



Digestor: device-independent access to the World Wide Web

Timothy W. Bickmore*, Bill N. Schilit¹

FX Palo Alto Laboratory, 3400 Hillview Avenue, Bldg. 4, Palo Alto, CA 94304, USA

Abstract

Digestor is a software system which automatically re-authors arbitrary documents from the world-wide web to display appropriately on small screen devices such as PDAs and cellular phones, providing device-independent access to the web. Digestor is implemented as an HTTP proxy which dynamically re-authors requested web pages using a heuristic planning algorithm and a set of structural page transformations to achieve the best looking document for a given display size. © 1997 Published by Elsevier Science B.V.

Keywords: Device-independent access; PDA; Document re-structuring; WWW

1. Introduction

Access to world-wide web documents from personal electronic devices has been demonstrated in research projects [2,10,17,18], and is now becoming a commercial reality. General Magic's Presto!Links for Sony's MagicLink, AllPen's NetHopper for the Newton and Sharp's MI-10 (shown in Fig. 1), all provide WWW browsers for PDA class devices, while the Nokia 9000 Communicator and Samsung's Duett provide web access capabilities from cellular phones (Fig. 2).

Unfortunately, most pages on the world-wide web are designed for display on desktop computers with color monitors having at least 640x480 resolution, and many pages are designed with even larger monitors in mind. This disconnect can lead to a 4-to-1 (or greater) ratio of designed vs. available screen area,



Fig. 1.

making direct presentation of most WWW pages on these small devices aesthetically unpleasant, un-navigable, and in the worst case, completely illegible. This presents the central problem addressed by this work; how to display arbitrary web pages which have been designed for desktop systems on personal electronic devices which have much more limited I/O capabilities.

* Corresponding author, e-mail: bickmore@pal.xerox.com

¹ E-mail: schilit@pal.xerox.com




Fig. 2.

2. Approaches to device-independent access

There are four general approaches to displaying WWW pages on small screen devices: *device-specific authoring*; *multiple-device authoring*; *client-side navigation*; and *automatic re-authoring*.

Device-specific authoring involves authoring a set of WWW pages for a particular display device, for example a cellular phone outfitted with a display and communications software such as the Nokia 9000. The basic philosophy in this approach is that users of such specialty devices will only have access to a select set of services, and the pages for these services can all be designed up-front for the device's particular display. Information may be provided from the world-wide web at large, but the desired pages must be pre-defined and custom information extraction and page formatting software must be written to deliver the information to the small device. This is the approach taken in Unwired Planet's UP.Link service [16] which uses a proprietary mark-up language (HDML).

In multiple-device authoring, a range of target devices is identified, and mappings from a single source document to a set of rendered documents are defined to cover the devices within the range. One example of this is the StretchText approach [8], in which portions of the document (potentially down to the word level) can be tagged with a "level of abstraction" measure. Upon receiving the document, users can specify the level of abstraction they wish to view and are presented with the corresponding detail or lack thereof. Another example of multiple-device authoring is HTML cascading style sheets (CSS)

[13]. In CSS, a single style sheet defines a set of display attributes for different structural portions of a document (e.g., all top-level section headings are to be displayed in red 18-point Times font). A series of style sheets may be attached to a document, each with a weight describing its desirability to the document's author. The user can also specify a style sheet, as can the WWW browser (the "default" style sheet). Although the author's style sheets normally override the user's, the user can selectively enable or disable the author's, providing the ability to tailor the rendering of the document to their particular display.

In client-side navigation, the user is given the ability to interactively navigate a single web page by altering the portion of it that is displayed at any given time. A very trivial example of this is the use of scroll bars on the document display area. A much more sophisticated approach is that taken in the PAD++ system [3], in which the user is free to zoom and pan the device display over the document with infinite resolution. Active Outlining [11] has also been implemented as a client-side navigation technique, in which the user can dynamically expand and collapse sections of the document under their respective section headings. Other techniques which fall into this category include semi-transparent widgets [12] and the Magic Lens system [4].

Finally, automatic re-authoring involves developing software which can take an arbitrary web document designed for the desktop, along with characteristics of the target display device, and re-author the document through a series of transformations so that it can be appropriately displayed on the device. This process can be performed either on the client, on the server, or on an intermediary HTTP proxy server (as in [7]) which exists solely for the purpose of providing these transformation services. An example of this latter approach is the UC Berkeley Pythia proxy [9], which performs transformations on web page images, although the focus of this work is on minimizing page retrieval time, not on producing the most appropriate page layout for the display device.

There are many possible automatic re-authoring techniques, which can be categorized along two dimensions: syntactic vs. semantic and transformation vs. elision. Syntactic techniques operate on the structure of the page, while semantic techniques rely on some understanding of the content. Elision

techniques basically remove some information, leaving everything else untouched, while transformation techniques involve modifying some aspect of the page's presentation or content. The table below illustrates these dimensions, along with examples of each category.

	Elide	Transform
Syntactic	Section outlining	Image reduction
Semantic	Removing "irrelevant content"	Text summarization

Each of the four approaches — device-specific authoring, multiple-device authoring, client-side navigation, and automatic re-authoring — has benefits and drawbacks. Device-specific authoring will typically yield the best-looking results, but limits the user's access to a small select set of web pages. Multiple-device authoring, while less total effort per document than device-specific authoring, still requires significantly more manual design work than simply authoring for a single desktop platform. Client-side navigation holds promise if a good set of techniques can be developed, but the "peephole" approach taken in PAD++ seems very awkward to use for large documents, and the active outlining technique has limited applicability since most web pages do not use a strict section/sub-section organization. Automatic re-authoring is thus the ideal approach to providing broad access to the web from a wide range of devices, if it can be made to produce legible, navigable and aesthetically pleasing re-authored documents without loss of information.

3. Understanding the re-authoring process

In order to gain an understanding of the process required of an automated re-authoring system, a study was conducted to assess the characteristics of typical web pages, and to identify candidate re-authoring techniques through the process of re-authoring several web pages by hand.

A relatively small sample of "typical" web pages — the Xerox Corporate web site [20] — was initially selected to focus the study. This collection of 3,188 web pages is representative of a state-of-the-art, professionally-designed site. A variety of statistics were collected on these pages using a web crawler,

to help gain an understanding of the structure and content of a typical page. These statistics generally agree with other, larger-scale studies that have been performed across the entire web [6,19].

Next, a subset of the pages in the Xerox web site was selected for manual re-authoring. A set of pages from the Xerox 1995 Annual Report were selected and converted by hand for display on a Sharp Zaurus PDA with a 320 × 240 pixel screen. Detailed notes were kept of the design strategies and techniques used.

Some of the design heuristics learned during this process were:

- Keeping at least some of the original images is important to maintain the look and feel of the original document. Common techniques include keeping only the first, or only the first and last image (*bookend images*) and eliding the rest.
- Section headers (H1–H6 tags) are not often used correctly. They are more frequently used to achieve a particular font size and style (e.g., bold), if they are used at all. Thus, they cannot be relied upon to provide a structural outline for most documents. Instead, documents with many text blocks can be reduced by replacing each text block with the first sentence or phrase of each block (*first sentence elision*).
- An initial rule of thumb for images is to reduce them all in size by a standard percentage, dictated by the ratio of the display area that the document was authored for to the display area of the target device. Images which contain text or numbers can only be reduced by a small amount before their contents become illegible.
- Semantic elision can be performed on sidebars which present information which is tangential to the main concepts presented in a page. Many of the Xerox pages had such sidebars which were simply eliminated in the reduced versions.
- Semantic elision can also be performed on images which do not contribute any information to the page, but serve only to improve its aesthetics.
- Pages can be categorized, and then re-authored based on their category. Two examples of these are *banners* and *link tables*. Banners primarily contain a set of images and a small number of navigation links (often only one) which serve to establish an aesthetic look, but contain little or

no content. When space is at a premium, these can usually be omitted entirely. Link table pages consist primarily of a set of hypertext links to other pages, and very little additional content. These pages can usually be re-formatted into a more compact form which just lists the links in a text block.

- Whitespace, which is taken for granted on a large display, is at a premium on small devices. Several techniques were discovered for reducing the amount of whitespace in a page. Sequences of paragraphs (P tags) or breaks (BR tags) can be collapsed into one. Lists (UL, OL, and DL tags) take up valuable horizontal space with their indenting and bullets, and can be re-formatted into simple text blocks with breaks between successive items (also observed in [8]).

In conclusion, to perform document re-authoring two things are required: a set of re-authoring techniques (a “bag of tricks”), and a strategy for applying them. Of the techniques used in the manual re-authoring study, those most amenable to codification were the syntactic elision techniques (section outlining, first sentence elision, image elision) and the syntactic transformation techniques (image size reduction, font size reduction). The design strategy learned during the study consisted of a ranking of the transformation techniques (i.e., *try this before that*) and a set of conditions under which each transformation or combination of transformations should be applied.

4. Digestor system design

Following the results of the study discussed above, there are two major elements to Digestor’s design: a collection of individual re-authoring techniques which transform documents in various ways; and an automated re-authoring system which implements a design strategy by selecting the best combination of techniques for a given document/display size pair.

4.1. Re-authoring techniques

4.1.1. Outlining

Section header outlining techniques provide a very good method for reducing the required display size for structured documents, such as technical papers and reports. The outlining process is depicted in Fig. 3. The contents of each section is elided from the document and the section header is converted into a hypertext link which, when selected, loads the elided content into the browser. When confronted with multiple section levels (sections, sub-sections, sub-sub-sections, etc.), there are two approaches to performing the elision. The first — *full outlining* — works by keeping only the section headers and eliding all content, with the results looking like a table of contents for a book. In the second approach — *to-level outlining* — a cutoff level in the section hierarchy is determined and all content below that level (including lower-level section headers) is elided, but all content above that level is kept.

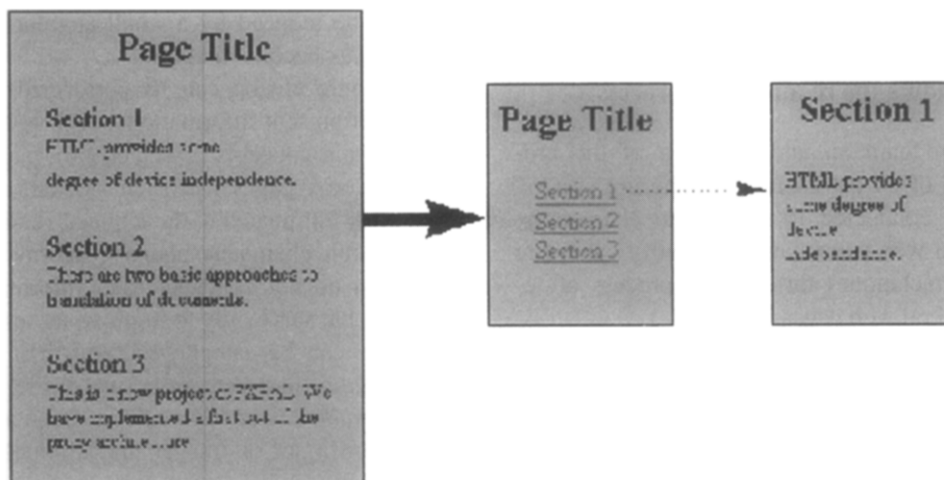


Fig. 3.

4.1.2. First sentence elision

Since most pages have text blocks, even when no section headers are present, first sentence elision can be a good way of reducing required screen area. In this technique, each text block is replaced with its first sentence (or phrase up to some natural break point), and this sentence is also made into a hypertext link to the original text block.

4.1.3. Image reduction and elision

Images present one of the most difficult problems for automatic re-authoring, because the decision of whether to keep, reduce, or eliminate a given image should be based on an understanding of the content and role of the image on the page. However, image reduction and elision can be applied without content understanding, as long as users are provided a mechanism by which they can retrieve the original image. The approach taken in *Digestor* is to provide a set of techniques which transform all images in a page by pre-defined scaling factors (25%, 50%, and 75%), and making the reduced images hypertext links back to the originals. In addition to image reduction, three syntactic elision techniques were also developed for images — elide all, first image only, and bookends — in which all images, all but the first, and all but the first and last are elided from the document, respectively. Elided images are replaced with their ALT text when available, or with a standard icon when no ALT text is available, which is also made into a hypertext link to the original image.

4.2. Automated re-authoring system

The overall process of deciding which combination of techniques to apply to a given page for a given client display seems at first to require some form of human artistic ability. However, an automatic re-authoring algorithm has been developed which captures many of the heuristics used in the manual re-authoring exercise, and which does a fairly good job of producing good looking pages for a given display.

The basic approach is that of a heuristic planner [1] which searches a document transformation space in a best-first manner. Each state in this space represents a version of the document, with the initial state representing the original “as-authored” docu-




Fig. 4.

ment. The measure of merit (evaluation function) for each state is a rough estimate of the screen area required for the document as it exists in that state (Fig. 4). A state can be expanded into a successor state through the application of single transformation technique. At every step in the search process the most promising state (that with the smallest display area requirements) is selected and a transformation applied to it, if possible. As soon as a state is created containing a document version that is “good enough”, the search is halted and that document is returned to the client for rendering. If the search is exhausted and no document version can be found that is good enough, then the best document found during the search is returned.

Heuristic information is used in several places in the planner, including: the order in which transformation techniques are applied to a given state; the pre-conditions for each transformation technique; and the determination of when a document version is “good enough” to halt the search. In general, transformations which make minor changes to the document are preferred over those which make more extensive changes (e.g., reducing images by 25% is preferable to reducing them by 75%). The pre-conditions for each transformation technique specify

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.