

Off the Beaten Tracks: Exploring Three Aspects of Web Navigation

Harald Weinreich¹, Hartmut Obendorf², Eelco Herder³, Matthias Mayer²

¹VSIS, Department of Informatics
University of Hamburg, Germany

harald@weinreichs.de

²ASI, Department of Informatics
University of Hamburg, Germany

obendorf@informatik.uni-hamburg.de
mayer@informatik.uni-hamburg.de

³Department of Computer Science
University of Twente, The Netherlands

herder@cs.utwente.nl

ABSTRACT

This paper presents results of a long-term client-side Web usage study, updating previous studies that range in age from five to ten years. We focus on three aspects of Web navigation: changes in the distribution of navigation actions, speed of navigation and within-page navigation.

“Navigation actions” corresponding to users’ individual page requests are discussed by type. We reconfirm links to be the most important navigation element, while backtracking has lost more than half of its previously reported share and form submission has become far more common. Changes of the Web and the browser interfaces are candidates for causing these changes.

Analyzing the time users stayed on pages, we confirm Web navigation to be a rapidly interactive activity. A breakdown of page characteristics shows that users often do not take the time to read the available text or consider all links. The performance of the Web is analyzed and reassessed against the resulting requirements.

Finally, habits of within-page navigation are presented. Although most selected hyperlinks are located in the top left corner of the screen, in nearly a quarter of all cases people choose links that require scrolling. We analyzed the available browser real estate to gain insights for the design of non-scrolling Web pages.

Categories and Subject Descriptors

H.5.4 Hypertext/Hypermedia: User issues.

General Terms

Human Factors, Design.

Keywords

Hypertext, navigation, user modeling, clickstream study, browser interfaces

1. INTRODUCTION

The World Wide Web is a system where the only constant is change. New technologies and services have emerged on the server side, allowing people to use their Web browser for manifold activities and tasks: content management systems allow the reuse of layout information and the dynamic creation of documents; in-

formation can be updated as frequently as required for news pages and online shops. Functionality that used to reside on desktops is made accessible through Web interfaces, covering an application range from email, chat, and bulletin boards to complex applications like travel agencies, libraries and shops. Many of these changes could not be foreseen when the Web and its client software were designed as a hypertext information system.

While the technology serving the Web behind the scenes has seen dramatic changes, the user interface of current Web browsers and their integrated navigation support tools closely resemble those of browsers from the early days of the Web. Most of the changes are of technical nature, such as the integration of plug-ins and JavaScript, allowing new content types and a greater degree of interaction, but not directly relating to browsing support. Although change here was slow and subtle, some new interaction methods have been introduced that might influence user habits, such as tabbed browsing (Figure 3) and mouse gestures [27].

There is a similar imbalance in focus between server-side and client-side analyses of Web browsing behavior. Whereas page requests to Web servers and search engines are commonly logged and used for subsequent analysis of usage patterns, access to cross-site browsing patterns is barred as this data is only available on the client. Studies collecting clickstream data on client-side are surprisingly scarce: the last reported long-term studies of browser usage are at least five years old. Consequently, many aspects of how people use their Web browsers today are either unknown or unproven. As changes of the Web and in browser features could have a major influence on user requirements that might be visible in their actions, we felt the necessity for an update study. The existing long-term studies are of excellent quality but also of impressive age; newer quantitative studies relate to specific tasks, were often performed under laboratory conditions or had fairly small sample sizes and thus fail to give an account of everyday Web use. The study presented here captures a detailed stream of user actions and page requests over a long period directly in the working environment, which enabled us to give a fairly consistent account of periods of users’ browsing activity.

Based on this data some recommendations for Web design are reconsidered. Navigation actions indicate that the Web has evolved to a hybrid between information system and online application. Users’ speed of navigation is high, even on pages that are rich in content and links. This rapid activity is set in relation to the measured responsiveness of the Web to identify limiting factors. The positions of selected links provide insights in user habits and issues with current page design. We finally find that in contrast to common belief, Web browsing is an activity that needs to share the users’ attention – and screen space – with other applications.

2. RELATED STUDIES

How users browse the Web has been subject to research almost from the start. Many previous studies are based on data found in server logs, analyzing various aspects of user navigation [18; 19; 35; 40]. However, server logs only report on user actions for a single site, caching mechanisms hide some of the interesting revisits, and interaction with the browsing application remains concealed. In contrast, several observational short-term studies examined in detail on how people use the Web; screen captures, video coverage, and Web diaries have been successfully applied to extract contextual information of Web browsing activity [5; 17; 25; 38]. However, as these studies are very time-consuming, they can cover only a very limited period of user activity and the size of the user sample is usually small. Also, these studies can hardly report on quantitative changes and may be biased by the test environment setting.

Since 1994, only three long-term studies have tried to analyze the clickstream created by the interaction of user and Web browser. Of these studies, only two recorded navigation actions – and both are over 10 years old by now. In 1994, the first long-term client usage study was described by Catledge and Pitkow [3]. They used an instrumented version of XMOsaic on their departmental Sun workstations to analyze the actions of users that lead to page requests. Their study covered the Web browsing activity of 107 users for 21 days. Catledge and Pitkow identified different navigation strategies and found that users operate only on small areas within sites. The back button was found to be the second most important user interface element accounting for 41% of navigation actions, second only to hyperlinks (52%). Other actions, such as “archiving” actions that lead to saving or printing a page, were unexpectedly rare.

In 1995, Tauscher and Greenberg focused on history support and analyzed the revisitation behavior of their participants [42]. They defined the “recurrence rate” as the probability of the next visited page having been visited before by the same user. A measured rate of 58% led them to the conclusion that the Web is a ‘recurrent system’, although they also found that most revisits are short-term: only 15% of page revisits are not in the list of the last 10 visited pages. Finally, in 1999, Cockburn and McKenzie retroactively analyzed the Netscape history and bookmark files of their participants [7]. Their users visited more pages per day, but at the same time the revisitation rate increased to 81%. They also found that Web use is rapidly interactive and users often visit many pages within seconds successively. They gained new insights into homepage and bookmark use. People tend to have many bookmarks, but use only few of them.

3. THE WEB BROWSING STUDY

25 unpaid volunteers participated in the Web logging study. Two 1.5 hour interviews at the beginning and end of the study were used to gather demographical data and information on general use, but also to confirm some interpretations of the captured data, e.g. concerning the use of multiple windows. Six (24%) of our participants were female. Ages ranged from 24 to 52 years (mean: 30.5). All our participants were experienced, browsing the Web for 3 to 12 years (mean: 8). Most of them came from Germany and the Netherlands (three Germans worked in Ireland and New Zealand) and all interviews were conducted in their native language. While all eight participants from the Netherlands (32%) worked as university employees in computer science, nine participants (36%) had a different background: two worked in psy-

chology, and one each in sociology, geology, electrical engineering, trading, coaching, history, and photography. Seven additional participants began the study, but dropped out for personal or technical reasons and had to be excluded from the analysis. The length of the study varied individually from 52 to 195 days (mean: 105 days). We were able to confirm (see Section 3.3) 137,272 user-initiated page visits to 65,643 distinct URIs and 9,741 different domains. Personal Web usage varied widely in browsing style and activity: the participants visited 19.5 to 204.8 pages per active day (every day in which at least one event was logged).

3.1 Pre-study challenges

Given the importance of the Web, it might seem surprising that only so few quantitative long-term studies analyzed the browsing behavior of Web users. This may be explained by the social and technical issues we encountered during the preparation of the study. Today, browsing is considered a personal activity, even if logging only takes place at the workplace. The Web is now used for many confidential tasks, such as online banking, shopping or writing e-mails. After initial informal surveys, it became clear we had to make use of a capturing system that did not record usernames or passwords and ignored secure connections. The participants were also given the option to screen the content of all log files before transmitting them to us.

Some prospective participants were also concerned that the installation of such ‘spyware’ might have negative influence on the stability or performance of their personal computer. These concerns were not unfounded: in one of our pilot studies instrumenting the Internet Explorer to record user actions and page requests led to stability issues when different Explorer versions were used or new plug-ins were installed – unacceptable for a long-term study where the browser is used daily as a production tool.

The potential participants of this study used many different browsers with different browser extensions. In consequence, the software recording user actions had to be compatible with at least most of these systems. We opted for a solution based on an intermediary intercepting the traffic between browser and Web. The release of Firefox 1.0 in 2004 finally provided a platform for instrumenting a popular Web browser; its interface is familiar to users of the Internet Explorer and new security risks of the internet motivated several of our participants to ‘upgrade’. For us, the open source status of the browser made it possible to implement a logging mechanism for all user interface actions.

3.2 Browser logging environment

The browser logging environment consisted of two complementing extensions: every participant had an intermediary installed to filter all transferred pages. It added JavaScript code to every page; when executed, this code read several browser parameters and communicated them to the intermediary. Thus, detailed knowledge about every link anchor selected, every form submitted and the current state of the history of the browser window was obtained. Furthermore, it allowed discriminating all windows and frames of the client and their current dimensions. The intermediary also analyzed all transferred documents and recorded descriptive data about their length, links and contents. Technically, this system was based on IBM’s WBI [2] and the Scone framework [32].

15 of the 25 participants used an instrumented version of Firefox that was modified to record the interaction with all user interface

widgets. These participants were either already using Firefox as their preferred browser or took the opportunity to switch. Using exact timestamps, this second clickstream log could be merged with the navigation log of the intermediary to gain more detailed and accurate data.

3.3 Data processing and consolidation

The first analyses of the recorded log files showed that serious data pre-processing was necessary to get valid results, as many of the entries were not directly related to user actions [cf. 45]. We found several different reasons responsible for these artifacts: HTML framesets break the document metaphor – what is visible for the user does not originate from a single HTML document and many unusable events are created as every subframe creates a page request to a different html file. We identified frames by the name of the frame and the parent window, which was read with embedded JavaScript code. The interpretation of these data required not only to collate the requests to one user action, the events were also often ambiguous, as we could not define which address was the most important one for a user action: if the user selected a link in the left navigation frame and a page was loaded in the right frame, the action occurred on the left page but the reaction affected the right page with another URI. The interpretation has consequences on load times, revisitation rate, recorded document sizes as well as the link positions. Therefore, we excluded frame pages for most analyses.




Figure 1: Proportions of artifacts for users w/o ad blockers

Another significant problem was caused by advertisements¹. JavaScript-initiated advertisements in pop-up windows are not deliberate user actions, events relating to such windows and page requests were therefore excluded from the study. A statistically even more relevant advertisement technique is based on iFrames that allow embedding other HTML pages in a Web document. Our data showed that iFrames are currently mainly used to dynamically include advertisements. For the group of participants that did not apply any kind of ad-blocker (8 users) frame and advertisement artifacts represented over 28 percent of page requests (Figure 1). This is remarkable, as it does not even consider online promotion realized as pure text, embedded images or flash animations.

A third source of non-user initiated page requests were automatic page reloads, mainly caused by news pages. These became visible as peaks in the stay time distribution of certain users. In some

¹ We identified advertisements by different lists of known servers, typical URI patterns and equivocal frame names.

cases, embedded JavaScript code let the browser refresh a page contents after a certain interval, in other cases even external applications like instant messaging agents were responsible for the artifacts. These contributed nearly four percent of the page requests. However, the ratio differed severely between participants: some did not show any periodically reloaded pages, others over 20%.

As became clear during the analysis of the comprehensive datasets we had gathered, data cleaning and confirmation of user-initiated events were important to be able to relate recorded events to user actions. Previous studies did not use similar data consolidation methods, probably because the amount of such ‘noise’ was lower in the past: in 1995, advertisements were still hardly known on the Web, and Bruce McKenzie [7] told us that even in 2000 the effect of such requests could be neglected.

4. RESULTS

The prominence of advertisements can safely be interpreted as an indicator for the increasing commercialization of the Web: sites providing e-commerce, news and entertainment have become the most popular destinations, whereas ten years ago Web usage was focused on information and content delivery [9]. Based on the collected data, we investigated whether these changes of the Web also induced changes in user navigation behavior.

4.1 The new character of Web navigation

In this section, we examine the actions that our participants employed to initiate page visits. We call these events *navigation actions*. Apart from selecting links, users can trigger navigation actions in different ways: entering URIs directly into the address bar of the browser, using different browser history mechanisms to revisit pages seen before, or submitting information via forms to interactive Web services, such as search engine.

The latest reported distributions between the various navigation actions that are based on long-term data date back to studies from 1995 and 1996. The comparison chart (Table 1) shows some major differences, which reflect both the changed nature of the Web and the way users interact with browser interfaces.

Link following has remained the most common navigation action, accounting for about 45% of all page transitions. Direct access to pages via the bookmark menu, bookmark toolbar (which was not present in previous studies; see Figure 2), home page button or the address bar has remained stable at about 10% as well. The detailed Firefox log as well as the interviews revealed, however, that our users had different preferences to access frequently used pages: some mainly used the bookmark menu, others solely preferred the bookmark toolbar and a few had the custom to type the URI of their favorite pages into the address bar, using its auto-completion function when available. These different behaviors show that customization of the interaction with the browser is necessary. On the other hand it also indicates that none of the current revisitation tools is entirely satisfying [1; 20].

The comparison chart shows a major increase in navigation actions that lead to opening a new browser window. In the mid-nineties, this event accounted for less than 1% of all navigation actions, compared to over 10% nowadays. However, while formerly only the explicit action of opening a new window using the associated menu item was registered², in this study other actions

² This follows from the much higher number of ‘close window’ and ‘exit program’ than ‘new window’ events reported in [3].

Table 1: Comparison chart of three long-term studies

	Catledge & Pitkow ³	Tauscher & Greenberg ³	This Study
Time of study	1994	1995-1996	2004-2005
No. of users	107	23	25
Length (days)	21	35-42	52-195, $\sigma=105$
No. of visits	31,134	84,841	137,272
Recurrence rate	61%	58%	45.6%
Link	45.7%	43.4%	43.5%
Back	35.7%	31.7%	14.3%
Submit	-	4.4%	15.3%
New window	0.2%	0.8%	10.5%
Direct access	12.6%	13.2%	9.4%
Reload	4.3%	3.3%	1.7%
Forward	1.5%	0.8%	0.6%
Other	-	2.3%	4.8%

could also result in opening a new browser window. These actions include following hyperlinks with `target="_blank"` as an attribute, starting the browser manually or from other applications and the using the ‘open link in new window’ or ‘open link in new tab’ entries of the browser’s context menu (Figure 3). Nevertheless, this confirmed that it has become common behavior to open more than one window while browsing the Web.

Accounting for over 15% of all navigation actions, form submission has become a key feature of user navigation as well, as it is a required interaction mechanism with service-oriented Web sites. 43% of all form submissions involve queries to search engines, followed in popularity by an online dictionary and travel planners.

By contrast, the share of back button actions has dropped from over 30% in the mid-nineties to less than 15% presently. This number includes backtracking multiple steps via the back button’s pull-down menu, which contributed less than 4% to all backtracking actions. The browser history is not specifically listed in the comparison chart, as it is hardly used – merely 0.2% of all page requests were initiated from the browser history. Only two of our twenty-five participants stated to use it from time to time, while ten participants even weren’t aware of the function at all.

The reduced usage of the back button, in combination with an increase of ‘forward navigation actions’ – following links, submitting forms and opening new windows – might indicate that users return less frequently to previously visited pages. However, the *recurrence rate* – the percentage of page revisits [42] – decreased to a much lesser extent; from about 60% to 46% (excluding frame pages; see above, [45]). One explanation is that most sites now offer structural links on every page that allow returning to the home page or a landmark page without using the back button. However, there might also be a relation with the increased amount of submit and new window actions. This issue will be explored in the following two subsections.

³ We recalculated the values of [3] and [42], as their notion of navigation action differed from each other. To make all values comparable we applied the definition used in this study.

Form submission and backtracking

The increased number of form submissions confirms another change of the Web: the move from a hypertext information system with mainly static pages to a hybrid between ‘classical hypertext’ and service oriented interactive systems, such as search engines, dictionaries and travel planners. The latter category of sites is often more similar to desktop applications than to information-centered hypertext: whereas navigation in hypertext involves orienteering behavior with frequent backtracking, interactive applications are mainly used for ‘getting things done’. This would imply that backtracking is less prominent during these activities. In order to confirm this hypothesis, we compared the backtracking usage of the top third form submitters of our participants with the remaining participants. The frequent submitters used the back button less frequently (9.2%) than the others (16.2%), a difference that is marginally significant ($t=2.715$, $p=0.012$).

Dynamic, interactive pages pose several challenges to the browser’s history mechanisms. First, the browser history does not take form submissions using the http POST method into account. Unless the URL of the resulting page explicitly contains the issued parameters, users cannot revisit earlier created documents such as travel plans without going through the process of entering the data again. Once the browser window is closed, the travel plan is lost. The same problem arises for documents that users might want to keep for future reference, such as order confirmations and flight reservations. Unlike in static hypertext, these pages are volatile, even if they contain information that will remain most relevant in the near or more distant future. Whereas users can save or print the information, most participants reported that they almost never did so. The Internet Explorer for Mac OS X features the Scrapbook, which provides an integrated interface for storing an exact copy of the Web page as it appears in the browser window. With the advent of service-oriented sites and volatile pages, similar functionality might be needed in other Web browsers as well. Travel plans, flight reservations and order confirmations should be treated as documents; context-sensitive functionality for storing, retrieving, opening and printing – like in regular office applications – appear to be essential in these situations.

To follow the analogy with office applications even further: as mentioned earlier in this section, the concept of hypermedia navigation is often replaced by a concept of interaction with an application. Hence, navigation tools such as the back button lose their original meaning in these contexts. While using an interactive Web service, the back button bears more similarity to the *undo button*: users press back to correct errors. To avoid potential disruption of the interaction, some online services disable the use of the back button by opening a pop-up window without navigation toolbars for sequences of interactive forms, or they explicitly advise the user not to use the back button.

In conclusion, browser interfaces do not provide appropriate functions for service-oriented sites, although these sites play a prominent role in Web usage. We think that one major challenge for the next generation of Web browsers is to reconcile the two different Web usage contexts – hypermedia navigation and interaction with Web-based services.

Windows, tabs and backtracking

As reported in the beginning of this section, our participants tended to use multiple windows to a large extent. The strategy of




Figure 2: Bookmark toolbar and browser history in sidebar

opening new windows seems to offer several advantages. Our users reported that multiple windows allowed them to ‘compare search results side by side’ and that ‘pages can be loaded in the background’ while they continue their navigation activities. Keeping search results and resulting navigation trails in separate windows also reduces the risk of losing the path to a decisive page because of backtracking to the result page. Several participants also had the habit of keeping a browser window with their favorite news site open in the background.

A comparison of the frequency of backtracking with the usage of multiple windows or tabs showed a correlation between these two navigation actions. The group of participants with the top third of new window events employed the back button to a lesser extent (10.2%) than the bottom third (16.4%); it is plausible that multiple windows are used as an alternative to backtracking ($t=2.509$, $p=0.026$). In addition to multiple windows, Firefox provides ‘tabbed browsing’ – several pages can be opened simultaneously in different browser tabs (Figure 3). Six participants who used tabs frequently, were backtracking less often (9.9%) than the remaining seven Firefox users (18.3%) that hardly opened any tabs ($t=2.311$, $p=0.038$). One participant explained he used ‘new tabs for closely related tasks and new windows for parallel tasks’.

A more disturbing consequence of the use of multiple windows is that it disrupts the concept of the back button. Its principal functionality is to return to recently visited pages. If users followed trails in multiple windows or tabs, the recent visit history is split into separate stacks, with no temporal relation. Moreover, each individual stack does not include actions from the originating window. Hence, users need to remember what actions they performed in which window or tab in order to relocate a previously visited page. This places a high cognitive burden on the user, in addition to the already demanding task of keeping track of their location in the Web [10]. Handling multiple windows in information systems was already reported to cause disorientation in pre-web studies [13]. The above concerns were confirmed by our participants; several said that they find many open Web documents hard to manage, in particular because the page titles displayed in task bar and tabs are often not helpful.

The more prominent use of multiple windows requires a major rethinking of the history mechanisms of browsers. It provides new challenges to the often criticized [6; 12; 21] yet often used back button. A linear history of most recent revisits, as proposed by




Figure 3: Tabbed browsing: opening a new tab from a link

[42], does not reflect the character of parallel trails and the unrelated back button stacks do not take the temporal relations between the trails into account.

4.2 The velocity of Web navigation

The speed of interaction with the Web browser is another aspect of Web navigation: how much time do users spend on Web pages, i.e., how much time do they take to read the page and think about the available options, before they perform their next action? Although the time between page requests can be gained from server logs [35] and the browser history [7], the data recorded with a client side logging software is more exact. Our software recorded the time between the display of the first parts of the HTML document and any subsequent navigation action in the same window that would lead to the request of another page. In consequence, delays – such as the time before the browser begins to load a page – could be differentiated from the stay time. Navigation actions that did not lead to a new request to the same site were also considered, like the selection of an external link, as well as backtracking – which is usually hidden in server logs since the page is loaded from the browser’s cache, and leaving a page by closing it. The capturing software did also distinguish between multiple windows and tabs, so it could be identified, when a user opened several pages at once from a hub page, but read them one after the other. Although this method removed incorrect values for the time a page is opened in a browser, there will be a bias towards long stay times, as open windows that currently do not have the user’s attention could not be identified and excluded.

This study confirms the rapid interaction behavior with heavy tailed distribution already reported by previous studies [3; 7; 11]. Our participants stayed only for a short period on most pages: 25% of all documents were displayed for less than 4 seconds and 52% of all visits were shorter than 10 seconds (median: 9.4s). However, nearly 10% of the page visits were longer than two minutes. Figure 4 shows the distribution of stay times grouped in intervals of one second. The peak value is located at stay times between 2 and 3 seconds; they contribute 8.6% of all visits.

We first assumed that most of these short and very short stay times represented *revisits*, e.g., visits to pages that had been seen recently or that were used frequently and therefore well known, like the home page of the browser. To analyze how much time users take to read new Web documents, all revisited pages were




Figure 4: Distribution of stay times for all participants

excluded from the statistics. The effect on the distribution of stay times was not as strong as expected (Figure 4): still over 17% of all new pages are visited for less than 4 seconds, nearly 50% are shown for less than 12 seconds and 11.6% are displayed for more than 2 minutes (median: 12.4s). However, a fifth of the 11.6% were visits of over 30 minutes to up to 5 days – most of these events are most likely created by unattended browser windows that were left open in the background of the desktop.

The tendency for very short page visits on the Web might have two reasons: either it expresses a cursorily and scanning usage behavior or it might characterize that many of the visited pages offer only little information and few navigational options. To verify this, the average number of words and links of the documents was calculated and compared with the stay times.

These results are based on nearly 60,000 first-page visits. The average number of words per page (only depicted text, not considering any markup code or any embedded objects or graphics) is 551 words ($\sigma=811$)⁴. The page stay times are dependent on the page size, but less than expected: pages visited for less than 12 seconds (they contribute about 50% of all requests) had an average number of 430 words. This is significantly lower than the mean size of documents with a longer retention time ($t=36.197$, $p=0.000$), but it is apparent that no person can read a full page of this length that quickly. Figure 4 (upper graph) illustrates the average number of words per page grouped in intervals of 2s staying time. Such a difference was also found for the number of navigational options per page (Figure 5). On average all visited pages had 53 hyperlinks⁵ ($\sigma=58$). For pages with a staying time of less than 12s the average number of links was 46 and significantly lower than for the remaining documents ($t=30.659$, $p=0.000$).

In consequence our participants often did not take the time to completely read every page, but they regularly just seemed to glimpse over most of the information offered, before they perform their next navigation action. Such a scanning behavior of Web

⁴ Outliers were removed (using a 3σ limit), as very few Web pages were atypically long and biased the sample. Average with outliers: 648 words ($\sigma=2342$).

⁵ Again, outliers were removed. All visited pages had an average of 61 hyperlinks ($\sigma=122$).




Figure 5: Correlation between stay time and page contents

users was already reported by [26] and [41] who observed it at several controlled Web usability studies. The data of this long-term study supports these results, indicating that the ‘scannability’ of information and hyperlinks as well as their intelligibility seem to be essential for the usability of Web pages (see Section 4.3).

Does the responsiveness of the Web suffice?

The participants of this study viewed many Web pages only for a very short time before taking the next action. This rapid interactive behavior requires a system that reacts quickly, too. Though the performance of the Web was often criticized in the nineties and even called ‘the most common usability problem of the Web’ [22], this shortcoming was *less* of a topic for the last five years. While the *digital divide* excludes many low-income families and the citizens of developing countries from the resources offered by internet and Web, a continually increasing number of users have broadband access [14]. It seems the Web has either become responsive enough to meet its users’ needs or they have got accustomed to inevitable delays, as new studies show that people are now more concerned with quality, privacy and security on the Web [8, 44]. We recorded transfer delays to estimate if the Web was fast enough to meet current user interface standards.

As the application areas of the Web changes from an information resource to an online application system, the performance requirements for hypertext systems as well as office software should be met. Many studies have confirmed a maximum reply time of 2 to 4 seconds; otherwise productivity degrades significantly [37; 16, p.74; 39, p.297].

Our participants had different types of internet connections. 14 participants made use of fast leased lines (at least OC-3), eight people used ADSL broadband access and three had only an ISDN dial-up connection. The logging software⁶ recorded two values within the browser for every page request: First, the delay between the user action that led to the new page request and the start

⁶ The intermediary added JavaScript code to every Web page. The script took timestamps of user and browser events, allowing us to get the actual response times as perceived by the participants [36]. This data was transmitted to the intermediary and stored in the event log.

of the page rendering was taken; second, the time to download the entire document including all embedded objects was registered.

The delay between a user action to request a new page and the arrival of the first data packages at the browser is critical, as the user has to pause and wait until this process is finished, before he may continue to use that browser window. We regarded only first-time page visits, as revisited documents are usually loaded from a proxy or the browser cache. The distribution (Figure 6, upper graph) shows that most requests are answered quite quickly: the median delay until the browser received the first parts of the document was 1.0 second. However, 90% of all replies took longer than 0.7 seconds. Consequently, data processing of the browser and the three way handshake of TCP/IP are already limiting factors for the performance of the Web. Furthermore, over 9% of the requests could not be answered within 4 seconds, and about 4% took more than 10 seconds or were not answered at all. These long delays are not only caused by low bandwidth connections and slow responding servers, due to the way the times were detected in the browser, they also stand for requests that led to an erroneous response (for instance caused by a removed document). 4% is actually the average number of broken links on the Web [23].




Figure 6: Response times for new pages and sites

A user request to a new server requires that the domain name is resolved, making these requests presumably slower. The lower graph of Figure 6 shows the average response times of all first-time visits to Web servers. The median delay was already 1.7 seconds and even 7% were not properly answered within 10 seconds. This indicates that transfer delays as well as broken links are a more severe problem for site-external than local references.

The ‘pure download time’ for Web pages, i.e., the time from the arrival of the first packet of the HTML document to the last part of every embedded object, shows a quite similar characteristic: 49% of all documents were downloaded in less than one second, with a tendency to very short times, and nearly 86% were fully transmitted in less than 4 seconds (Figure 7). However, this does also mean that over 14% of the page transfers could not be completed in 4 seconds.

As this result might be distorted by the three ISDN users, the download times were discriminated by bandwidth (Figure 7). Expectedly, the performance of the Web is poor for ISDN users: less than 60% of the documents could be completely downloaded within 4 seconds. But also for broadband users, the delays were often insufficient: 9% (leased line) or 16% (ADSL) of the trans-




Figure 7: Download times by bandwidth

fers were not completed in 4 seconds – even not considering the additional median round trip and processing time of one second.

At first glance, a delay of about 1 second before the browser starts to display the new page seems to be acceptable; however, when the rapid navigation behavior of our participants and the requirements of interactive online applications are considered, this is comparatively long and may certainly have negative influence on the overall user performance. Especially external links are potentially critical, as they bear longer delays and more erroneous references. Further research seems to be necessary.

4.3 Within-page navigation

If a Web page is too long to be displayed on one screen, the user has to *navigate* on the page by *scrolling* it. Scrolling is often considered problematic, as it can result in user disorientation: the reader may lose track of the context as the main headers, the site identifier and the main navigation elements move off screen. Furthermore, scrolling increases the cognitive burden: while long pages require the reader to remember information that scrolled off the screen, short pages allow comparing all available options side by side. Therefore, especially for *entry pages* and *navigation pages* guidelines recommend to fitting the page on one screen and show all options immediately [24; 30].

On the other hand, for *content pages* long documents can also bring advantages, as they can be read and printed without the need to flip from one page to another. Still, wide pages that require horizontal scrolling should be avoided for any kind of Web document, since they bear the risk that users have to move their viewport in two dimensions, wide text lines have a low readability, and the printout of these pages may be cropped [24; 28].

The short stay times on Web pages seem to justify short, non-scrolling pages; they pose the question whether users take the time to scroll at all. As scrolling is especially critical for navigation pages, we concentrated on pages that were used to navigate to another page and analyzed all clicks on hyperlinks as they comprise the most important navigation action on the Web.

A comparison of the link click positions and the browser viewport size shows that most links (76.5%) were selected within the region visible on load time (Table 2). Although users hardly scroll horizontally to select a link (altogether 0.4%), over 23% of the link clicks were below the initially visible region. Even for pages with stay times below 12 seconds, over 20% of the selected links required scrolling. Actually, the number of clicked links ‘below

Table 2: Location of selected links

	Visible Area	Right of Visible Area
Visible Area	76.5%	0.3%
Below Visible Area	23.1%	0.1%

the fold’ is higher than expected considering the short median stay time on pages.

However, the position of selected links is not only influenced by the scrolling habits of users, but also by the location of the available links. To get an overview of the navigation activity in the different screen regions, a map of all link clicks was created by grouping them in sections of 40 by 40 pixels. From over 27,000 recorded clicks, 93% were within an area of 1040×1600 pixels, which is depicted in Figure 8.

The most actively selected area is located in the upper left corner of the map: over 45% of all user clicks occur in a region of the browser window that is slightly larger than the upper left quadrant of the initially visible page area (520×400 pixels). Furthermore, in Figure 8 the horizontal and vertical navigation bars that provide the main structural hyperlinks on most major sites become visible, albeit the emphasis is on the left part of the horizontal and the upper part of the vertical bar. It should be noted this does not mean that these areas are generally suited best to place links, yet pages with a matching layout will meet the expectations of many users and increase the consistency of the Web.

The image shows an unexpectedly active area in the lower part of the documents, at about 600×1000 pixel and below, which required scrolling on all of our participants’ screens. It turned out that these clicks are mainly caused by Google’s placement of the ‘next page’ link. In fact, this detail indicates how often users have to scroll on the Web, as all popular search engines provide result lists that do not fit on a single screen – if we remove all search engine pages from the dataset, the vertical scroll ratio drops from 23% to 18%. Without taking transfer time into account, requiring users to both scroll and to flip through pages seems to be inefficient.

The real browser real estate

Web page optimization for the users’ rapid navigation behavior includes incorporating scannability and providing all relevant information at a single glimpse. Creating Web pages that do not require scrolling, however, requires knowledge of the available browser real estate. Although Web content should in principle be accessible to all people with different abilities and arbitrary hardware [4], for an aesthetic appearance device-specific designs are often inevitable. Particular style sheets should not require particular screen dimensions, but in practice, style sheets and graphical elements (like bitmap graphics) need to consider the available space of the browser. Thus, many Web authors try to optimize layout and design for a specific resolution.

Over the last years, the average screen resolution of personal computers has increased [43]. Until 2004, many experts recommended to base Web design on a resolution of 800×600 pixels, whereas recently the migration of the layout for a resolution of 1024×768 pixels is being recommended⁷ more often. We wanted

⁷ Experts who support ‘wide screen designs’ are [33], [34] and for intranets [31].

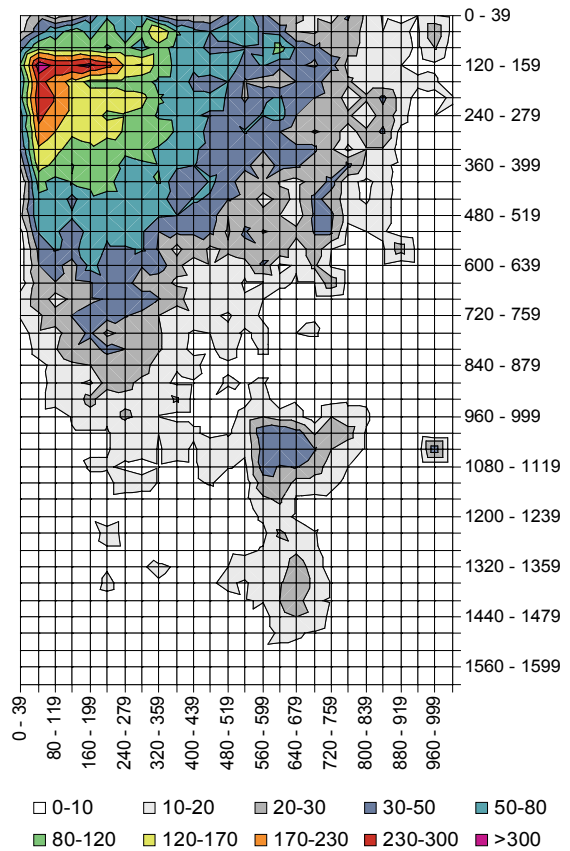


Figure 8: Link activity areas of Web users.

to find out whether the full resolution is really available for Web pages, or if technical or personal reasons limit the browser viewport. Due to browser internals, we could only record the browser viewport size for 20 participants. The size was recorded in more than 12,000 instances, usually after selecting a link. Frame pages were again excluded, as their sub-pages might adulterate the statistics. All participants used a screen resolution of at least 1024×768 pixels, twelve even had a higher resolution and three made use of two displays.

We could identify two groups: while eleven participants had the browser in full screen mode most of the time, nine participants preferred a smaller window size for at least half of the recorded time. However, the users of a maximized browser window also rarely had the full desktop resolution available: office toolbars, instant messaging clients, browser toolbar extensions as well as browser tabs and the side bar (Figure 2) took some of the vertical and horizontal space for most of them. Furthermore, some participants preferred to maximize the browser manually, leaving a border of several pixels around the window unused.

The nine participants with smaller browser windows (two of them with a screen width of 1024 pixels) left on average about 160 pixels of horizontal and 170 pixels of vertical space unused – scrollbars, pull-down menus, toolbars and the windows task bar already considered. For the users with a screen resolution of 1024 pixels, the average available document width was only about 890 pixels. They stated to prefer the windowed mode for several reasons: it permitted them to view and select other windows directly on the screen making the organization of several applications less com-

plicated. Furthermore, a narrower window would improve the readability of many documents as the lines of text were shorter.

Consequently, accessible sites should consider that people have different preferences using their desktop system and resizing their windows. If sites do not want to displease their visitors by forcing them to maximize their browser window or scroll horizontally, designers should not count on having exclusive rights to the screen real estate: flexible layouts leaving at least 15% of the screen width obtainable should instead be applied.

5. LIMITATIONS OF THE STUDY

Though the recorded data of this study is extensive and detailed, it has its limits, as it misses contextual information. The extent of logging was restricted technically, as data capture was limited to the browser, ignoring related software that was used in conjunction with the Web client, such as word processors, e-mail agents and other office applications. Furthermore, the data of clickstream logs have a limited expressiveness, as aims and tasks of the users often stay below the surface. This makes their contextual interpretation inherently difficult and additional qualitative information is needed to support a detailed task-related evaluation of the data. The two 90-minute interviews conducted at the beginning and end of the study could only deliver scarce data for a substantial qualitative analysis.

Another critical aspect of such a study concerns social and professional differences in user groups. Although we tried to recruit many different participants, all were frequent computer users with long internet experience. Still, the variance in the captured data was fairly large for almost all aspects of navigation: the number of visits, page vocabulary and use of search engines differed between the individuals. Also, the applied navigation habits, especially to directly access documents and to revisit pages varied. However, the remarkable differences in this small participant group did already reveal that Web browsers are used with various personal preferences and that individual users have particular demands. Web browsing is no longer simply navigation – and to a decreasing degree it is pure information and document retrieval. This should be considered for the design of future Web clients, Web sites and Web applications.

6. SUMMARY

This paper presents results of an extensive long-term study that captured the Web browsing behavior of 25 participants with diverse backgrounds and tasks. Although some results from former studies could be confirmed, we found evidence for a change of interaction with the Web. A strong increase in the proportion of submit events indicates the rising number of dynamic Web pages and ‘Web applications’; the increased number of new window events suggests that interaction with the Web client is changing from single-window hypertext navigation to a new mode where several paths are followed in parallel. Furthermore, some browser windows stay open for a long time or are dedicated for special online services like a news site or an online dictionary. Since navigation support of current browsers is still very similar to the early days of the Web, new problems surface: backtracking and history were not designed for dynamic pages and online applications with volatile contents; the concept of the back button fails for users of multiple windows or tabs as every document area has its own history stack. In consequence, users are faced with a new cognitive overhead when they browse the Web.

Our results confirm that browsing is a rapidly interactive activity. Even pages with plentiful information and many links are regularly viewed for a brief period – an interesting background for Web designers, who could focus on offering concise pages that load fast. Interface standards for the Web would help to make navigation on unfamiliar sites easier and quicker [15; 29]. Also, response times are often still not appropriate and Web developers should consider the additional delay before the browser starts to display a document when they estimate download times.

The analysis of link click positions shows that users do scroll frequently – even on navigation pages and on pages they visit only for a short period. Still, most selected links reside in the upper left quarter of the browser window. Placing the most important links in this area will increase consistency with this de facto standard. To avoid scrolling pages, Web designers need to consider the limitations of browser real estate even for users with a high screen resolution. If they make use of the full resolution of 1024×768, the resulting design restrains many users who might prefer to use their Web client in windowed mode or concurrently use tools that reduce their available screen space.

In conclusion, we see two challenges for the future: on the one hand, rapid interaction should be supported by user interface standards for Web pages and Web applications that meet the users’ expectations. On the other hand, browsers need to become more flexible and should be able to adapt to the type of Web site, the habits of the user and her tasks. Further research has to target these challenges.

REFERENCES

- [1] Abrams, D., Baecker, R., Chignell, M.: Information Archiving with Bookmarks: Personal Web Space Construction and Organization. Proc. CHI'98, 1998, pp. 41-48.
- [2] Barrett, R., Maglio, P.P., Kellem, D.C.: How to Personalize the Web. Proc. of CHI '97, Atlanta, GA., 1997
- [3] Catledge, L.D., Pitkow, J.E.: Characterizing browsing strategies in the World-Wide Web. Computer Networks and ISDN Systems, (27)6, 1995, pp. 1065-1073.
- [4] Chisholm, W. et al.: Web Content Accessibility Guidelines 1.0, World Wide Web Consortium, 1999.
<http://www.w3.org/TR/WAI-WEBCONTENT/>
- [5] Choo, C.W., Detlor, B., Turnbull, D.: Information Seeking on the Web: An Integrated Model of Browsing and Searching, First Monday Online Journal, (5)2, Feb. 2000
- [6] Cockburn, A., McKenzie, B., JasonSmith, M.: Pushing Back: Evaluating a New Behavior for the Back and Forward Buttons in Web Browsers. Int. J. of Human-Computer Studies. 57(5), 2002, pp. 397-414.
- [7] Cockburn, A., McKenzie, B.: What Do Web Users Do? An Empirical Analysis of Web Use. Int. Journal of Human-Computer Studies, 54(6), 2001, pp. 903-922.
- [8] Cole, J.I. et al.: The UCLA Internet Report - "Surveying the Digital Future". Los Angeles, UCLA Center for Communication Policy, 2003
www.forbes.com/fdc/mediasourcecenter/UCLA03.pdf
- [9] comScore Media Metrix: Surfing Down Memory Lane to January 1996: comScore Media Metrix Revisits First-Ever Web Site Rankings. February 25, 2004.
<http://www.comscore.com/press/release.asp?press=434>

- [10] Conklin, J.: Hypertext: An Introduction and Survey. IEEE Computer 20(9), 1987, pp. 17-41.
- [11] Cunha, C.R., Bestavros, A., Crovella, M.E.: Characteristics of WWW Client-based Traces, Tech. Rep. BU-CS-95-010, CS Dept, Boston University, 1995. <http://citeseer.ist.psu.edu/cunha95characteristics.html>
- [12] Greenberg, S., Cockburn, A.: Getting Back to Back: Alternate Behaviors for a Web Browser's Back Button. 5th Human Factors and the Web Conference, 1999.
- [13] Halasz, F. G.: Reflections on Notecards: Seven Issues for the Next Generation of Hypertext Systems. Communications of the ACM, 31(7), 1988, pp. 836-852.
- [14] Harris Interactive: More Than Four in Ten Internet Users Now Have Broadband. Poll #63, 8.9., 2004. www.harrisinteractive.com/harris_poll/index.asp?PID=492
- [15] Hoffman, M.: Enabling Extremely Rapid Navigation in Your Web or Document, 2nd Version, December 1997. <http://www.hypertextnavigation.com/infoaxcs.htm>
- [16] Horton, W.: Designing and Writing Online Documentation: Help Files to Hypertext. (NY) John Wiley & Sons, 1990.
- [17] Hyams, J., Sellen, A.: Gathering and Sharing Web-Based Information: Implications for "ePerson" concepts. HP Labs: Tech. Report: HPL-2003-19, Feb. 2003, 42 pages.
- [18] Jansen, B. J., Spink, A., Saracevic, T.: Real life, Real Users and Real Needs: A Study and Analysis of User Queries on the Web. Information Processing and Management, 36(2), 2000, pp. 207-227.
- [19] Jansen, B. J., Pooch, U. W.: A Review of Web Searching Studies and a Framework for Future Research. J. American Society of Information Science, 52(3), 2000, pp. 235-246.
- [20] Jones, W., Bruce, H., Dumais, S.: Keeping found things found on the Web. 10th Int. Conf. Information and Knowledge Management (ACM CIKM), 2001, pp. 119-134.
- [21] Kaasten, S., Greenberg, S.: Integrating Back, History and Bookmarks in Web Browsers. Ext. Abstracts CHI'01, 2001, pp. 379-380.
- [22] Kehoe, C., Pitkow, J., Sutton, K., Aggarwal, G., Rogers, J.: 10th GVU WWW User Survey, 1998. http://www.cc.gatech.edu/gvu/user_surveys/survey-1998-10
- [23] LinkQuality.com: The Web's Integrity Benchmark, Linkalarm Inc., 2005. <http://linkquality.com/>
- [24] Lynch, P. J., Horton S.: Web Style Guide: Basic Design Principles for Creating Web Sites. (Yale), Yale University Press, 2002. <http://www.webstyleguide.com/>
- [25] Milic-Frailing, N., Jones, R., Rodder, K., Smyth, G., Blackwell, A., Sommerer, R.: Smartback: Supporting Users in Back Navigation. 13th Int. Conf. on World Wide Web, NY, 2004, pp. 63-71.
- [26] Morkes, J., Nielsen, J.: Concise, SCANNABLE, and Objective: How to Write for the Web, Sun Microsystems. 1997. <http://www.useit.com/papers/webwriting/writing.html>
- [27] Moyle, M., Cockburn, A.: The Design and Evaluation of a Flick Gesture for 'Back' and 'Forward' in Web Browsers. 4th Australasian User Interface Conf., 2003, pp. 39-46.
- [28] Nielsen, J.: Scrolling and Scrollbars. Useit Alertbox, Nielsen Norman Group, Freemont, USA, August 2005. <http://www.useit.com/alertbox/20050711.html>
- [29] Nielsen, J.: The Need for Web Design Standards. Useit Alertbox, Nielsen Norman Group, Freemont, USA, Sept. 2004. <http://www.useit.com/alertbox/20040913.html>
- [30] Nielsen, J.: The Tyranny of the Page: Continued Lack of Decent Navigation Support in Version 4 Browsers. Useit Alertbox, Nielsen Norman Group, Freemont, USA, Nov. 1997. <http://www.useit.com/alertbox/9711a.html>
- [31] Nielsen, J.: Ten Best Intranets of 2005. Useit Alertbox, Nielsen Norman Group, Freemont, USA, February 2005. <http://www.useit.com/alertbox/20050228.html>
- [32] Obendorf, H., Weinreich, H., Hass, T.: Automatic Support for Web User Studies with SCONE and TEA. Ext. Abstracts of CHI '04, (Vienna), 2004, pp. 1135-1138. <http://www.scone.de/>
- [33] Ottaway, M.: Nielsen/Netratings Data Support TVNZ's Widescreen Web Decision. Nielsen/Netratings, May 2004. http://www.nielsen-netratings.com/pr/pr_040504_nz.pdf
- [34] Peterson, E.: Which screen resolution should you be using? Jupiterresearch Analyst Weblogs, 2004. <http://weblogs.jupiterresearch.com/analysts/peterson/archives/003965.html>
- [35] Pitkow, J.E.: Summary of WWW Characterizations. The Web Journal 2(1), 1998, pp. 3-13.
- [36] Rajamony, R., Elnozahy, M.: Measuring Client-Perceived Response Times on the WWW, 3rd USENIX Symposium on Internet Technologies and System, (SF), 2001.
- [37] Robertson, G., McCracken, D.L., Nevell, A.: The ZOG Approach to Man-Machine Communication. Carnegie Mellon Univ., CS Dept., Tech. Rep. CMU-CS-79-148, 1979
- [38] Sellen, A.J., Murphy, R. M., & Shaw, K.: How Knowledge Workers Use the Web. Proc. of CHI 2002, Minneapolis, ACM Press, 2002, pp. 227-234.
- [39] Shneiderman, B.: Designing the User Interface: Strategies for Effective Human-Computer Interaction. AW, 1992.
- [40] Spiliopoulou, M., Mobasher, B., Berendt, B., Nakagawa, M.: A Framework for the Evaluation of Session Reconstruction Heuristics in Web Usage Analysis. INFORMS Journal of Computing, 15(2), 2003, pp. 171-190.
- [41] Spool, J., et al.: Web Site Usability: A Designer's Guide. Morgan Kaufmann Publishers, (San Francisco, CA), 1998.
- [42] Tauscher, L., Greenberg, S.: How People Revisit Web Pages: Empirical Findings and Implications for the Design of History Systems. Int. J. of Human Computer Studies, 47(1), 1997, pp. 97-138.
- [43] The Counter: Global Screen Resolution Statistics. Jupiter-media Corporation, 2005. <http://www.thecounter.com/>
- [44] Weinreich, H., Lamersdorf, W.: Eine Umfrage zu Link- und Objekt-Attributen im Web. Proc. Mensch und Computer 2003. Teubner Verlag, 2003, pp. 387-390.
- [45] Weinreich, H., Obendorf, H., Herder, E.: Data Cleaning Methods for Client and Proxy Logs. WWW 2006 Workshop Proc.: Logging Traces of Web Activity, 2006