body's lines, without any overlap.

On trying its first strategy, moving rigidly, the text body determines that a space of 200x22 pixels can be created. This does not quite meet the annotation's preferred size. Since it does exceed the annotation's minimum size, the text body could choose to pursue this strategy. However, it has another strategy it can try.

The second strategy, interline compression, can yield a space of 200x25 pixels. This still does not meet the annotation's preferred size, but it is closer. The text body chooses to pursue this strategy. In order to provide contrast with the text body's font (12pt Times), it requires that the annotation make itself a sans serif font that is at least 10pt, and preferably Helvetica at 12pt, colored red. Thus the text body returns the following proposal to the annotation:

SpecificGChar

size	170x25 pixels
font face	Helvetica
font size	12pt
color	red

RangedGChar

size	<= 200x25 pixels
font face	sans serif
font size	>= 10pt
color	*<any>*

Step 3 (supporting). The annotation considers this proposal. It can display itself in Helvetica and in red, but to fit within the maximum size allowed, the annotation must make itself 10pt. Using these characteristics, it calculates the actual size it will need, 150x20 pixels. The annotation then constructs a final graphical specification that it sends to the text body to confirm its choice:

SpecificGChar

size	150x20 pixels
font face	Helvetica
font size	10pt
color	red

Step 4 (primary). The text body then makes its final choice of strategy. In this case, since the final size of the annotation is compatible with the first strategy of moving rigidly (which yields space up to 200x22 pixels), it actually chooses to use this preferred strategy. The animation is initiated, resulting in the following new expanded state:

> Arise, fair sun, and kill the envious moon,
> Who is already sick and pale with grief
> That thou her maid art far more fair than she.
> the moon here is thought of as Diana,
> the goddess and patroness of virgins
> Be not her maid, since she is envious.

Alternative step 3. Consider a variation of the above example in which the text body requires that the annotation use a minimum font size of 12pt (rather than 10pt). Since displaying itself at 12pt Helvetica will cause the annotation to exceed the size limits set by the text body, it considers other strategies. One such alternative strategy is to display only its first line, which results in a size of 187x13 pixels. This is within the guidelines stipulated by the text body. The expanded annotation then looks like the following:

> Arise, fair sun, and kill the envious moon,
> Who is already sick and pale with grief
> That thou her maid art far more fair than she.
> the moon here is thought of as Diana…
> Be not her maid, since she is envious.

As these examples show, the negotiation is not merely an exchange of fixed layout needs. Both the primary and the supporting objects may select from multiple strategies in order to meet each other's mutual requirements.

Discussion

It is possible that the supporting object cannot accept the space offering and its attendant graphical conditions, and must then reject the offer. However, this case is intended to happen only rarely, and usually does not occur because the primary's strategies are designed to be permissive rather than restrictive. Thus, although the architecture allows for the possibility of quite specific demands by the primary and the supporting objects, our philosophy is that conditions should be generally broad and easily accepted by the fluid objects. We give up guaranteed success of the negotiation, but gain a simpler negotiation model that still allows flexibility in the kinds of strategies the participants can employ.

If the supporting object must reject the offer—for example, if there are many other already expanded supporting objects taking up space, with no way to make them smaller—the fluid document can handle the failure with some extreme fallback behavior, for example, opening up the supporting object in a new window.

The negotiation is not symmetrical—by design, the negotiation architecture favors the primary object. In fact, the primary object is able to make its salience guidelines proposal based on both its own needs and the supporting object's capabilities, while the supporting object is not directly informed of the primary's presentation. However, the general grounds for negotiation are set by the supporting object, and the supporting object still has great leeway in choosing conforming presentations. This asymmetrical design mirrors our notion of fluid documents in which the primary material is perceived as defining a document to which supporting objects contribute in localized fashions. In addition, and most importantly, having only the primary object propose the salience guidelines simplifies the negotiation and limits the number of negotiation rounds.

In fact, the negotiation architecture is guaranteed to terminate because it is structured as a four-step process. The first steps of the negotiation narrow the conditions of the RangedGChars, and concrete choices are mandated at steps 3 and 4. This architecture does limit the possibilities of the negotiation—in particular, if the objects have radically different strategies, they will not be able to follow a different course if the initial negotiations "go awry" by not settling on a mutually agreeable amount of space or graphical presentation. To allow that kind of extended negotiation, however, would require an iterative negotiation sequence that would likely require complex conditions to guarantee convergence, if any such guarantee were possible.

One possible course of action after a failed negotiation is to have supporting objects start a new negotiation, but with a new initial proposal representing a different presentation strategy. The system would only resort to the extreme fallback after several such re-negotiations failed. We are investigating this approach, although in practice there has yet been no need to systematically address failed negotiations.

Our negotiation sequence seems to represent the minimum number of rounds for effective exchange of information. The first proposal is made operating on knowledge of only one party's needs (the supporting object). Thus it can only be considered a tentative proposal. The second proposal is made with knowledge of both parties' needs—but from the perspective of only one party (the primary object). Finally, the supporting object has an opportunity to make a choice based on information from both the primary and supporting objects, in the third step. Of course, more rounds of negotiation would inform each participant's decisions even more. However, the three exchanges in the negotiation seem to be effective for the cooperating fluid objects we have designed, and is a simpler design.

SUMMARY

Fluid documents provide lightweight access to supporting information while maintaining the context of the primary information. We have explored a variety of techniques for adjusting the layout and graphical presentation of textual material to accommodate expanding annotations, supported by a negotiation architecture for selecting mutually satisfactory presentations. This system allows different kinds of fluid components to be created, each of which has its own strategies for making space and for presenting itself, and allows these components to co-exist in a fluid document. The negotiation architecture presented in this paper incorporates detailed knowledge about space and salience as well as a negotiation protocol that is guaranteed to terminate. We are employing this architecture in our implementations of non-textual fluid documents that include illustrations, tables, and maps.

Acknowledgments

Ken Fishkin contributed significantly to many early conversations on fluid documents. Many thanks to the UIST reviewers and John Maloney, whose comments on earlier drafts of this paper have improved it considerably. We also thank the other members of the PDR team for their inspiration and enthusiasm.

REFERENCES

1. Lyn Bartram, Albert Ho, John Dill, Frank Henigman. The Continuous Zoom: A constrained fisheye technique for viewing and navigating large information spaces. *UIST'95 Proceedings*, 1995, 207-215.
2. Benjamin B. Bederson, James D. Hollan. Pad++: A zooming graphical interface for exploring alternate interface physics. *UIST'94 Proceedings*, 1994, 17-26.
3. Jacques Bertin. *Semiology of Graphics*. University of Wisconsin Press, 1983.
4. Alan Borning, Kim Marriott, Peter Stuckey, Yi Xiao. Solving linear arithmetic constraints for user interface applications. *UIST'97 Proceedings*, 1997, 87-96.
5. Grace Colby, Laura Scholl. Transparency and blur as selective cues for complex information. *SPIE'91 Proceedings*, 1991, 114-125.
6. George W. Furnas. Generalized fisheye views. *CHI'86 Proceedings*, 1986, 16-23.
7. S. Greenberg, C. Gutwin, A. Cockburn. Awareness through fisheye views in relaxed-WYSIWIS groupware. *Graphics Interface '96 Proceedings*, 1996, 28-38.
8. Beverly L. Harrison, Kim J. Vicente. An experimental evaluation of transparent menu usage. *CHI'96 Proceedings*, 1996, 391-398.
9. Takeo Igarashi, Jock D. Mackinlay, Bay-Wei Chang, Polle T. Zellweger. Fluid visualization of spreadsheet structures. *Visual Languages'98 Proceedings*, 1998.
10. Eser Kandogan, Ben Shneiderman. Elastic Windows: A hierarchical multi-window World-Wide Web browser. *UIST'98 Proceedings*, 1998, 169-177.
11. Y. K. Leung, M. D. Apperley. A review and taxonomy of distortion-oriented presentation techniques. *ACM Transactions on Computer-Human Interaction*, 1 (2), 1994, 126-160.
12. Toshiyuki Masui, Mitsuru Minakuchi, George R. Borden IV, Kouichi Kashiwagi. Multiple-view approach for smooth information retrieval. *UIST'95 Proceedings*, 1995, 199-206.
13. Mark T. Maybury, Wolfgang Wahlster (eds). *Readings in Intelligent User Interfaces*. Morgan Kaufman, 1998.
14. Michael Murtaugh. NiF Elastic Catalog. http://nif.www.media.mit.edu/ecat/.
15. Christine M. Neuwirth, Ravinder Chandhok, David S. Kaufer, Paul Erion, James Morris, Dale Miller. Flexible diff-ing in a collaborative writing system. *CSCW'92 Proceedings*, 1992, 147-154.
16. Ken Perlin, David Fox. Pad - An alternate approach to the computer interface. *SIGGRAPH'93 Proceedings*, 1993, 57-64.
17. George G. Robertson, Stuart K. Card, Jock D. Mackinlay. Information visualization using 3D interactive animation. *Communications of the ACM*, 36 (4), 1993, 57-71.
18. Shumin Zhai, Julie Wright, Ted Selker, Sabra-Anne Kelin. Graphical means of directing users' attention in the visual interface. *INTERACT'97 Proceedings*, 1997, 59-66.
19. Polle T. Zellweger, Bay-Wei Chang, Jock D. Mackinlay. Fluid links for informed and incremental link transitions. *Hypertext'98 Proceedings*, 1998, 50-57.