

Adaptive Layout for Dynamically Aggregated Documents

Evan Schrier¹

Mira Dontcheva¹

Charles Jacobs²

Geraldine Wade³

David Salesin^{1,4}

¹Computer Science & Engineering
University of Washington
Seattle, WA 98105-4615

evans0@msn.com, mirad@cs.washington.edu

²Microsoft Research, ³Microsoft
One Microsoft Way
Redmond, WA 98052-6399

{cjacobs, gwade}@microsoft.com

⁴Adobe Systems
801 N. 34th Street
Seattle, WA 98103
salesin@adobe.com

ABSTRACT

We present a system for designing and displaying grid-based document designs that adapt to many different viewing conditions and content selections. Our system can display traditional, static documents, or it can assemble dynamic documents “on the fly” from many disparate sources via the Internet. Our adaptive layouts for aggregated documents are inspired by traditional newspaper design. Furthermore, our system allows documents to be interactive so that readers can customize documents as they read them. Our system builds on previous work on adaptive documents, using constraint-based templates to specify content-independent page designs. The new templates we describe are much more flexible in their ability to adapt to different types of content and viewing situations. This flexibility comes from allowing the individual components, or “elements,” of the templates to be mixed and matched, according to the content being displayed. We demonstrate our system with two example applications: an interactive news reader for the New York Times, and an Internet news aggregator based on MSN Newsbot.

ACM Classification H5.2 [Information interfaces and presentation]: User Interfaces. - Graphical user interfaces.

General Terms Design, Human Factors

Author Keywords

Adaptive layout, grid-based layout, constraints, PDF, XML, XSL, CSS.

INTRODUCTION

The Internet makes it possible to easily assemble documents using information from many disparate sources. As a result, these documents, which we call *aggregated documents*, are becoming more and more common on the World Wide Web. Today, we see these documents at news aggregators, search engines, and RSS readers. However, aggregated documents currently remain nothing more than a list of hyperlinks, text, and images. We present a system that allows designers to create rich page layouts that re-introduce the graphic designs

Permission to make digital or hard copies of part or all 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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
IUI'08, January 13-16, 2008, Maspalomas, Gran Canaria, Spain.
Copyright 2008 ACM 978-1-59593-987-6/08/0001 \$5.00



Figure 1. Using templates that adapt to both content and display size we present the New York Times RSS feed as a newspaper.

usually found in print media, and apply those designs to a collection of content that is unknown at design time. Our system also facilitates the design’s adaptation to the variety of devices and screen dimensions available today.

For our layouts, we draw inspiration from newspaper design, which provides an existing, well developed methodology for displaying aggregated content, both efficiently and attractively. Arnold [2] describes newspaper design as the combination of typography and layout for quick, accurate transmission of information. As a result, typographic elements such as fonts, spacing, columnation, and juxtaposition are used to create compelling displays of content. In our work we introduce these typographic elements to the design and layout of digital aggregated documents (see Figure 1).

We build on prior work on adaptive grid-based document layout (AGBDL) [15], which allows users to design grid-based document layouts that adapt to different window dimensions. That work showed that high-quality, adaptive document layout is possible through constraint-based layout templates. In this work, we extend those adaptive, grid-based layout ideas to the rendering of dynamically aggregated documents. We also present a novel interaction model for viewing and reading these documents rendered with our system.

The previous system [15] focused on high-quality layout of content with a fixed document structure, such as content in magazines. However, aggregated documents include information from many different sources that do not have a fixed

structure. Because these documents are assembled dynamically, we do not know the content that will be present; therefore, we cannot design specific layout templates ahead of time. Even if we were to try to design templates ahead of time, we would have to enumerate all possible content in all possible combinations, resulting in an explosion of layout templates.

In this paper we present a new template language for building templates that include multiple possibilities for displaying each piece of content, with rules for choosing between them; a designer might have portrait images and landscape images laid out differently, for example. These new templates can also include optional elements, which are displayed only if appropriate content is present. This new template specification vastly increases the applicability of each template, and the range of content that can be displayed with a relatively small collection of templates. Furthermore, the new templates are much easier to write because they are specified with high-level primitives that instruct the system to generate complex constraint systems automatically. The new template language compares to the original as a higher-level programming language compares to assembly language. With our new template specification we are able to display content with the style of the New York Times, as shown in Figure 1, with just five templates. With the previous AGBDL system, this display would require over fifty layout templates to handle all possible content.

As part of our system, we also present a set of interactions for manipulating these documents. Users can now drag and drop stories into different regions of the document. They can page through individual stories without altering others, and they can display hierarchical content such as newspaper sections side by side. In addition to adapting to a variety of display sizes, our documents can adapt by flowing content around overlapping windows of other applications, changing some basic assumptions about the desktop metaphor. We believe our system is the first to include such functionality.

In summary, we present these new ideas:

- placing layout decisions on the individual elements of a page, rather than on the page in its entirety, enabling much richer combinations of elements;
- a new high-level language enabling an easier and more concise template specification by designers who are not programmers;
- “smart” overlapping windows, which reflow content around other applications on the desktop;
- and novel user interactions that allow users to drag and drop hyperlinked content, read content in context, and interactively grow and shrink regions of interest.

We also demonstrate a new application of adaptive layouts for advertising, an innovation that should make this style of reading more commercially viable. Indeed, the research prototype described in this paper was inspiration for the New York Times commercial product, *Times Reader*.

RELATED WORK

There are many services on the Internet that dynamically aggregate Web content. RSS feeds, which allow users to easily stay up to date with favorite news or blog sites, have become so common that RSS readers, which aggregate multiple RSS feeds, are now regularly found in Web portals and mail clients. Search engines also aggregate content when they present users with search results for a query. Finally, news aggregators crawl the Web and compose aggregated documents that display current news around the world. All of these aggregators display their aggregated documents into simple lists using traditional HTML. We chose to implement our own layout renderer because we found HTML too limiting. For example, in HTML, text cannot flow from one element into another, making it difficult to create articles with more than one column. The recent CSS3 specification [22] includes multi-column layout within a single element, but it does not allow for flow between arbitrary elements, flow around image elements, and pagination. Using our adaptive templates, we could provide a richer front-end for RSS clients and news aggregators by simply translating their XML output, as we do in the Newsbot example.

In the HCI community, researchers have explored a variety of displays for aggregated documents. Kandogan and Shneiderman [17] built a specialized browser, Elastic Windows, which allows users to interactively build a custom layout for their Web browser by aggregating visited pages into a hierarchical display. More recently, Chickenfoot [6] and Grease-monkey [7] have allowed developers to write scripts that can automate, aggregate, and customize user interactions with specific websites. Dontcheva *et al.* [10] built a system that aggregates individual pieces of content from multiple visited websites. Their system displays these selected webpage features in template-based layouts, providing interactive summaries of relevant information from the browse session. In subsequent work, Dontcheva *et al.* [11] have extended their approach to create search templates that organize search results using layout templates. To the best of our knowledge, however, our system remains the only one that does adaptive document layout for dynamically aggregated documents. Kamba *et al.* [16] introduced one of the first Web-based adaptive newspapers, The Krakatoa Chronicle, which supports many of the adaptations present in our system, but does not have the ability to render modern grid-based layouts.

One important property of our system is the separation of content and style, which allows designs to accommodate unknown content. This is not a new idea; in fact, several World-Wide-Web Consortium (W3C) standards support this separation, e.g., the Extensible Stylesheet Language (XSL) [1] and Cascading Style Sheets (CSS) [18]. Borning *et al.* [8] were the first to suggest applying constraint systems to the layout of Web documents and applets. In their system, constraints were specified by both the author and the reader, and the final layout was negotiated by priorities on individual constraints. Constraint Cascading Style Sheets [3] built upon this work and then proposed the addition of constraints to the Scalable Vector Graphics language [21], which they called

Constraint Scalable Vector Graphics (CSVG) [4, 19]. CSVG allowed constraints to control the properties of diagram components. They also proposed extensions supporting better interactivity and constraint-based animation. All of these techniques control the layout of a single HTML webpage, while ours can control the layout of content from a collection of sources, integrating them into a single, possibly hierarchical document.

There have been many systems that support the design and construction of constraint-based graphical user interfaces, Garnet [20] and subArctic [14] among others. These toolkits are generally designed for traditional, widget-based interfaces, rather than document-based interfaces such as ours. Gajos and Weld [13] automatically generate adaptive interfaces in the SUPPLE system. Their approach is to find an optimal arrangement of widgets to layout an interface that meets the target device's constraints and is estimated to be easy to use by an objective standard. While document layout is related to widget layout, it addresses a different task and presents a different set of challenges.

BACKGROUND

Since dynamically aggregated documents vary in content structure, we must maintain a clean separation between content and style. Towards this end, we further develop the representation introduced in the AGBDL system. To provide the necessary background for the rest of the paper, we first describe this representation.

Document structure

In our system, a *document* is represented as a set of parallel *streams* of content. For instance, the body text of an article could be one of the streams in the document. Other streams might contain figures, headlines, or sidebars. Within a stream, an individual piece of content could be text, an image, or a hierarchical collection of content. These hierarchical collections, or *subdocuments*, mirror the structure of a document itself. Thus, a figure might be represented as a subdocument containing separate streams for the image, title, caption, and photo credit.

To specify stylistic properties like color, font size, and font face, text can be tagged with a style name. This style name is then looked up in a *stylesheet*, which maps style names to actual style properties. The stylesheets themselves are provided by the layout templates, described below.

Content can be also be annotated with attributes that alter the way it is treated during layout. For example, an image could be marked with an "importance" attribute, and this attribute can be used in computing the desired size of the image in the final layout.

Layout templates

A document's visual style is encoded as a set of *templates*. Each template organizes a collection of content at a range of display sizes. Figure 2 shows a template that displays a headline, an image, and some body text. Each template is composed of *elements* that display content or other visual



Figure 2. A template with its elements outlined in blue. Content can flow from one element to another. We use a layering paradigm to flow the text around the image.

features, *constraints* that define relationships between elements, and *preconditions* that characterize the suitability of the template for the particular content. A template can also reference a stylesheet that is used to map style tags into concrete style properties.

Elements

The basic building blocks of a page layout are *elements*: rectangular regions that are arranged on the page and filled with content. Each element receives content from one of the document streams. More than one element may consume content from the same stream, in which case the content flows from one element into the next. For example, the two-column layout in Figure 2 has two elements, each corresponding to a column. Text flows from the element on the left into the element on the right. To place the image on the page and flow the text around it, we use a layering paradigm. A page can be represented as a $2\frac{1}{2}$ -D scene, where each element also has a *z-order*. Elements that have a higher *z-order* sit atop lower elements, and the area of the higher elements (in this case, the image) is subtracted from the area of the elements underneath (the text).

Since content can be hierarchically organized into subdocuments, templates can also be hierarchical. Elements that are filled by subdocuments can specify a set of templates to format the hierarchical collection of content displayed by that element. When rendering the page, the layout engine chooses the template that best fits the particular subdocument.

In order to allow text styling to vary across the different parts of the page, each element may reference its own stylesheet. Styles present in an element's stylesheet override the styles found in the template's stylesheet.

Constraints

The size and placement of each element in a template is determined by the evaluation of a set of interdependent constraints that form a directed acyclic graph. In Figure 2, the title is constrained to begin at the top of the page and al-

lowed to grow as tall as required in order to fit the title text. The body elements are constrained to begin at the bottom of the title element, and to end at the bottom of the page. The image element is constrained to the middle of the body text. In our implementation, the constraint system is comprised of a pool of constraint variables whose values are computed by a mathematical expression in terms of the other constraint variables. This configuration is known as a one-way constraint system.

A template's constraint variables include both input and output variables. The input variables correspond to the context of the layout template, such as the page dimensions. The output variables correspond to the bounding rectangle for each element and the overall template score, which allows a template to express its fitness in terms of the inputs. If content is tagged with attributes, those values are reflected as additional variables in the constraint system.

Choosing templates: preconditions and scoring

In order to render a particular document, the system must choose an appropriate set of templates. It can make this choice using preconditions or a template's overall score. *Preconditions* define when a template is valid. *Content preconditions* specify content availability requirements, e.g., "there must be at least one item in the figure stream for this template to be considered." *Value preconditions* express a valid range for one of the constraint variables, e.g., "the page width must be between 400 and 600 points." If multiple templates are valid, the layout engine uses the output score to select the best fitting template. Typically a template will produce a low score if the constraints that place the elements are unable to produce a good layout. For example, if a figure is so tall that it doesn't fit on the page, the score might be negative.

The document content and layout templates feed into a document paginator and layout engine to produce a collection of potential page layouts. The paginator can optionally globally optimize the layout by slightly under-filling pages and altering template choices. We use a simpler, "First-Fit" paginator for the examples in this paper. For more details on this process, please see a description of AGBDL [15].

LAYOUT TEMPLATE SPECIFICATION

The principal challenge behind adaptive document layout for dynamically aggregated content lies in the design of layout templates that are flexible to many different assortments of content. Since a designer cannot be expected to design templates for all possible content combinations, template primitives must allow for variable content without explicit enumeration. Designers must also be able to incorporate layout preferences and conditions with minimal effort. Finally, the layout engine must be able to make reasonable choices about content display without intervention. To achieve these goals, we move design decisions to the element level from the overall template level and propose a set of high-level template primitives that are able to position and size available content appropriately. The new primitives we propose here include: *conditional elements*, which automatically choose whether to appear in a given layout based upon the availabil-

ity of content and any other arbitrary constraints; *conditional groups*, which choose a single element from a group of possible elements, also based upon content and constraints; and *adaptive grids*, which automatically choose column count, sizes and margins. The constraint system always sizes and locates any elements that appear on a page.

Conditional elements

As described in the previous section, the layout engine selects layout templates using preconditions that express when a specific template is valid. Preconditions can be defined with respect to content, such as the availability of an image, or with respect to attribute values, such as the possible range of page sizes. In AGBDL, preconditions were used to invalidate entire templates from consideration. Thus, separate templates were necessary for different configurations of content, such as news articles with one, two, or three images. In this work, in order to make the layout templates adapt to the variable structure of the available content, we add preconditions to the individual display elements and call such elements *conditional elements*. Conditional elements can have content and value preconditions. Indeed, placing a content precondition on an element that consumes content, requiring that some content be available, is so natural and useful that we made it implicit. Only one template is now necessary to properly display multiple articles, regardless of the number and orientation of images or the existence of a headline. Figure 3 shows three different articles rendered by the same layout template. Value preconditions can specify the minimum or maximum page size for which an element should be displayed.

In AGBDL, every constraint in the layout template had to be specified by the author. To make positioning of elements easier, we modified the element specification to let a designer specify that an element, such as an image, should always appear fixed at a particular point on the y -axis, or under another element, and always have a specific aspect ratio. The system now generates the necessary constraint system for the designer, which computes the appropriate dimensions and element positions.

Conditional groups

In addition to constraining elements, the designer can also define conditional groups. A *conditional group* defines a set of elements of which only one will be used in a given layout. A *conditional group* is used in the template in Figure 3 to choose from a two-column, one-column, or half-column element depending upon both the aspect ratio of the image and the width of the columns. The designer can specify a conditional group containing a layout element for each orientation along with a policy for selecting the most appropriate element. The policies our system currently includes are best-fit, first-fit, and first-good-fit. A *best-fit* group will select the element that produces the best score. A *first-fit* group will select the first element for which content is available. Finally, a *first-good-fit* group selects the first element it examines with a non-negative score.



Figure 3. We can render articles with a variable number of images and image aspect ratios using only one layout template, by placing the possible image locations in a *conditional group*.

Adaptive grids

Many in the design community [9] believe that long lines of text degrade readability; thus, as the width of an element or page expands, adding more columns improves a document's readability. In AGBDL, dynamic growth of columns was implemented by defining templates for multiple-column layouts and switching to a template with more columns when the page grew beyond a certain width. However, creating templates for different numbers of columns along with templates for different combinations of content leads to a combinatorial explosion in the number of templates necessary to handle different content and viewing situations.

To alleviate this growth in the number of templates, we define an adaptive column grid. The grid divides the page into columns based on parameters provided by the designer. Depending on the page width and the grid specification, the system determines the number of columns to use, and sizes the columns, margins, and gutters whenever the template is rendered.

By default, the system will generate as many columns as possible of at least the minimum width specified in the template. Alternatively, the designer can specify that the system include as few columns as possible without going over the maximum width specified. The individual columns in a grid can also be configured to a fixed width, combining fixed- and variable-width columns of different widths and proportions. When the page is resized, columns will be added in left-to-right order as the page gets wider, and will be dropped from right to left when the width decreases. The designer may, however, change the order in which columns are added and dropped by assigning priority values to each column.

Example

We illustrate the interaction between different primitives in the context of an example template. A complete description of the template language for specifying rich displays for aggregated documents is beyond the scope of this paper. This is a code fragment that shows an adaptive grid and a *best-fit* group.

```
<template id="sample">

  <!-- adaptive grid -->
  <grid type="variable" margins="5">
    <adaptive minColumnWidth="210" maxColumns="4"/>
  </grid>

  <elements>

    <!-- conditional element -->
    <element id="body" column="col1,col2" under="headline">
      <content src="bodyText"/>
    </element>

    <!-- best-fit conditional group -->
    <group id="photoGroup" type="bestfit">
      <element id="photo" template="imageTemplate">
        <content src="image"/>

        <instance column="col1" bottom="page.bottom">
          <pass-constraint var="target.aspect"
            value="portrait"/>
        </instance>

        <instance column="col1,col2" bottom="page.bottom">
          <pass-constraint var="target.aspect"
            value="landscape"/>
        </instance>

        <constraint var="photo.score"
          value="photo.template.score"/>
      </element>
    </group>
  </elements>
</template>
```

This template specifies an adaptive grid with up to four columns of at least 210 points width. It then defines a conditional element and a best-fit conditional group. The group contains one element description for an image, which uses a template, *imageTemplate*, to lay out the content. The layout template *imageTemplate* includes a bitmap and a caption stream. The best-fit group specifies two possibilities, listed as instances of the element. The first instance occupies the first column and the second instance spans the first two columns. We pass each instance a parameter, *target.aspect*, to calculate a layout goodness score. The first instance generates a positive score if it is a portrait image, and the second instance generates a positive score if it is a landscape image. A constraint with the name of the element and a suffix “score”, in this case *photo.score*, evaluates the best-fit. Here it is defined to be the score generated by the template that lays out the element.

HIERARCHICAL DOCUMENTS

When entire articles are encapsulated in an aggregated document, it becomes hierarchical in nature. News aggregators can have deep hierarchies as they include different sections, such as “World”, “National”, “Health”, etc. As we described previously, a document can include other documents, which we call subdocuments, thus creating a hierarchy. Each of these subdocuments will have its own templates, and each one can be recursively paginated and displayed in an element of the parent document’s template. But the layout of each subdocument may also be affected by the overall layout of the parent document. To allow for the interaction between elements at different levels in the hierarchy, we allow the designer to pass input and output variables between a parent and a child template.

Input variables enable the top-level template to exercise control over the appearance of the subdocuments, while still letting them handle most layout functions independently. Returned values allow the parent template to determine if the given subdocument is a good fit in a particular spot. They can further be used to align other display elements with features in the child’s layout. Any constraint in the parent or child template can be passed as a parameter or returned as a result from an adjacent template in the hierarchy. Some of the input variables that we have explored include the number of columns, the amount of padding, and flags for suppressing images or borders. Some of the output variables we have used are layout quality scores and alignment coordinates.

Another way a parent can influence the layout of a child template is by passing it a stylesheet. A designer can override a subdocument’s default stylesheet at the element instance. For example, newspapers will typically use different fonts for nearby headlines so that readers can easily distinguish among them. In our New York Times front page examples (see Figure 8), we avoid displaying abutting headlines in the same style by passing a different stylesheet to adjacent story elements.

For visual consistency, it is often desirable to have a subdocument conform to the grid underlying the parent document. Once a grid is defined for the overall document, each element that displays a subdocument can be passed the number of columns that the element spans in the parent. This allows the subdocuments to use the same number of columns in its own layout, as does the layout in Figure 4. If the parent grid is irregular, a more detailed specification can be provided.

Typically, a subdocument is laid out entirely within a single element on the page. However, if the desired layout does not confine the subdocument to a single rectangular region, the designer may specify that different parts of the subdocument be placed in separate elements on the page. For example, in Figure 4 the large front-page image is in one element, and the story that corresponds to that image is in another element directly below the image.

By default, each element consumes content from a separate subdocument, but the designer can constrain any number of



Figure 4. We use recursive masking to adapt subdocuments to overlap. The right-hand story flows around the bottom story.

elements on the page to consume content from the same subdocument. This can be useful when the desired layout of a subdocument does not fit into a rectangular region on the parent template. For example, in Figure 8 the large front-page image is in one element, and the story that corresponds to that image is in another element directly below the image.

A final consideration for displaying hierarchical documents is handling overlapping elements. As described earlier, we accommodate overlap in a single document by letting the designer assign a *z*-order to the affected elements. The system renders the element with the highest *z*-order first and creates an occlusion mask for the rendered region. Then it renders the remaining elements, adding to the mask as new regions are rendered. We extended this algorithm to hierarchical documents by making the masking recursive and passing the occupied regions of higher-layer elements down the hierarchy. This allows the righthand story in Figure 4 to flow its text around the overlapping story at the bottom of the page. If a subdocument has multiple possible image locations, the template can favor a location that is not occluded.

USER INTERACTION

In addition to adapting to different screen sizes and constantly changing collections of aggregated content, electronic documents should also be responsive to the reading needs of the viewer. For example, the front page of a typical newspaper seldom contains articles in their entirety. The stories generally continue on an interior page, forcing the reader to scan through the paper before finishing them. With our system, users can read articles without scanning through the entire newspaper. They can simply position the cursor over the story of interest and click on the “page down” button. This will display the second page of the story in place of the first. Alternatively, they can click on the headline or the “continue” link, which will take them to a page featuring the entire article, if it fits in the window. If it doesn’t fit, then they will see a paginated view of the article, and they can use the “page down” key again to see the additional pages. A third possibility is to activate an “expand” link that will alter the



Figure 5. The user clicks on the upper left story. The story expands in place and occupies a larger portion of the overall reading window. The other stories are formatted accordingly.

layout on the front page, so that the story of interest has a larger footprint on the page, making it more convenient to read (and page through if necessary) in place (see Figure 5). When expanding a section of a page, the system animates the transition. A smooth, animated transition makes it clear to the user what is being changed and where various features on the page are moving. We implement animation by varying the constraints that size the elements over time. When one story grows the other stories shrink to accommodate it.

While our system creates documents that look like print media, these documents still have digital qualities and include digital elements such as hyperlinks. For example, the table of contents is a list of hyperlinks. To follow hyperlinks, the user can click on them in the conventional fashion. If the link is in a subdocument, such as in an article displayed on the front page, then the subdocument is typically replaced by the linked content.

Alternatively, the document, which in our example is the newspaper's front page, can be replaced by the linked content, which is useful when the linked document is another main page rather than a single article. The hyperlinked content can also replace different subdocuments. For example, a table of contents can be set up to always load the hyperlinked content in a specific location on the front page. The viewer can also interactively specify where the hyperlinked content should be displayed by dragging hyperlinks and dropping them into a display element. The target element can be any element on the page that displays another subdocument. A user can also delete stories from the page that have already been read and are not of interest.

In addition to flowing content recursively around content elements, we have made it possible for our documents to adapt to overlapping windows of other applications (see Figure 6). Text is reflowed in the available regions, and images are relocated if the template allows for alternate locations. We are able to achieve this flexible layout using the masking we described in the previous section, by initializing the top-level mask with the Windows client region of the document viewer. Adapting to overlapping windows is functionality that is well suited to the multitasking environments common on the desktop today.



Figure 6. Our aggregated documents are sensitive to overlap and can flow around any desktop window. The two articles are rendered with the same template, which moves the photo in the left-most article out of the lower-right corner where it would be occluded by the first article and the filesystem folder.

IMPLEMENTATION

We implemented two versions of the system described in this paper, one as a stand-alone client application and the second as a server-side application that delivers HTML webpages. The stand-alone application is far more responsive and provides a better interactive experience, whereas the Web server allows documents to be viewed in any modern Web browser. The main part of our system, the layout engine, is the same for both systems. The only difference is that the server version produces an HTML webpage instead of drawing text and images to the screen. We decided against implementing our work with DHTML because it is difficult to flow text across page elements. Our layout engine implementation is similar to that of AGBDL, described previously in [15]. In this section, we describe the implementation of the new parts of the system and then show results.

We gather the content for our aggregated documents using a variety of techniques. Most content streams provide XML content, which we transform using XSLT (Extensible Stylesheet Language Transformations), a language designed to translate XML content. A manually authored translation script provides rules that govern the transformation. Any arbitrary XML data can be transformed into a document that our system can display, provided a translation file exists. Once a document is translated, it can be displayed as a stand-alone document, or its content can be included in whole or in part in another document.

High-level template primitives

This work adds many new, high-level primitives to the design template language that make it much easier to author templates. In some cases, these new language features ease the production of highly complicated constraint systems, and are similar to macros that get translated into lower-level constraint "code." In other cases, the new language features reflect fundamental new capabilities of the layout system.

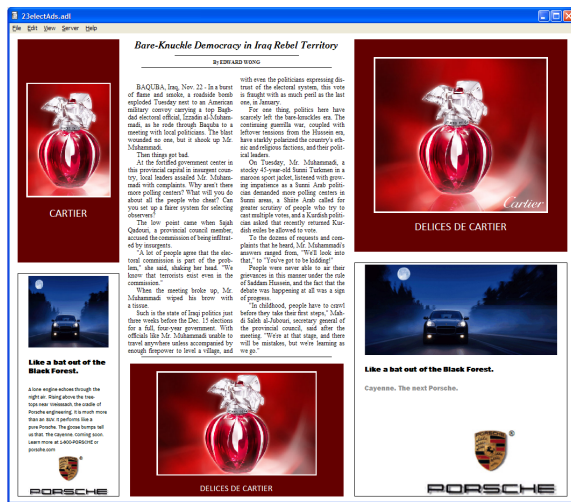


Figure 7. Adaptive advertisements can be designed using our layout templates. Two mock ads are displayed in different size elements.

The most complicated of the new template primitives is the adaptive grid. When the grid is loaded, the system calculates the transition thresholds where columns should be added and removed. Since the designer has a lot of control over the adaptation to various display widths, the system generates a complicated constraint system that activates and sizes the appropriate number of columns, as well as margins and gutters. Columns that are inactive have their width collapsed to zero so that other constraints that depend upon the column are evaluated normally. Sometimes a perfect transition point does not exist. For example, if the allowable column widths are between 200 and 300 points, then there are some pages (e.g., one that is 350 points wide) that are too big for one column and too small for two. The best solution, we believe, is to absorb any extra width into the margins and gutters, which our system will do by default. Alternatively, the layout can be left- or right-justified and all additional space is left in one margin or the other.

While conditional groups could be implemented using constraint systems alone, a more efficient and powerful approach is to place support for them in the layout engine. The system recognizes these structures, iterates through the given layouts, and selects the appropriate element using the policy defined by the designer and the constraint values generated by the layouts. This also allows the element choice to depend upon the results of a given layout, rather than just the initial conditions.

RESULTS

To demonstrate our ideas, we have developed several example applications. The applications include an adaptive version of the online magazine, Slate; an interactive news-reader for the New York Times; and a front-end for the news-aggregator, MSN Newsbot. Figures 8 and 9 show examples of the New York Times and MSN Newsbot applications. Additionally, for the New York Times demonstration we created a set of adaptive mock advertisements shown in Figure 7. These mock ads are just subdocuments that contain the ad

content, along with templates specialized for that content. These ads present a unified visual theme at a wide range of sizes and aspect ratios. We see a lot of potential for this idea, because print media advertising is often sold as a percentage of a page. Our technology allows this established model to be carried into the online world where non-standard page sizes are becoming more and more common.

Our New York Times application uses our new template system and displays the stories of the website's RSS feed as a broadsheet newspaper. To gather the content we combine the RSS feed of the website and a Web scraper that collects the full content of each article from their website. The RSS feed is translated using an XSLT transformation, while the scraped article is converted to our document format using a helper application. The New York Times application displays a set of one- to six- column layouts using only five templates.

We built a Slate magazine application using the website's XML Web feed. We used the earlier AGBDL template system, and it required 74 templates to display four different basic designs; one-, two-, and three-column layouts and a PDA layout. Each one of these templates was four to five times larger than the New York Times templates. It was this initial application that motivated our research in the adaptive layout of dynamically aggregated documents and the development of more powerful templates.

The MSN Newsbot application makes database requests over the Internet to retrieve current breaking news updates from thousands of Web sources and displays them in an attractive adaptive document. This application uses just two top-level templates: one for a small-screen display, and one for all other displays. Each of the stories and each navigation menu of the MSN Newsbot application are subdocuments and use one of a total of four possible subdocument templates.

There are still some major limitations of our system: 1) we currently require designers to author templates in XML code, which will be a barrier for some users; 2) we use a one-way constraint solver for the layout engine (rather than a more general solver such as Cassowary [5] that can handle cycles in the dependency graph), which limits to some extent the page designs that can be described; and 3) the need to write a translation script to parse the raw content from websites into a suitable format for our layout engine, which requires some knowledge of the structure of the targeted website.

The most common application that one-way constraints cannot realize is column-balancing. Multi-column pages that are not entirely full look best when all of the columns are the same length. Since this is common in our domain, we implemented a special primitive that performs iterative balancing on a group of text elements.

FUTURE WORK

There are several ways we would like to expand the system further. We plan to build a graphical template design tool,



Figure 8. We show the New York Times newsreader at two different sizes. The content is automatically fetched live from the RSS feed and displayed using a single template. If the featured story includes a photo in landscape orientation and the browser window is wide enough, that photo is pulled up above the story and placed over two columns. The stories appearing as headlines in the table of contents can be placed directly into any of the story boxes interactively.

ideally one that could work alongside a text-based XML editing tool so that templates could be created and edited in either tool interchangeably.

The template language we have described contains a number of high-level primitives that allow constraint systems to be generated by the system for common situations in adaptive page designs. We plan to identify more of these primitives to make the language more powerful and the descriptions more concise. For instance, we have an automated adaptive column grid, but there may be a similar useful construct in the vertical dimension. We would also like to provide support for letting users define new primitives to generate constraint systems that they use repeatedly in their designs. We intend to get feedback from designers to further refine the definition language.

We hope to evaluate the interaction we propose with a user study. We plan to explore more user interactions, such as allowing a user to modify layouts by directly manipulating boundaries of columns and elements. There is a trade-off, however, between the degree of control a designer maintains over the layouts and the latitude allowed the viewer. Finally, we would like to investigate using a more general constraint solver in the layout engine.

SUMMARY AND CONCLUSION

In this paper, we have presented a number of new ideas. These include: 1) the placement of preconditions and design-evaluation metrics on the individual elements of a page, rather than on just the page in its entirety, thereby enabling much richer combinations of page design elements without elaborating each one as a separate template; 2) a new set of high-level primitives that generate constraint systems automatically given some parameters, thereby enabling a more concise template specification by designers who are not programmers; 3) recursive masking for adaptation to multiple

desktop windows; 4) novel user interactions that allow users to drag and drop hyperlinked content, read content in context, and interactively grow and shrink regions of interest. We also demonstrated a new application of adaptive layouts for advertising (see Figure 7), an innovation that should make this style of reading more commercially viable. However, the paper's biggest contribution perhaps is in the successful and non-obvious assembly of many smaller ideas into a larger whole — the creation of an innovative, working system that had not been possible before.

Virtually all commercial print media today use high-quality grid-based designs, yet current electronic publication tools support these designs only for documents that are static — both in layout and content. Our system addresses this situation by allowing designs that adapt to different viewing sizes, adapt to different dynamically aggregated content, and can be interactively modified by the reader. As on-line reading becomes more and more prevalent, these capabilities will become increasingly important.

The work presented in this paper, we believe, may have broad relevance in the larger field of user interface design. First, on-line reading is, in and of itself, a user-interface issue. Indeed, our adaptive layout templates embody the UI for online reading. They dictate not only how the reading material will appear to the user, but also how the user will navigate through that material or customize the view. In this regard, our templates define adaptive, graphical user interfaces, and the template specification language provides a tool for building them. Moreover, many of the ideas we present may prove relevant beyond just reading: we expect they may be useful for adaptive presentation of other dynamically created or assembled content as well—synthetic UI controls, results of database searches, or numerical data laid out as charts and graphs — to name a few.



Figure 9. We show the MSN Newsbot reader at two different sizes. At the larger size, navigation menus appear on both the left- and right-hand sides; at the PDA size, these appear on the left only. Images are omitted from the navigation pane of the PDA-size layout. Each of the navigation menus is a paginated document, which can be paged through to see more stories.

ACKNOWLEDGEMENTS

We would like to thank Steven Drucker, Patrick Baudisch, Dan Weld, and Wilmot Li for helping us with this paper. Funding and research facilities were provided by Microsoft Research, Adobe Systems, and the University of Washington GRAIL lab.

REFERENCES

- Adler, S., Extensible stylesheet language xsl:Version, 2000.
- Arnold, E., Modern newspaper design, Harper & Row, Publishers, New York, NY, 1969.
- Badros, G. J., Borning, A., Marriott, K., and Stuckey, P. 1999. Constraint cascading style sheets for the Web. In *Proceedings of the 12th Annual ACM Symposium on User interface Software and Technology* (Asheville, North Carolina, United States, November 07 - 10, 1999). UIST '99. ACM Press, New York, NY, 73-82.
- Badros, G. J., Tirtowidjojo, J. J., Marriott, K., Meyer, B., Portnoy, W., and Borning, A. 2001. A constraint extension to scalable vector graphics. In *Proceedings of the 10th international Conference on World Wide Web* (Hong Kong, Hong Kong, May 01 - 05, 2001). WWW '01. ACM Press, New York, NY, 489-498.
- Badros, G. J., Borning, A., and Stuckey, P. J. 2001. The Cassowary linear arithmetic constraint solving algorithm. In *ACM Trans. Comput.-Hum. Interact.* 8, 4 (Dec. 2001), 267-306.
- Bolin, M., Webber, M., Rha, P., Wilson, T., and Miller, R. C. 2005. Automation and customization of rendered web pages. In *Proceedings of the 18th Annual ACM Symposium on User interface Software and*

Technology (Seattle, WA, USA, October 23 - 26, 2005). UIST '05. ACM Press, New York, NY, 163-172.

- Boodman, A. *www.greasespot.net*
- Borning, A., Lin, R., and Marriott, K. 1997. Constraints for the web. In *Proceedings of the Fifth ACM international Conference on Multimedia* (Seattle, Washington, United States, November 09 - 13, 1997), 173-182.
- Brighurst, R., The Elements of Typographic Style. Hartley & Marks, Publishers, Vancouver, BC, Canada, 1996.
- Dontcheva, M., Drucker, S. M., Wade, G., Salesin, D., and Cohen, M. F. 2006. Summarizing personal web browsing sessions. In *Proceedings of the 19th Annual ACM Symposium on User interface Software and Technology* (Montreux, Switzerland, October 15 - 18, 2006). UIST '06. ACM Press, New York, NY, 115-124.
- Dontcheva, M., Drucker, S. M., Salesin, D., and Cohen, M. F. 2007. Relations, Cards, and Search Templates: User-Guided Data Integration and Layout. In *Proceedings of the 20th Annual ACM Symposium on User interface Software and Technology* (Newport, Rhode Island, October 7 - 10, 2007). UIST '07. ACM Press, New York, NY.
- Feiner, S. K. 1988. A grid-based approach to automating display layout. In *Proceedings on Graphics Interface '88* (Edmonton, Alberta, Canada). Canadian Information Processing Society, Toronto, Ont., Canada, 192-197.
- Gajos, K. and Weld, D. S. 2004. SUPPLE: automatically generating user interfaces. In *Proceedings of the 9th international Conference on intelligent User interfaces* (Funchal, Madeira, Portugal, January 13 - 16, 2004). IUI '04. ACM Press, New York, NY, 93-100.
- Henry, T. R., Hudson, S. E., and Newell, G. L. 1990. Integrating gesture and snapping into a user interface toolkit. In *Proceedings of the 3rd Annual ACM SIGGRAPH Symposium on User interface Software and Technology* (Snowbird, Utah, United States, October 03 - 05, 1990).
- Jacobs, C., Li, W., Schrier, E., Barger, D., and Salesin, D. 2003. Adaptive grid-based document layout. In *ACM SIGGRAPH 2003 Papers* (San Diego, California, July 27 - 31, 2003). SIGGRAPH '03. ACM Press, New York, NY, 838-847.
- Kamba, T., Bharat, K., Albers, M. 1995. The Krakatoa Chronicle - An Interactive, Personalized, Newspaper on the Web. In *Proceedings of the Fourth International World Wide Web Conference*, 159 - 170. (Boston, MA, December 11-14, 1995.)
- Kandogan, E. and Shneiderman, B. 1997. Elastic Windows: a hierarchical multi-window World-Wide Web browser. In *Proceedings of the 10th Annual ACM Symposium on User interface Software and Technology* (Banff, Alberta, Canada, October 14 - 17, 1997).
- Lie, H.W., and Box, B. 1996. Cascading style sheets, level 1. W3C recommendation. <http://www.w3.org/style/CSS/>.
- Marriott, K., Meyer, B., and Tardif, L. 2002. Fast and efficient client-side adaptivity for SVG. In *Proceedings of the 11th international Conference on World Wide Web* (Honolulu, Hawaii, USA, May 07 - 11, 2002).
- Myers, B.A., Giuse, D., Dannenberg, R.B., Vander Zanden, B., Kosbie, D., Pervin, E., Mickish, A., and Marchal, P. Garnet: Comprehensive Support for Graphical, Highly-Interactive User Interfaces. *IEEE Computer* 23, 11 (November 1990), 71-85.
- Quint, A. 2003. Scalable vector graphics In *IEEE Multimedia*, Vol.10, Iss.3, July-Sept. 2003 Pages: 99- 102.
- W3C CCS3 Working Draft, 6 June, 2007. Lie, H.K., editor. <http://www.w3.org/TR/css3-multicol>.