WebQuilt: A Framework for Capturing and Visualizing the Web Experience

Jason I. Hong, Jeffrey Heer, Sarah Waterson, and James A. Landay
Group for User Interface Research, Computer Science Division
University of California at Berkeley
Berkeley, CA 94720-1776 USA
+1 510 643 7354

{jasonh, waterson, landay}@cs.berkeley.edu,   jheer@hkn.eecs.berkeley.edu

ABSTRACT

WebQuilt is a web logging and visualization system that helps web design teams run usability tests (both local and remote) and analyze the collected data. Logging is done through a proxy, overcoming many of the problems with server-side and client-side logging. Captured usage traces can be aggregated and visualized in a zooming interface that shows the web pages people viewed. The visualization also shows the most common paths taken through the website for a given task, as well as the optimal path for that task as designated by the designer. This paper discusses the architecture of WebQuilt and also describes how it can be extended for new kinds of analyses and visualizations.

Categories and Subject Descriptors

H.1.2 [**Models and Principles**]: User/Machine Systems – *Human factors;* H.3.5 [**Information Storage and Retrieval**] Online Information Services – *Web-based services;* H.5.2 [**Information Interfaces and Presentation**] User Interfaces – *Evaluation / methodology;* H.5.4 [**Information Interfaces and Presentation**] Hypertext/Hypermedia – *User issues*

General Terms

Measurement, Design, Experimentation, Human Factors

Keywords

usability evaluation, log file analysis, web visualization, web proxy, WebQuilt

1. INTRODUCTION

1.1 Background

There are two usability problems all web designers face: understanding what tasks people are trying to accomplish on a website and figuring out what difficulties people encounter in completing these tasks. Just knowing one or the other is insufficient. For example, a web designer could know that someone wants to find and purchase gifts, but this isn't useful unless the web designer also knows what problems are preventing the individual from completing the task. Likewise, the web designer could know that this person left the site at the checkout process, but this isn't meaningful unless the designer also knows

that he truly intended to buy something and is not simply browsing.

There are a variety of methods for discovering what people want to do on a website, such as structured interviews, ethnographic observations, and questionnaires (for example, see [1]). Instead, we focus here on techniques designers can use for tackling the other problem, that is, understanding what obstacles people are facing on a website in the context of a specific task.

Through interviews with a number of web designers, we identified a few important indicators to look for when analyzing the results of a task-based usability test. These indicators include identifying the various paths users take, recognizing and classifying the differences in browsing behavior, knowing key entry and exit pages, and understand various time-based metrics (e.g. average time spent on a page, time to download, etc.). All of this data, when given the framework of a task and the means to analyze it, would be useful for designers.

Traditionally, this kind of information is gathered by running usability tests on a website. A usability specialist brings in several participants to a usability lab and asks them to complete a few predefined tasks. The usability engineer observes what stumbling blocks people come across and follows up with a survey and an interview to gain more insights into the issues.

The drawback to this traditional approach is that it is very time consuming to run usability tests with large numbers of people: it takes a considerable amount of work to schedule participants, observe them, and analyze the results. Consequently, the data tends to reflect only a few people and is mostly qualitative. These small numbers also make it hard to cover all of the possible tasks on a site. Furthermore, small samples are less convincing when asking management to make potentially expensive changes to a site. Lastly, a small set of participants may not find the majority of usability problems. Despite previous claims that around five participants are enough to find the majority of usability problems [2, 3], a recent study by Spool and Schroeder suggests that this number may be nowhere near enough [4]. Better tools and techniques are needed to increase the number of participants and


Figure 1. Server-side logging is done on the web server, but the data is available only to the owners of the server.

tasks that can be managed for a usability test.

In contrast to traditional usability testing, server log analysis (See Figure 1) is one way of quantitatively understanding what large numbers of people are doing on a website. Nearly every web server logs page requests, making server log analysis quite popular. In fact, there are over 90 research, commercial, and freeware tools currently available [5]. Server logging also has the advantage of letting test participants work remotely in their own environments: instead of coming to a single place, usability test participants can evaluate a website from any location on their own time, using their own equipment and network connection.

However, from the perspective of the web design team, there are some problems with server logs. Access to server logs are often restricted to just the owners of the web server, making it difficult to analyze subsites that exist on a server. For example, a company may own a single web server with different subsites owned by separate divisions. Similarly, it is also impractical to do a log file analysis of a competitor's website. A competitive analysis is important in understanding what features people consider important, as well as learning what parts of your site are easy-to-use and which are not.

Figure 2. Client-side logging is done on the client computer, but requires special software running in the background or having a special web browser.

Client-side logging has been developed to overcome these deployment problems. In this approach, participants remotely test a website by downloading special software that records web usage (See Figure 2). However, client-side logging has two weaknesses. First, the design team must deploy the special software and have end-users install it. Second, this technique makes it hard to achieve compatibility with a range of operating systems and web browsers. What is needed is a logging technique that is easy to deploy for any website and is compatible with a number of operating systems and browsers.

Another problem with using either server- or client-side web logs to inform web design is that existing server log analysis tools do not help web designers understand *what* visitors are trying to do on a website. Most of these tools produce aggregate reports, such as "number of transfers by date" and "most popular pages." This kind of information resembles footsteps in the forest: you know someone has been there and where they went, but you have no idea what they were trying to do and whether they were successful. To better understand usability problems, designers need logging tools that can be used in conjunction with known tasks, as well as sophisticated methods for analyzing the logged data.

1.2 WebQuilt

As pointed out, gathering web usability information is not a simple task with current tools. Furthermore, a best practice

industry has learned is that the earlier usability feedback can be incorporated into the design, the easier and less costly it is to fix the problems. To recap, there are four things that could greatly streamline current practices in web usability evaluations:

1. A way of logging web usage that is fast and easy to deploy on *any* website
2. A way of logging that is *compatible* with a range of operating systems and web browsers
3. A way of logging where the *task* is already known
4. Tools for *analyzing* and *visualizing* the captured data

To address these needs, we developed WebQuilt, a tool for capturing, analyzing, and visualizing web usage. To address the first and second needs, we developed a proxy-based approach to logging that is faster and easier to deploy than traditional log analysis techniques (See Figure 3). This proxy has better compatibility with existing operating systems and browsers and requires no downloads on the part of end-users. It will also be easier to make compatible with future operating systems and browsers, such as those found on handheld devices and cellular phones.
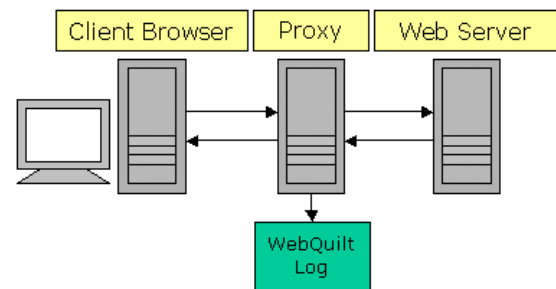
Figure 3. Proxy-based logging is done on an intermediate computer, and avoids many of the deployment problems faced by client-side and server-side logging.

To address the third need, we designed the proxy to be flexible enough that it can be used in conjunction with existing tools, such as those offering participant recruitment and online surveys. With these existing tools, we can know who the users are, what tasks they are trying to accomplish, and whether they were satisfied with how the site supported these tasks (for example, tools like these are provided by NetRaker [6] and Vividence [7]).

To address the fourth need, we designed a visualization that takes the aggregated data from several test sessions and displays the web pages people viewed, as well as the paths they took. However, knowing that we would not immediately have all of the solutions for analyzing the resulting data, WebQuilt was designed to be extensible enough so that new tools and visualizations could be implemented to help web designers understand the captured data.

WebQuilt is intended for task-based usability tests. Test participants are given specific tasks to perform, such as browsing for a specific piece of information or finding and purchasing an item. The WebQuilt proxy can track the participants' actions, whether they are local or remote. After a number of web usage traces have been captured, tools developed with the WebQuilt framework can be used to analyze and visualize the results, pointing to both problem areas and successful parts of the site. It is important that a task be attached to the test participants'

Figure 4. WebQuilt dataflow overview. The proxy logger captures web sessions, generating one log file per session. Each log file is processed by the Action Inferencer, which converts the log of page transactions into a log of actions. The results are combined by the Graph Merger, laid out by the Graph Layout, and visualized by the Visualization component.

interactions, because otherwise one must interpret the intent of visitors, something that is difficult to do based on web usage traces alone. Though WebQuilt can certainly be used to capture any general browsing behavior, the visual analysis provided in this paper is structured to support a task-basked framework.

In the rest of this paper, we describe the architecture of WebQuilt and give a description of our current visualization tool. We then close with a discussion of related work and directions we plan to take in the future.

2. WEBQUILT ARCHITECTURE

WebQuilt is separated into five independent components: the Proxy Logger, the Action Inferencer, the Graph Merger, the Graph Layout, and the Visualization (See Figure 4). The *Proxy Logger* mediates between the client browser and the web server and logs all communication between the two. The *Action Inferencer* takes a log file for a single session and converts it into a list of actions, such as "clicked on a link" or "hit the back button." The *Graph Merger* combines multiple lists of actions, aggregating what multiple people did on a website into a directed graph where the nodes represent web pages and the edges represent page requests. The *Graph Layout* component takes the combined graph of actions and assigns a location to each node. The *Visualization* component takes the results from the Graph Layout component and provides an interactive display.

Each of these components was designed to be as independent of each other as possible. There is a minimal amount of communication between each component, to make it as easy as possible to replace components as better algorithms and techniques are developed. In the rest of this section, we describe each of these components in detail.

2.1 Proxy Logger

The goal of the proxy logger is to capture user actions on the web. As a proxy, it lies between clients and servers, with the assumption that clients will make all requests through the proxy. Proxies have been used in a number of applications, from mining user "trails," caching, and web cataloguing [8, 9]. WebQuilt uses a proxy to log user sessions. In this section we first discuss problems with current logging techniques, describe how WebQuilt's proxy approach addresses these problems, and then continue with a description of the proxy's architecture.

2.1.1 Problems with Existing Logging Techniques

Currently, there are two common ways of capturing and generating web usage logs: server-side and client-side logging. Server-side logs have the advantage of being easy to capture and generate, since all transactions go through the server. However, there are several downsides to server-side logging, as pointed out by Etgen and Cantor [10] and by Davison [11]. One problem is that web caches, both client browser caches and Intranet or ISP caches, can intercept requests for web pages. If the requested page is in the cache then the request will never reach the server and is thus not logged. Another problem is that multiple people can also share the same IP address, making it difficult to distinguish who is requesting what pages (for example, America Online, the United States' largest ISP, does this). A third problem with server-side logging is with dynamically assigned IP addresses, where a computer's IP address changes every time it connects to the Internet. This can make it quite difficult to determine what an individual user is doing since IP addresses are often used as identifiers. While researchers have found novel ways of extracting useful user path data from server logs on a statistical level [12], the exact paths of individual users still remain elusive. Furthermore, with standard server logs users' tasks and goals (or lack thereof) are highly ambiguous.

One alternative to gathering data on the server is to collect it on the client. Clients are instrumented with special software so that all usage transactions will be captured. Clients can be modified either by running software that transparently records user actions whenever the web browser is being used (as in [13]), by modifying an existing web browser (as in [14] and [15]), or by creating a custom web browser specifically for capturing usage information (as with [7]).

The advantage to client-side logging is that literally everything can be recorded, from low-level events such as keystrokes and mouse clicks to higher-level events such as page requests. All of this is valuable usability information. However, there are several drawbacks to client-side logging. First, special software must be installed on the client, which end-users may be unwilling or unable to do. This can severely limit the usability test participants to experienced users, which may not be representative of the target audience. Second, there needs to be some mechanism for sending the logged data back to the team that wants to collect the logs. Third, the software is platform dependent, meaning that the software only works for a specific operating system or specific browser.

WebQuilt's logging software differs from the server-side and client-side approaches by using a proxy for logging instead. The proxy approach has three key advantages over the server-side approach. First, the proxy represents a separation of concerns. Any special modifications needed for tracking purposes can be done on the proxy, leaving the server to deal with just serving content, making it easier to deploy, as the server and its content do not have to be modified in any way.

Second, the proxy allows anyone to run usability tests on any website, even if they do not own that website. One can simply set up a proxy and ask testers to go through the proxy first. The proxy simply modifies the URL of the targeted site to instead go through the proxy. End users do not have to change any settings to get started. Again, this makes it easy to run and log usability tests on a competitor's site.

Finally, having testers go through a proxy allows web designers to "tag" and uniquely identify each test participant. This way designers can know *who* the tester was, *what* they were trying to do, and afterwards can ask them *how well* they thought the site supported them in accomplishing their task.

A proxy logger also has advantages over client-side logging. It does not require any special software on the client beyond a web browser, making it faster and much simpler to deploy. The proxy also makes it easier to test a site with a wide variety of test participants, including novice users who may be unable or afraid to download special software. It is also more compatible with a wider range of operating systems and web browsers than a client-side logger would be, as it works by modifying the HTML in a platform-independent way. Again, this permits testing with a more realistic sample of participants, devices, and browsers.

It is important to note that this approach is slightly different from traditional HTTP proxies. Traditional proxies (e.g. a corporate firewall) serve as a relay point for all of a user's web traffic, and the user's browser must be configured to send all requests through the proxy. The WebQuilt proxy differs in that it is URL based – it redirects all links so that the URLs themselves point to the proxy, and the intended destination is encoded within the URL's query string. This avoids the need for users to manually configure their browsers to route requests through the WebQuilt proxy, and so allows for the easy deployment of remote usability tests by simply providing the proper link.

2.1.2 WebQuilt Proxy Logger Implementation
The current WebQuilt proxy logger implementation uses Java Servlet technology. The heart of this component, though, is the log file format, as it is the log files that are processed by the Action Inferencer in the next step. To use the WebQuilt analysis tools, it actually does not matter what technologies are used for logging or whether the logger lies on the server, on a proxy, or on the client, as long as the log format is followed. Presently, the WebQuilt Proxy Logger creates one log file per test participant session.

WebQuilt Log File Format
Table 1 shows a sample log. The *Time* field is the time in milliseconds the page is first returned to a client, where 0 is the start time of the session. The *From TID* and the *To TID* fields are transaction identifiers. In WebQuilt, a transaction ID represents the Nth page that a person has requested. The From TID field represents the page that a person came from, and the To TID field represents the current page the person is at. The transaction ID numbers are used by the Action Inferencer for inferring when a person used the browser back button and where they went back. The *Parent ID* field specifies the frame parent of the current page. This number is the TID of the frameset to which the current page belongs, or –1 if the current page is not a frame.

The *HTTP Response* field is just the response from the server, such as "200 ok" and "404 not found." The *Link ID* field specifies which link was clicked on according to the Document Object Model (DOM). In this representation, the first link in the HTML has link ID of 0, the second has link ID of 1, and so on. Both <A> and <AREA> tags are considered links. This data is useful for understanding which links people are following on a given page.

The *Frame ID* field indicates which frame in an enclosing frameset the current page is in. These are numbered similarly to link IDs – a frame ID of 0 indicates the first frame in the frameset, and so on. If the current page is not a frame, a value of –1 is used. The *HTTP Method* field specifies which HTTP method was used to request the current page. Currently the proxy supports the GET and POST methods. The last fields are the *URL + Query* fields, which represent the current page the person is at and any query data (e.g. CGI parameters) that was sent along with the request.

The WebQuilt log format supports the same features that other log formats do. For example, the first row shows a start time of 6062 msec and the second row 11191 msec. This means that the person spent about 5 seconds on the page http://www.google.com. However, it has two additional features other logging tools and formats do not. The first is the Link ID. Without this information, it can be difficult to tell which link a person clicked on if there are redundant links to the same page, which is a common practice in web design. This can be important in understanding which links users are following and which are being ignored. The second is finding where a person used the back button. The highlighted cells in Table 1 show an example of where the person used the back button to go from transaction ID of 4 back to transaction ID of 2, and then forward again, this time to a different destination.

Time	From TID	To TID	Parent ID	HTTP Response	Frame ID	Link ID	HTTP Method	URL + Query
6062	0	1	-1	200	-1	-1	GET	http://www.google.com
11191	1	2	-1	200	-1	-1	GET	http://www.phish.com/index.html q=Phish&btnI=I%27m+Feeling+Lucky
167525	2	3	-1	200	-1	1	GET	http://www.phish.com/bios.html
31043	3	4	-1	200	-1	2	GET	https://www.phish.com/bin/catalog.cgi
68772	2	5	-1	200	-1	15	GET	http://www.emusic.com/features/phish

Table 1. Sample WebQuilt log file in tabular format. The highlighted cells show where a person went back from the fourth requested page to the second, and then forward again.
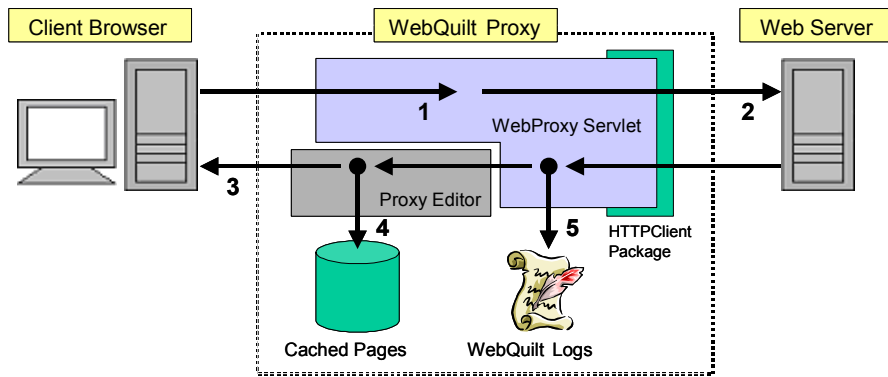
Figure 5. Proxy architecture overview.

WebQuilt Proxy Logger Architecture

Figure 5 illustrates the Proxy Logger's architecture. As mentioned before, the WebQuilt Proxy is built using Java Servlet technology. The central component of the system is the WebProxy servlet, which must be run within a Servlet and JSP engine (e.g. Jakarta Tomcat or IBM WebSphere). The Servlet engine provides most of the facilities for communicating with the client – it intercepts HTTP requests and hands them off to the WebProxy servlet, handles session management, and provides output streams for sending data back to the client. The WebProxy component processes client's requests and performs caching and logging of page transactions. It is aided by the ProxyEditor module, which updates all the links in a document to point back to the proxy. Underlying the WebProxy component is the HTTPClient library [16], an extended Java networking library providing full support for HTTP connections, including cookie handling.

Phase (1) - Processing Client Requests

In the first phase, an HTTP request is received from the client by the proxy. All WebQuilt specific parameters (including the destination URL) are extracted and saved. At this time the proxy collects most of the data that is saved in the log file. For example, the time elapsed since the beginning of the session is calculated, and other various parameters such as transaction IDs, parent ID, and link ID are stored.

There are two ways a person can start using the proxy. The first is by requesting the proxy's default web page and submitting a URL to the proxy (See Figure 6). The other way a person could start using the proxy is by using a link to a proxied page. For example, suppose you wanted to run a usability study on Yahoo's website. If the proxy's URL was:
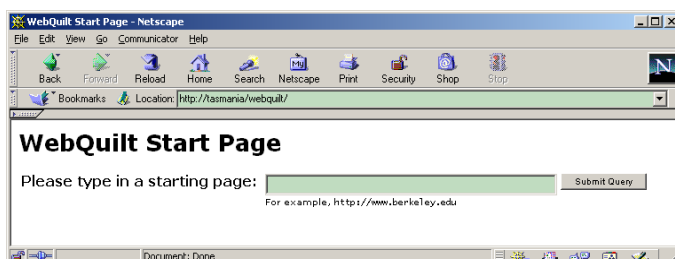
http://tasmania.cs.berkeley.edu/webquilt,

then participants could just use the following link:

http://tasmania.cs.berkeley.edu/webquilt/webproxy?replace=http://www.yahoo.com.

This method makes it easy to deploy the proxy, as the link can just be sent via email to users. Again, we expect other tools to be used for recruiting participants and specifying tasks for them to do. The proxy is flexible enough that it can easily be integrated with such other usability evaluation tools.

Phase (2) - Retrieving the Requested Document

After the client's request has been received and analyzed, the proxy attempts to retrieve the document specified by the request. If no document has been requested (i.e. the *replace* parameter was absent) the proxy returns the default start page. Otherwise, the proxy opens an HTTP connection to the specified site and requests the document using either the GET or POST method, depending on which method the client used to request the page from the proxy. The proxy then downloads the document from the server. At this time, the rest of the data needed for the log entry is also stored, including the HTTP response code and the final destination URL (in case we were redirected by the server).

Phase (3) – Redirecting Links to the Proxy

Before the downloaded page is sent back to the client, it must be edited so that all the links on the page are redirected through the proxy. This work is done by the Proxy Editor module. Initially, the proxy checks the content type of the page. If the content type is provided by the server and is not of the form text/html, the proxy assumes that the page is not an HTML document and returns it to the client without editing. Otherwise, the proxy runs the page through the proxy editor.

The editor works by dynamically modifying all requested pages, so that future requests and actions will be made through the proxy. The document's base HREF is updated and all links, including page hyperlinks, frames, and form actions, are redirected through the proxy.

First, the <BASE> tag is updated or added to the page within the page's enclosing <HEAD> tags. This tag's HREF field points to the document base – a location against which to resolve all relative links. This allows the client browser to then request stylesheets, images, and other embedded page items from the correct web location rather than through the proxy.

Figure 6. Default page for the WebQuilt proxy. The proxy will retrieve and dynamically modify the URL that is entered.

The proxy editor also modifies all link URLs in the page to use the proxy again on the next page request. Thus, once a person has started to use the proxy, all of the links thereafter will automatically be rewritten to continue using the proxy. The editor also adds Transaction IDs to each link. Again, Transaction IDs represent the Nth page that a person has requested. Embedding the transaction ID into a link's URL lets the proxy identify exactly what page a person came from. Link IDs are also added to each link URL. This allows the proxy to identify exactly which link in the page a person clicked on. If the current page is a frame, the proper Parent ID and Frame ID parameters are also included.

The link URL and other parameters are included as variables within a query string for the link. The link URL itself is first rewritten as an absolute link and then added, along with the other parameters, to a query that will be read by the proxy. For example, if you are viewing the page www.yahoo.com through the proxy, the link

```
<A HREF="computers.html">
```

could be rewritten as

```
<A HREF="http://tasmania.cs/webquilt/webproxy?
replace=http://www.yahoo.com/computers.html&ti
d=1&linkid=12">
```

HTML <FRAME> tags are dealt with similarly – the URLs for the target frames are rewritten to pass through the proxy and extra information such as the frame parent's TID and the frame ID are included. <FORM> tags are dealt with a little differently. The <FORM> tag's ACTION field is set to point back to the proxy as usual, but the actual target URL, the current TID, and, if necessary, the Parent ID and Frame ID, are encapsulated in <INPUT> tags with input type "hidden". These tags are inserted directly after the enclosing <FORM> tag. Since they are of type "hidden", they do not appear to the user while browsing, but are included in the resulting query string upon a FORM submit. The proxy editor also handles tags of the form <META HTTP-EQUIV="refresh" …>. These tags cause the browser to load a new URL after a specified time duration. The editor updates these tags to make sure the new URL is requested through the proxy.

The Proxy Editor implementation uses a simple lexical analysis approach to edit the page. The editor linearly scans through the HTML; comments and plain text are passed along to the client unchanged. When a tag is encountered, the type of the tag (e.g. 'A' or 'TABLE') is compared against a set of tags that require editing. If the tag is not in this set (i.e. not a link) it is simply passed along to the client, otherwise it is handed off to the proper TagEditor module, which updates the tag contents as described above, before being sent along.

Phases (4) and (5) – Page Caching and Logging

After performing any necessary editing and sending the requested document to the client, the proxy then saves a cached copy of the HTML page. Before writing out to disk or to a database, the original document is run through the proxy editor again, but this time only the <BASE> tag is updated. This allows for the page to be opened locally and yet still appear as it would on the web, as long as none of the non-cached items, such as images, have not changed on the server. Finally, the log entry for the current transaction is written to the appropriate log file.

2.1.3 Additional Proxy Functionality

The base case of handling standard HTTP and HTML is straightforward. However, there are also some special cases that must be dealt with. For example, cookies are typically sent from web servers to client browsers. These cookies are sent back to the web server whenever a client browser makes a page request. The problem is that, for security and privacy reasons, web browsers only send cookies to certain web servers (ones in the same domain as the web server that created the cookie in the first place). To address this, the proxy logger manages all cookies for a user throughout a session. It keeps a table of cookies, mapping from users to domains. When a page request is made through the proxy, it simply looks up the user, sees if there are any cookies associated with the requested web server or page, and forwards these cookies along in its request to the web server. This is currently handled within the proxy by the HTTPClient library, with modifications made to ensure separate cookie tables are used for each active user session.

Another special case that must be dealt with is the HTTPS protocol for secure communication. HTTPS uses SSL (Secure Socket Layer) to encrypt page requests and page data. The proxy logger handles HTTPS connections by using two separate secure connections. When a client connects to the proxy over a secure connection, the proxy in turn creates a new HTTPS connection to the destination server, ensuring that all network communication remains secure. Our implementation uses Sun's freely available JSSE (Java Secure Socket Extension) [17] in the underlying network layer to enable encrypted communication both to and from the proxy.

2.1.4 Proxy Logger Limitations

Trapping every possible user action on the web is a daunting task, and there are still limitations on what the WebQuilt proxy logger can capture. The most pressing of these cases is links or redirects created dynamically by JavaScript and other browser scripting languages. As a consequence, the JavaScript generated pop-up windows and DHTML menus popular on many websites are not captured by the proxy. Other elusive cases include server-side image maps and embedded page components such as Java applets and Flash animations.

One obvious way to overcome these limitations is to use a traditional proxy approach, where all requests are transparently routed through a proxy. While this would certainly allow one to capture all user interactions, it introduces some serious deployment issues. Most significantly, the traditional proxy approach would require users to configure their browsers to use the proxy and then undo this setting after performing usability tests. This would seriously hamper the ease with which remote usability tests could be performed. Furthermore, any users who currently sit behind a firewall would be unable to participate, as changes to their proxy settings could render them unable to connect to the internet.

2.2 Action Inferencer

Action Inferencers transform a log of page requests into a log of inferred actions, where an action is currently defined as either requesting a page, going back by hitting the back button, or going forward by hitting the forward button. The reason the actions must be inferred is that the log generated by the proxy only captures

page requests. The proxy cannot capture where a person uses the back or forward buttons of the browser to do navigation, since pages are loaded from the local browser cache.

WebQuilt comes with a default Action Inferencer, but the architecture is designed such that developers can create and plug in new ones. It should be noted that given our logging approach, the inferencer can be certain of when pages were requested and can be certain of when the back button was used, but cannot be certain of back and forward combinations. Additionally, the current implementation does not specifically identify when a user clicks on the browser's refresh button.

As an example, figure 7 shows a graph of a sample log file. Figure 8 shows how the default Action Inferencer interprets the actions in the log file. We know that this person had to have gone back to Transaction ID 1, but we don't know exactly how many times they hit the back and forward buttons. Figure 8 shows what happens if we assume that the person went directly back from TID 3 to TID 1, before going on to TID 4.

Figure 9 shows another valid way of inferring what happened with the same log file. The person could have gone back and forth between TID 2 and 3 a few times before returning to TID 1.

2.3 Graph Merger
The Graph Merger takes all of the actions inferred by the Action Inferencer and merges them together. In other words, it merges multiple log files together, aggregating all of the actions that test participants did. A graph of web pages (nodes) and actions (edges) for the task is available once this step is completed.

2.4 Graph Layout
Once the log files have been aggregated, they are passed to the Graph Layout component, which prepares the data for visualization. The goal of this step is to give an (x,y) location to all of the web pages. Since there are a variety of graph layout algorithms available, we have simply defined a way for developers to plug-in new algorithms. Currently, WebQuilt uses an edge-weighted depth-first traversal of the graph, displaying the most trafficked path along the top, and incrementally placing the less and less followed paths below. This algorithm also uses a grid positioning to help organize and align the distances between the nodes. Another possible algorithm that has been attempted is a simple force-directed layout of the graph. This algorithm tries to place connected pages a fixed distance apart, and tries to spread out unconnected web pages at a reasonable distance.

2.5 Visualization
The final part of the WebQuilt framework is the visualization component. There are many ways of visualizing the information. We have built one visualization that shows the web pages traversed and paths taken (See Figures 10 and 11).

Web pages are represented by screenshots of that page as rendered in a web browser. Arrows are used to indicate traversed links and where people hit the back button. Thicker arrows indicate more heavily traversed paths. Color is used to indicate the average amount of time spent before traversing a link, with colors closer to white meaning short amounts of time and colors closer to red meaning longer amounts of time. Zooming is used to see the URL

Figure 7. A graphical version of the log file in Table 1. The letters 'A', 'B', 'C', and 'D' are for this graph only and are not part of the log file.

Figure 8. One possible way of interpreting the log file in Table 1. This one assumes that a person repeatedly hit the back button before clicking on a new link.

Figure 9. Another way of interpreting the log file in Table 1. This one assumes that a person uses the back and forward buttons a few times before clicking on a new link.

for a web page and to see a detailed image of the individual pages (See Figure 10).

Figure 10 shows an example visualization of twelve usage traces, where the task was to find a specific piece of information on the U.C. Berkeley website. The pages along the highlighted path at the top represent the optimal path. By looking at the thickness of the lines, one can see that many people took the optimal path, but about the same number of people took a longer path to get to the same place. Following some of these longer paths, one can also see where users come to a page, and decide to backtrack, either via the back button or a link. Figure 11 shows a zoomed in view of one of the pages.

There are also several red arrows, which indicate that people took a long time before going to the next page. However, none of the red arrows are along the optimal path, meaning that people that took that path did not have to spend a large amount of time to get to the next page. One key feature of this visualization is the ability to zoom in and provide various levels of detail. For example, from the overview of the entire task, a viewer can see a red arrow indicating a long time spent on the page, but upon zooming in on that page the viewer would see that it is perhaps a very text-heavy page the user probably spent time reading. Providing the context of the task and a framework to add more details when needed, this visualization offers a number of simple, but very useful and quick analysis of the user experience.
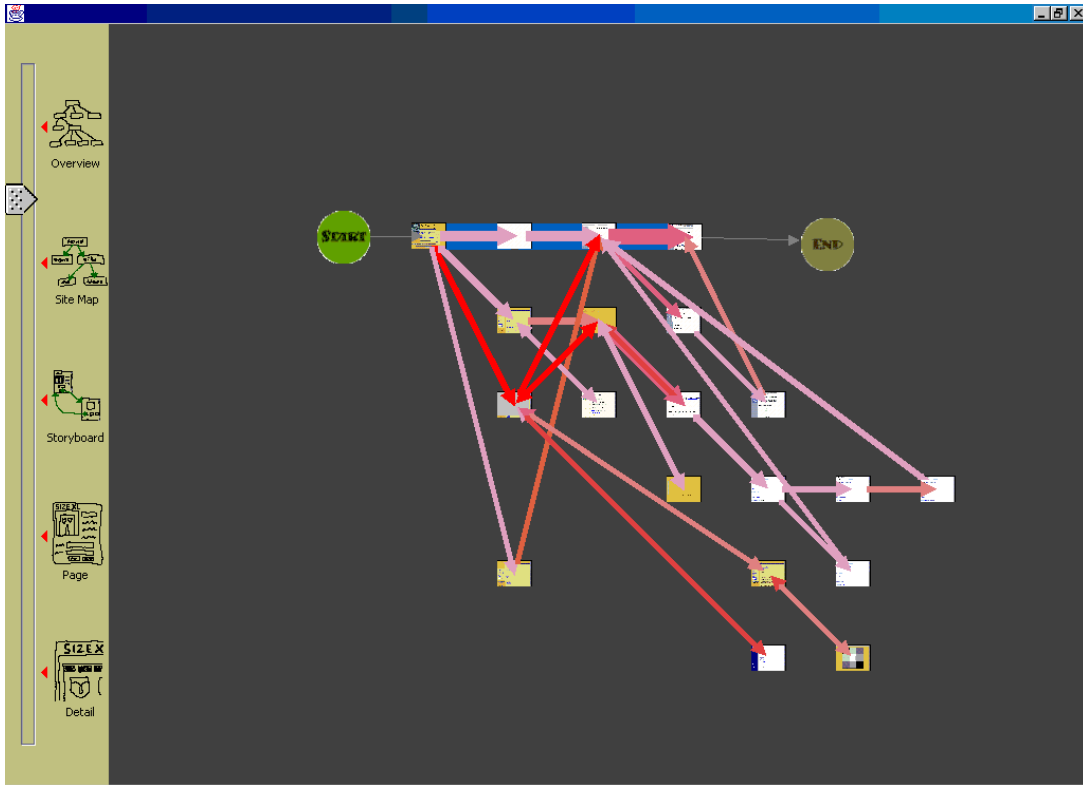
Figure 10. An example visualization of twelve usage traces for a single defined task. The circle on the top-left shows the start of the task. The circle on the top-right shows the end of the task. Thicker arrows indicate more heavily traversed paths (i.e., more users). Thick blue arrows mark a designer indicated optimal path. Darker red arrows indicate that users spent more time on a page before clicking a link, while the lighter pink arrows indicate less time.
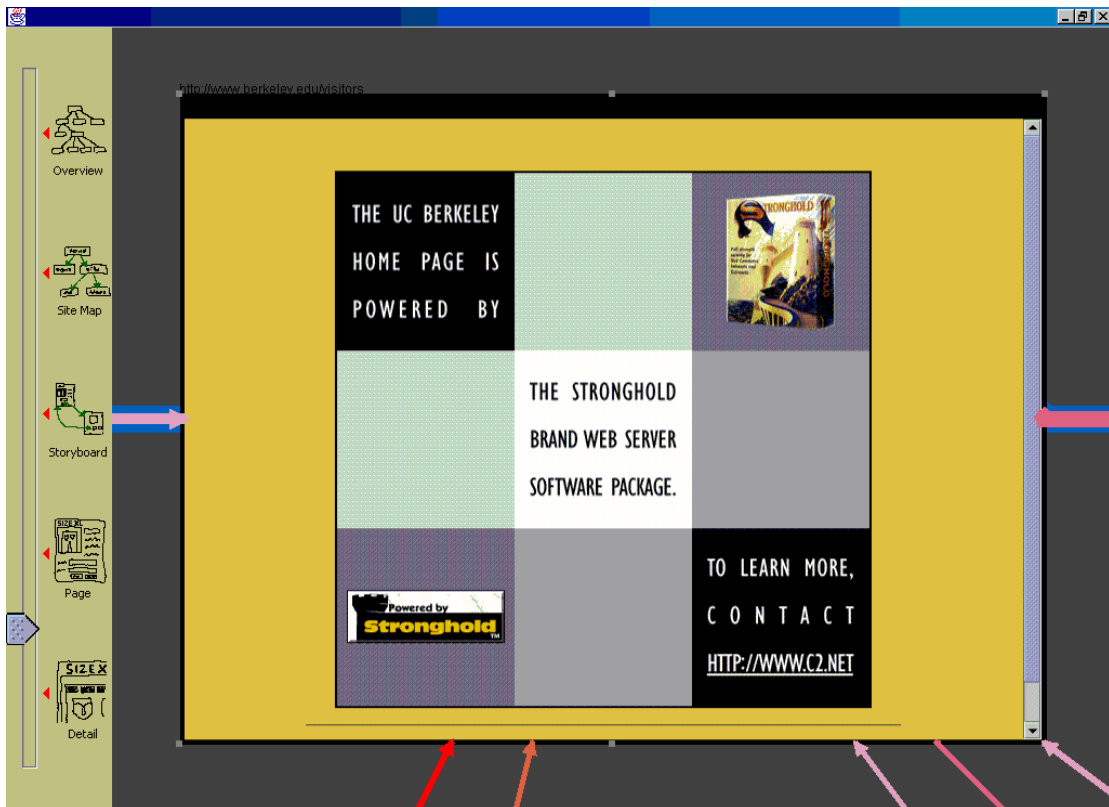

Figure 11. The zoom slider on the left is used to change the zoom level. Individual pages can be selected and zoomed-in on to the actual page and URL people went to.

3. RELATED WORK

In order to address some of the problems faced by client-side and server-side logging, the National Institute for Standards and Technology (NIST) has recently developed WebVIP [18]. Intended as a tool to help run usability tests, WebVIP makes a local copy of an entire site and instruments each link with special identifiers and event handling code. This code is activated when a link is clicked on. WebVIP shares some of the same advantages that WebQuilt's logging software has, such as better compatibility with existing operating systems and browsers (since only the HTML is modified) and some ability to run logged usability tests on sites one does not own. However, WebQuilt's proxy approach to logging lets it work without having to download an entire site, which is more realistic for many situations. WebQuilt avoids the problems of stale content and of invalid path specifications for complex sites, and also works with database-backed sites that dynamically generate HTML when page requests are made.

Another system that is similar to WebQuilt's logging software is Etgen and Cantor's Web-Event Logging Technique (WET) [10]. WET is an automated usability testing technique that works by modifying every page on the server. It can automatically and remotely track user interactions. WET takes advantage of the event handling capabilities built into the Netscape and Microsoft browsers. WET has more sophisticated event logging than WebQuilt currently supports, though there are plans for merging WET's advanced event handling capabilities into WebQuilt (see Future Work section). However, again, WebQuilt differs with its proxy approach, which again does not require ownership of the server and requires no changes to the server. These last two advantages are important when trying to accomplish web evaluations – the designers and usability team might not be *allowed* to make changes to or be given access to a production server.

Usability log visualization has a long history. Guzdial and colleagues review and introduce several desktop-based systems in [19]. There are other recent visualizations that use the notion of paths. For example, the Footprints web history system [20] displays aggregate user paths as hints to what pages have been followed by other people. VISVIP [21] extends the work in WebVIP and shows individual paths overlaid on top of a visualization of the website.

Two commercially available tools share similar goals to WebQuilt. Vividence Clickstreams [22] visualizes individual and aggregate user paths through a website. However, it uses client-side logging and has all of the problems associated with that technique. Furthermore, WebQuilt's visualization differs in that it combines aggregate path information with designated optimal paths to make it easier to see which pages people had trouble with. WebQuilt also uses a zooming interface to show different portions of the website, including screenshots of individual pages to provide better context.

Visual Insight's eBizinsights [23] is a sophisticated server log visualization tool for generating custom reports and interacting with the log data. While providing a number of interesting and useful chart-style visualizations, eBizinsights targets a marketing and management audience, which eventually will affect the designer decisions, but isn't intended as a tool for the designer

to use as he or she works. eBizinsights also uses server side logging, which, again, has all of the problems mentioned earlier with this technique.

The closest visualization work is QUIP [24], a logging and visualization environment for Java applications. WebQuilt builds on QUIP by extending the logging and visualization to the web domain. WebQuilt also adds in zooming, and uses screenshots of web pages for detail instead of abstract circles.

4. FUTURE WORK

We would like to extend our proxy logging system to overcome some of its current limitations, especially the handling of JavaScript. One possible approach would be to include a full JavaScript parser within the Proxy Editor module. The JavaScript code itself could then be updated to ensure that dynamically created links and pop-up windows pass through the proxy. Care must be taken, however, to ensure that the parsing and editing can be done in a reasonable amount of time, so the user's browsing experience is not adversely affected.

We would also like to capture a richer set of user interactions as they surf or perform a task. There is a wealth of data available to client-based loggers that our proxy currently does not collect. This includes page locations that are mistaken for links and clicked, page scrolling events, and more. By writing JavaScript code that captures these events and then inserting this code into pages as they are proxied, we may be able to remotely log some of these forms of interaction without any additional effort on the part of the user. We are currently investigating the approach used by WET [10], which captures low-level events on the Mozilla browser and on Microsoft's Internet Explorer browser, as well as work done by Edmonds on tracking user actions [25]. However, using some of these technologies could limit the variety of internet-enabled devices on which to use this proxy framework. Ideally, we would like WebQuilt to gather data from more than just traditional web browsers. We would like to see to our framework configured for interaction logging on hand-held computers, web-enabled cell phones, and other emerging internet-ready devices.

We are also looking at additional visualizations for displaying and interacting with the traces. There has been some work done in visualizing server logs [26-28] as well as visualizing individual and aggregate user paths [20-22]. We plan to re-implement some of the ideas demonstrated by these visualizations, and add in interactions and visualizations that are more useful for web designers. For example, the visualization in Figure 11 could be modified such that when zoomed in, the arrows could be re-anchored to show exactly which link was clicked on from a given page. This is an example of semantic zooming [29], where the details of the visualization changes depending on zoom level. Along the same lines, we would like to include methods for accessing more details, or filtering of information, depending upon the needs of the individual designer.

Lastly, we plan on doing a rigorous evaluation of WebQuilt with real designers and usability experts to assess the quality of the tool and explore of larger, more complex data sets.

5. CONCLUSIONS

We have described WebQuilt, an extensible framework for helping web designers capture, analyze, and visualize web usage where the task is known. WebQuilt's proxy-based approach to logging overcomes many of the problems encountered with server-side and client-side logging, in that it is fast and easy to deploy, can be used on any site, can be used with other usability tools such as online surveys, and is compatible with a wide range of operating systems and web browsers.

We have also described the architecture for WebQuilt, and shown how new algorithms and visualizations can be built using the framework. Again, we knew that we would not have all the solutions for analyzing and visualizing the captured data, so the system must be extensible enough so that new tools can be easily built. We have demonstrated one simple zooming interface for displaying the aggregated results of captured web traces, and are currently building more sophisticated visualizations and interactions for understanding the data.

WebQuilt can be found at:

`http://guir.berkeley.edu/projects/webquilt`

6. ACKNOWLEDGMENTS

7. REFERENCES

1. Beyer, H. and K. Holtzblatt, *Contextual Design: Defining Customer-Centered Systems*. 1998, San Francisco: Morgan Kaufmann.

2. Virzi, R.A., *Refining the Test Phase of Usability Evaluation: How Many Subjects Is Enough?* Human Factors, 1992. **34**(4): p. 457-468.

3. Nielsen, J. and T. Landauer. *A mathematical model of the finding of usability problems*. in *ACM INTERCHI'93*. 1993. Amsterdam, The Netherlands.

4. Spool, J. and W. Schroeder. *Testing Web Sites: Five Users is Nowhere Near Enough*. in *ACM CHI Conference on Human Factors in Computing Systems*. 2001. Seattle, WA.

5. *Access Log Analyzers*. http://www.uu.se/Software/Analyzers/Access-analyzers.html

6. NetRaker Corporation, *NetRaker Suite*. 2001.

7. Vividence, *Vividence Browser*. 2000.

8. Pikrakis, A. *MEMOIR - Software Agents for Finding Similar Users by Trails*. in *Third International Conference and Exhibition on The Practical Application of Intelligent Agents and Multi-Agents*. 1998. London, UK.

9. Chris Dodge, B.M., Hans Pfeiffenberger. *Web Cataloguing Through Cache Exploitation and Steps Toward Consistency Maintenance*. in *Third International World Wide Web Conference*. 1995.

10. Etgen, M. and J. Cantor. *What Does Getting WET (Web Event-Logging Tool) Mean for Web Usability?* in *Fifth Human Factors and the Web Conference*. 1999.

11. Davison, B. *Web Traffic Logs: An Imperfect Resource for Evaluation*. in *Ninth Annual Conference of the Internet Society (INET'99)*. 1999. San Jose.

12. Pitkow, J.E. and P. Piroli. *Mining Longest Repeated Subsequences to Predict World Wide Web Surfing*. Second USENIX Symposium on Internet Technologies and Systems. 1999.

13. Choo, C.W., B. Detlor, and D. Turnbull. *A Behavioral Model of Information Seeking on the Web -- Preliminary Results of a Study of How Managers and IT Specialists Use the Web*. in *1998 ASIS Annual Meeting*. 1998.

14. Tauscher, L.M., *Evaluating History Mechanisms: An Empirical Study of Reuse Patterns in WWW Navigation*, in *Department of Computer Science*. 1999, University of Calgary: Calgary.

15. ErgoSoft Laboratories, *ErgoBrowser*. 2001. http://www.ergolabs.com/ergoBrowser/ergoBrowser.htm

16. Tschalar, R., *HTTPClient V0.3-2*. 2000. http://www.inovation.ch/java/HTTPClient/index.hml

17. Sun, *Java Secure Socket Extension*. 2001. http://java.sun.com/products/jsse/index.html

18. NIST, *WebVIP*. 1999. http://zing.ncsl.nist.gov/webmet/vip/webvip-process.html

19. Guzdial, M., *et al.*, *Analyzing and Visualizing Log Files: A Computational Science of Usability*. 1994, Georgia Institute of Technology.

20. Wexelblat, A. and P. Maes. *Footprints: History-Rich Tools for Information Foraging*. in *CHI99*. 1999. Pittsburgh PA.

21. Cugini, J. and J. Scholtz. *VISVIP: 3D Visualization of Paths through Web Sites*. in *International Workshop on Web-Based Information Visualization (WebVis'99)*. 1999. Florence, Italy: IEEE Computer Society.

22. Vividence, *Vividence Clickstreams*. 2000. http://www.vividence.com/resources/public/solutions/demo/demo-print.htm

23. Visual Insights, *eBizinsights*. 2001. http://www.visualinsights.com/

24. Helfrich, B. and J.A. Landay, *QUIP: Quantitative User Interface Profiling*. 1999. http://home.earthlink.net/~bhelfrich/quip/

25. Edmonds, A., *Lucidity Project*. 2001. http://sourceforge.net/projects/lucidity

26. Chi, E., *et al.* *Visualizing the Evolution of Web Ecologies*. in *ACM CHI Conference on Human Factors in Computing Systems*. 1998.

27. Hochheiser, H. and B. Shneiderman, *Understanding Patterns of User Visits to Web Sites: Interactive Starfield*

Visualizations of WWW Log Data. 1999, University of Maryland, College Park.

28. Pitkow, J. and K. Bharat. *WebViz: A Tool for World-Wide Web Access Log Analysis*. in *First International Conference on the World-Wide Web*. 1994.

29. Bederson, B.B. and J.D. Hollan. *Pad++: A Zooming Graphical Interface for Exploring Alternative Interface Physics*. in *the ACM Symposium on User Interface Software and Technology: UIST '94*. 1994. Marina del Rey, CA.