Not Quite the Average: An Empirical Study of Web Use

HARALD WEINREICH and HARTMUT OBENDORF
University of Hamburg
EELCO HERDER
University of Hannover
and
MATTHIAS MAYER
University of Hamburg

In the past decade, the World Wide Web has been subject to dramatic changes. Web sites have evolved from static information resources to dynamic and interactive applications that are used for a broad scope of activities on a daily basis. To examine the consequences of these changes on user behavior, we conducted a long-term client-side Web usage study with twenty-five participants. This report presents results of this study and compares the user behavior with previous long-term browser usage studies, which range in age from seven to thirteen years. Based on the empirical data and the interview results, various implications for the interface design of browsers and Web sites are discussed.

A major finding is the decreasing prominence of backtracking in Web navigation. This can largely be attributed to the increasing importance of dynamic, service-oriented Web sites. Users do not navigate on these sites searching for information, but rather interact with an online application to complete certain tasks. Furthermore, the usage of multiple windows and tabs has partly replaced back button usage, posing new challenges for user orientation and backtracking. We found that Web browsing is a rapid activity even for pages with substantial content, which calls for page designs that allow for cursory reading. Click maps provide additional information on how users interact with the Web on page level. Finally, substantial differences were observed between users, and characteristic usage patterns for different types of Web sites emphasize the need for more adaptive and customizable Web browsers.

Categories and Subject Descriptors: H.5.4 [**Information Interfaces and Presentation**]: Hypertext/Hypermedia—*User issues*; H.5.2 [**Information Interfaces and Presentation**]: User Interfaces—*Evaluation/methodology*

This article is an extended version of Weinreich, H., Obendorf, H., Herder, E., and Mayer, M. 2006. Off the beaten tracks: exploring three aspects of Web navigation. In *Proceedings of the 15th International Conference on the World Wide Web*, Edinburgh, UK, 133–142.
Author's address: Department of Informatics, University of Hamburg, Germany; email: harald@weinreichs.de.

General Terms: Human Factors, Design, Experimentation

Additional Key Words and Phrases: Navigation, browser interfaces, web browsing, web design, user study, usability, hypertext, WWW, Web

ACM Reference Format:

Weinreich, H., Obendorf, H., Herder, E., and Mayer, M. 2008. Not quite the average: An empirical study of Web use. ACM Trans. Web, 2, 1, Article 5 (February 2008), 31 pages. DOI = 10.1145/1326561.1326566 http://doi.acm.org/10.1145/1326561.1326566

1. INTRODUCTION

The World Wide Web has become one of our primary means of information and communication, a space for expressing both private and professional interests, and at the same time a huge marketplace and economic factor. Today, it is a key platform for news and entertainment, e-commerce, research, communication and collaboration (USC 2007). The importance of the usability of the related user interfaces can hardly be overestimated. And yet, an important factor is omitted when content, structure and experience are designed and evaluated on the Web: we know surprisingly little about the way people interact with their browsers during their daily use of the Web, or about ways in which they revisit pages after a longer period. While user navigation on single Web sites is commonly logged and used for subsequent analysis of user behavior, the exact nature of the users' interaction with the browser and cross-site browsing patterns remain inaccessible, as they can only be observed on client side. Studies analyzing personal use of the Web over a long term are surprisingly scarce: newer studies focus on specific tasks, were performed under laboratory conditions, or dealt with search rather than Web navigation in general—and consequently can only give a limited insight into everyday Web use (see Section 2).

The most recently reported client-side long-term studies are more than 7 years old—and thus represent the 1990s, a time in which the World Wide Web was still in its incipiency: the user population was dominated by researchers, most documents had static content and the focus lay on information delivery. With the increasing commercialization and the growing number of people accessing the Web using home connections, its user population became more and more diverse and new requirements emerged. Great efforts have been made to standardize technological infrastructure, and the inventiveness of those designing interactive experiences on the Web—within and beyond technological limitations—is astounding. Since the end of the nineties, new Web applications have gained popularity providing functionality which used to reside on desktops, covering a wide range of tasks from email, chat and bulletin boards to complex services such as travel agencies, libraries, and shops. This development was hardly predictable when the first Web browsers were developed. Yet, current browser interfaces and their navigation tools still closely resemble those of the browsers from the early Web days, mainly focusing on information retrieval and hypertext navigation. This mismatch illustrates the need for updating and extending findings on how users interact with the Web and what problems they encounter today.

This paper presents findings of a naturalistic long-term client-side Web usage study that provides empirical evidence for the problems of current users and identifies critical areas for a future improvement of the Web experience as a whole. We recorded a detailed stream of user actions with the Web browser directly in the working environment of our participants over a period of two to four months. This fairly comprehensive account of browsing activities was accompanied by interviews to obtain qualitative data about the tasks and habits of our participants and selected browsing sessions. We were able to identify changes compared to previous studies, analyze effects on different user interaction practices with Web browsers and pages, and identify new demands for browser navigation tools and the design of usable Web applications.

In the following sections, we first relate our study to previous work and recapture some results of the preceding long-term client-side studies that were conducted between 1994 and 1999 (Section 2). The study set-up is described in Section 3, including the selection of participants, interview procedures, our browser logging environment, and data processing and consolidation. The client log data and the interview results form the basis for our quantitative and qualitative analyses, which are presented in Section 4: We first discuss how user interaction with the Web (navigation actions) reflect the evolution of the Web to a hybrid between information system and online application (Section 4.1). This development poses new challenges to browser interface design and especially the history functions of current browsers (Section 4.2). Whereas in the previous studies Web browsing was mainly limited to one browser window, switching between multiple windows and browser tabs has become commonplace for a major share of our participants—a change that places higher demands on browser support for Web page revisitation (Section 4.3). We then substantiate the observations of other researchers that users often spend only very little time on Web pages, even on those rich in content and links (Section 4.4). Subsequently, our analysis of link click positions provides insights on within-page navigation habits and the interaction with search result pages. Although scrolling seems to be quite common even on navigation pages, nearly half of the clicked links were located in the upper left quadrant of the initially visible Web pages (Section 4.5). The recorded browser window sizes demonstrate that Web browsing is an activity that needs to share the users' attention—and screen space—with other applications (Section 4.6). Finally, we explore the differences in navigation habits of our participants. Personal preferences and the types of Web sites visited regularly have remarkable influence on interaction behavior. Consequently, statistical data on the "average" Web user can hardly represent the diversity of user needs (Section 4.7).

The article concludes with a discussion of the possible effects of our results on the development of future browsers and Web design guidelines.

2. RELATED STUDIES

How users browse the Web has been subject to research almost from its inception. Many studies are based on data from *server access logs*, analyzing various aspects of user navigation [Pitkow 1998; Spiliopoulou et al. 2003; Nicholas et al.

2006]. However, these logs have their limitations: they only report on user actions within a single site, browser caching mechanisms may hide some page revisits, and various specifics on the interaction with the browsing application remain concealed. Other researchers analyzed *search-engine transaction logs* [Jansen et al. 2000; Jansen and Pooch 2000; Jansen and Spink 2005; Rose and Levonson 2004; Schmidt-Maenz and Koch 2006; Teevan et al. 2006] to gain insights into query behavior and selected elements of the result list. The users' navigation paths after leaving the result page, however, remain unknown.

White and Drucker [2007] recently filled this gap by using a browser logging system and analyzing the navigation trails of over 2500 participants after having used a search engine. The results indicated that a certain type of users tended to "navigate" from a search result page, whereas others displayed a more "exploratory" behavior by submitting many queries during a search session, jumping between search results, and visiting numerous new sites. They suggested different strategies for coping with these different kinds of users.

Another perspective is provided by observational short-term studies. Screen captures, video coverage, and Web diaries have been successfully applied to understand Web browsing activities in various specific contexts [Baker 2003; Bernard et al. 2002; Choo et al. 2000; Hyams and Sellen 2003; Knight et al. 2007; Milic-Frailing 2004; Sellen et al. 2002]. Teevan et al. [2004] demonstrated the value of link following in addition to querying a search engine. This process of *orienteering* provides the user with contextual information for further specifying the search and interpreting the results. Orienteering has found to be important in *refinding* information as well; in a laboratory study, Capra and Pérez-Quiñones [2003] discovered that users heavily relied on recovering *waypoints*, which they used to remember the trail to the desired location.

Whereas observational studies provide us with qualitative insights on Web navigation, it remains unclear to what extent the results—which may be biased by the test-environment setting—can be observed in everyday Web usage. Quantitative evidence can only be provided by long-term client-side studies. Unfortunately, since 1994, only three long-term studies have been published that focus on the user's interaction with the Web browser. Of these studies, only two recorded exact navigation actions—and both are more than 10 years old by now.

In 1994, the first long-term client usage study was described by Catledge and Pitkow [1995]. They used an instrumented version of XMosaic on their departmental Sun workstations to record the browsing activity of 107 users for 21 days. Catledge and Pitkow identified several *navigation strategies*—for instance a "hub-and-spoke pattern" as a result of the frequent use of backtracking—and found that users operate only on small areas within Web sites. The back button ranked second in this study regarding frequency of use, accounting for 41% of navigation actions; only hyperlinks were used more often (52%). Other actions, such as "archiving" actions (e.g., saving or printing a page), were comparatively rare.

In 1995, Tauscher and Greenberg focused on history support and analyzed the *revisitation behavior* of their participants [Tauscher and Greenberg 1997]. 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." They also found two power law distributions within the revisitation patterns: most page revisits were to pages seen *recently* before, and furthermore a small number of highly *popular* pages received a high share of all revisits per user.

Finally, in 1999, Cockburn and McKenzie studied Web use by retroactively analyzing the Netscape history and bookmark files of 17 users [Cockburn and McKenzie 2001]. Their participants visited more pages per day than reported in previous studies, but at the same time the revisitation rate increased to 81%. Their data also indicated that Web use is rapidly interactive and that users often visit many pages within seconds successively. Finally, new insights on homepage and bookmark use were obtained: their participants tended to take many bookmarks, but used only few of them.

Many researchers and Web designers still use these studies as a point of reference. However, as indicated in the introduction, the landscape of the World Wide Web has changed dramatically. In Obendorf et al. [2007], we examined the effect of these changes on how users revisit information and updated existing findings to create a more detailed account of revisitation on the Web. In this article, we identify several aspects of Web navigation and investigate to what extent current browsers cope with the new requirements of Web users. Design implications are given where appropriate.

3. THE WEB-BROWSING STUDY

The Web-usage study presented here was conducted in Winter 2004/2005 with 25 unpaid volunteers. The subjects were acquired in the personal and professional environments of the authors of this study. To obtain use data from a diverse group, we invited all users with a Web experience of at least three years and a minimal use frequency of one Web session per week.

Six of our participants were female (24%). Ages ranged from 24 to 52 years (mean: 30.5). All subjects were experienced with the Web, having used it for 3 to 12 years (mean: 8). Most of them came from Germany and the Netherlands (three Germans lived abroad in Ireland and New Zealand) and all interviews were conducted in their native language. While all eight participants from the Netherlands worked as university employees in computer science, nine Germans (36%) had a different background: two worked in psychology, and one each in sociology, geology, electrical engineering, trading, coaching, history, and photography. Seven additional participants began the study, but dropped out due to technical or personal 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 137,272 user-initiated page visits to 65,643 distinct URIs and 9,741 different domains (see Section 3.3).

Interpretation of the quantitative results from the Web logs was facilitated by two interviews that each lasted at least 90 minutes and that were held at the beginning and the end of the study. During the first interview we gathered demographical data and information on general use and habits. The second interview focused on the user interaction behavior and tasks during the study to

validate our interpretation of the captured data. For example, we asked them about their usage of multiple browser windows and tabs, and their intentions to visit certain pages or sites frequently. Furthermore, we presented graphical visualizations of selected browsing sessions, to understand the goals and problems of our participants. The qualitative data from the interviews thus contributed significantly to the study and the results reported here. Although the interviews limited the number of participants, they provided deeper insights into user practices and preferences that would have remained hidden in a purely quantitative study.

Twenty five subjects cannot provide representative mean values for all users of the Web; however, the long-term character of the study and the qualitative results allow identifying problems relevant for many Web users. Furthermore, we were able to elaborate reasons for the strong personal differences between users. Even in our small sample usage intensity, navigation strategies, and visited sites varied considerably between individuals. In fact, the diversity of browsing habits of our participants was significant enough to indicate that researchers should be careful to draw general interpretations from pure statistical data of Web use (see Section 4.7).

In our data, the only consistently popular site was Google: it was the most frequently visited site for 11 participants and within the top four sites for all others. Therefore, Google use received special attention in several of our analyses.

3.1 Pre-Study Challenges

Considering the importance of the Web, it might seem surprising that only so few long-term studies used client-side log data to analyze the browsing behavior of Web users. This may partly be explained by the social and technical difficulties we had to overcome during the preparation of the study. Today, browsing is considered a *private activity*—or at least, nonprivate browsing is constantly interleaved with private browsing activities, even if logging focuses on the workplace: The Web is now used for many confidential tasks, such as online banking, shopping or Web-based email. After initial informal surveys, it became clear we had to create a capturing system that does not record user names or passwords, and that ignores communication over secure connections. Furthermore, all participants were given the opportunity to screen their own log files before making them available to us.

Several prospective participants were also concerned that the installation of "spyware"—technically, a fairly appropriate description of our instrument—might have negative impact on the reliability of their personal computer. These concerns were not completely unfounded: in one of our pilot studies instrumenting Internet Explorer to record user actions and page requests led to compatibility issues when different Explorer versions were used or new plugins were installed, decreasing overall stability—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 used for 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. Furthermore, the release of Firefox 1.0 in 2004 provided a platform for instrumenting a popular Web browser. Its interface is familiar to users of Internet Explorer and new security risks of the internet motivated several of our participants to "upgrade." The open source status of the browser made it possible for us to implement a logging mechanism for all user interface actions, supplementing the data that was recorded by the intermediary system.

3.2 The Browser Logging Environment

The browser logging environment consisted of two complementing extensions: every participant had an intermediary installed that filtered all transferred pages, and 15 of the 25 participants additionally made use of an instrumented version of the Firefox browser.

The intermediary added JavaScript code to every page. When executed, this code assigned unique identifiers to windows and frames; links and submit elements got additional event handlers, and many browser parameters were registered, such as the load and stay time, the history state and the document window size. It was a major challenge to keep this code compatible with interactive Web pages, especially to prevent interference with present JavaScript code of all kinds of pages. The event data of each event were transmitted to the intermediary by requesting (hidden) image objects from the intermediary. The parameters were encoded in the request URI and the intermediary returned a small dummy image. This technique allowed us to record detailed data on all page requests using browser timestamps, including the selected link anchors, form submission data and the current browser status. Furthermore, the intermediary analyzed all transferred documents and recorded descriptive data about their contents, size and links. Our logging system was based on the Scone framework [Obendorf et al. 2004] and IBM's WBI [Barrett 1997].

The 15 participants who used Firefox during the study were either already using it as their preferred browser, or embraced the opportunity to switch. The instrumented browser recorded the users' interaction with all important user interface widgets. It was based on Firefox 1.0. We modified the source code of the user interface files written in ECMAScript [ECMA-262 1999] and XUL [Hyatt 2001]. Our browser install file already included some popular extensions, for example, All-in-One Gestures[1] for mouse gestures, so we were able to alter the code of these extensions as well. All relevant Interface elements called a function that wrote usage data into a protocol file in CSV format, including browser timestamps, window identifiers and event details. We were able to distinguish 76 different user actions with the browser.

Using browser timestamps and window identifiers, the Firefox log was merged with the enhanced click-stream log of the Scone intermediary to gain more detailed and accurate data.

[1]See: http://perso.orange.fr/marc.boullet/ext/extensions-en.html

3.3 Data Processing and Consolidation

The first analyses of the recorded log files showed that serious data preprocessing was necessary to get accurate results, since many of the entries were not directly related to user actions [cf. Weinreich et al. 2006]. We found several causes 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 extra events are created as every subframe causes a page request for a distinct HTML file. We identified frames by the name of the frame and the parent window, which was read by the embedded JavaScript code. The interpretation of these data required not only collating multiple page requests to one user action, the events were also often ambiguous, as we could not define which URI was the most important one for a user action as a universally valid rule. For example, 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 response affected the right page with another address [Weinreich et al. 2006]. The interpretation has consequences on load times, revisitation rate, recorded document sizes as well as the link positions. Unfortunately, no standard exists how to handle such frame events in log files, and no solution to this problem has been offered by previous studies. Therefore, we decided to exclude frame page requests for several of our analyses to prevent ambiguous results.

Another significant problem was caused by *advertisements*.[2] JavaScript-initiated advertisements in pop-up windows are not deliberate user actions, and consequently events relating to ad windows and page requests were excluded from the study. A statistically even more relevant advertisement technique is based on iFrames, which allow embedding other HTML pages in a Web document. According to our data, iFrames were mainly used to dynamically include advertisements. For the group of participants that did not apply any kind of ad-blocker (8 users), frame, iFrame, and advertisement artifacts represented about 28% of all HTML page requests. This is remarkable, as it does not even consider online promotion realized as plain text, embedded images or flash animations.

A third source of non-user-initiated page requests were *automatic page reloads*, mainly caused by news sites. In some cases, embedded JavaScript code refreshed a page contents after a certain interval, in other cases external applications like instant messaging agents were responsible for these artifacts. Such events became visible as peaks in the stay time distribution of some users. They contributed nearly four percent of all 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

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

known on the Web, and Bruce McKenzie [Cockburn and McKenzie 2001] let us know that in their study from 2000, the effect of such requests could still be neglected.

4. RESULTS

As discussed in the introduction, the Web has become increasingly diverse since the first long-term client-side studies were conducted in the mid-nineties: whereas ten years ago Web usage was focused on academic information exchange, nowadays the most popular sites are highly dynamic and cover multiple areas, such as ecommerce, entertainment and communication [USC 2007]. The role of the Web browser has moved from a hypertext viewer to a universal client for online services. Furthermore, several new browser features were introduced, like tabbed browsing and the back button's popup menu. But also personal workstations have evolved, Web access has become quicker and screen sizes have increased. However, the consequences of these changes on everyday Web use are only partially known and have not been backed by naturalistic long-term studies.

In this section, we will analyze key impacts of these manifold changes on the way users interact with the Web. Our log data of over 135,000 user-initiated page visits revealed significant shifts in user interaction since the last studies, and the interviews supplied qualitative data to identify reasons for these changes. Combined, they provide a background for developing future browsers, browser enhancement, and designing more usable Web pages and applications.

First, we address changes in Web navigation that can be induced from the recorded user interaction with browser interface widgets (Section 4.1). The most remarkable difference to previous studies is the declining use of backtracking that motivated further analyses. We continue with a discussion on challenges for site design and browser navigation support for increasingly dynamic Web sites (Section 4.2). In the third subsection, the merits and drawbacks of the application of multiple browser windows and tabs are analyzed. We show that the established concept of the back button does not correspond with the way many users apply windows and tabs (Section 4.3).

The fourth subsection focuses on the speed of user interaction with the Web and consequential design requirements for Web pages: even pages rich in content and links are frequently visited only for a very brief time (Section 4.4). Thereafter, user navigation *within* Web pages is explored—scrolling activities and link selection (Section 4.5). Most selected links are within the initially visible screen and users focus on an area in the upper left corner of the browser; this calls for a page layout consistent with established standards, and indicates the usefulness of fitting all options on a single screen. The next section shows, however, that user habits limit the available document space even on screens with high resolution (Section 4.6).

Finally we analyze individual differences between Web users, as a result of various user tasks, navigation strategies and site interests (Section 4.7). The significant differences between our participants demonstrate the required versatility of current browsers.

Table I. Comparing Three Long-Term Client-Side Web Studies

	Catledge & Pitkow[3]	Tauscher & Greenberg[4]	This Study
Time of study	1994	1995–1996	2004–2005
No. of users	107	23	25
Length (days)	21	35–42	52–195, ø=105
No. of visits	31,134	84,841	137,272
Recurrence rate	61%	58%	45.6%
Link	45.7%	43.4%	43.5% →
Direct access	12.6%	13.2%	9.4% ↘
New window	0.2%	0.8%	10.5% ↗
Submit	–	4.4%	15.3% ↗
Back	35.7%	31.7%	14.3% ↘
Reload	4.3%	3.3%	1.7%
Forward	1.5%	0.8%	0.6%
Other	–	2.3%	4.8%

4.1 The New Character of Web Navigation

Users interact with the Web browser to access resources on the Web. We labeled events that initiate a page visit *navigation actions*: they comprise all user actions that lead to a new entry in the browser history and therefore (usually) allow for returning to that browser state by using the back button or by bookmarking the page (see also Section 4.2). This definition of *navigation actions* is based on the current navigation model of Web browsers and the concepts used in related studies, enabling us to compare the results. It can be associated to browser *hard states* in the taxonomy of browser interaction states of Meschkat and Mittleman [2007], as opposed to *soft states* and *transient states* that are not represented by separate entries in the browser history and relate to within-page navigation like page scrolling, data input, or the interaction with Flash applications and AJAX-enabled pages.

Apart from selecting links, users can trigger *navigation actions* in different ways: entering URIs directly into the address bar of the browser, using[3] different browser history mechanisms to revisit pages, or submitting information via forms to interactive Web services, such as search engines.

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

Link following continues to be the most common navigation action, accounting for about 45% of all page transitions. 'Direct access' to pages—via the bookmark menu, bookmark toolbar,[4] home page button, or the address bar—has remained stable at about 10% as well. The most significant changes are the increased number of pages opened in *new browser windows*, the raised importance of *form submissions*, and a decrease in *back button usage*.

[3]Since previous studies recorded different sets of browser events, we recalculated some of the values of [Catledge and Pitkow 1995] and [Tauscher and Greenberg 1997], and used the notion of navigation action defined in this study.
[4]The bookmark toolbar (Figure 1) was not available in previous studies.

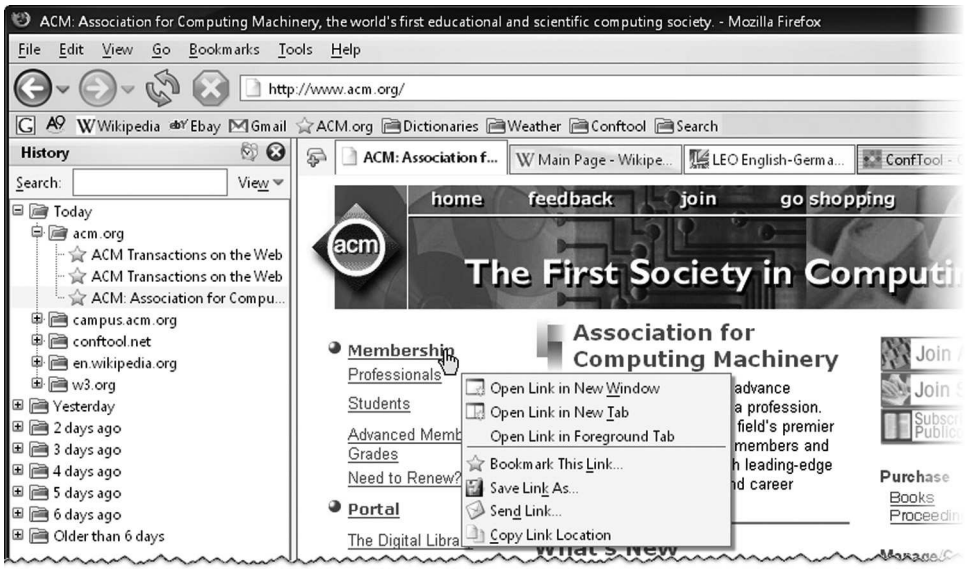

Fig. 1. Screenshot showing bookmark toolbar (top), history in sidebar (left), and browser tabs (right).

Several of our participants regularly opened new windows or tabs to display a new Web page. In the mid-nineties, such events 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 pull-down menu item was registered,[5] this study also considers previously unavailable user actions that result in opening a Web page in a new window. This includes following hyperlinks with target="_blank" as anchor attribute, starting the browser from the desktop or from another application, and using the "open link in new window" or "open link in new tab" entries of the browser's context menu (Figure 1). Nevertheless, the log data confirmed that it has become common behavior to have more than one document opened while browsing the Web (see Section 4.3)

Accounting for over 15% of all navigation actions, *form submission* has become a key feature of user navigation. By contrast, the share of back button actions has dropped from over 30% in the mid-nineties to less than 15% in our study. This number includes backtracking multiple steps via the back button's pull-down menu, which contributed only 3% to all backtracking actions and has therefore negligible influence on the decreased rate.

Another browser revisitation tool—the browser history (Figure 1)—is not explicitly listed in the comparison chart, as it was hardly used: merely 0.2% of all page requests were initiated from the history. Only two of our twenty-five participants reported to use it from time to time, while ten participants even were not aware of this tool at all.

[5]This follows from the much higher number of "close window" and "exit program" than "new window" events reported in Catledge and Pitkow [1995].

Reasons for the Reduced Backtracking Share. The reduced usage of the back button, in combination with an increase of "forward navigation actions"—following links, submitting forms and opening new windows—suggest that users return less frequently to previously visited pages. However, as the moderate decrease of the average *recurrence rate* [Tauscher and Greenberg 1997] indicates, the share of pages revisits has decreased to a much lesser extent; it dropped only from about 60% to 46%[6] (Table I). One explanation is that most Web sites nowadays offer structural links on every page that allow returning to the home page or landmark pages of the site, without using the back button. However, the reduced backtracking rate also relates to the increased amount of *submit* and *new window* actions.

The increased number of form submissions characterizes a fundamental change of the Web during the last decade: the move from an academic information system with primarily static hypertext documents to a hybrid between a common information source and service-oriented interactive systems, such as search engines, online shops, travel planners and Web-based email. Most service-oriented sites are more comparable to desktop applications than to information-centered hypertext: whereas hypertext navigation involves orienteering behavior with frequent backtracking, interactive applications are mainly used for completing certain workflows. Hence, backtracking should be 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 other participants (16.2%), a difference that is marginally significant ($t = 2.715$, $p = 0.012$). A closer look at the consequences of this change will be taken at the following Section 4.2.

We also expected to find a connection between the frequency of back-button usage and the increase of new window events, as some participants reported the habit of opening link targets in new windows or tabs to quickly return to an important page by selecting the corresponding tab or window instead of backtracking to it. Our data supports this statement: 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%), indicating that multiple windows are used as an alternative to backtracking ($t = 2.509$, $p = 0.026$). In addition to multiple windows, modern browsers provide "tabbed browsing"—several pages can be opened simultaneously in different browser tabs of one window (Figure 1). Our results suggest that browser tabs stimulate working with multiple Web sites simultaneously: six participants, who used tabs frequently, were backtracking less often (9.9%) than the remaining nine Firefox users (18.3%) that opened hardly any tabs ($t = 2.311$, $p = 0.038$). Merits and drawbacks of the usage of multiple browser windows and tabs are discussed in more detail in Section 4.3.

Another reason for a decreased backtracking rate might have been the emerging of AJAX-enabled pages that allow for navigation within a page without the

[6]Frame pages were excluded for consistency reasons [Weinreich et al. 2006]. We used an updated definition for the calculation of the recurrence rate [Obendorf et al. 2007].

need—or the option—to return to the last browser state by using the back button. Although the required browser technologies had been available since Internet Explorer 5 from 1999, our analysis of the 300 most frequently visited sites of this study showed that this technique still had no significant influence. However, we suppose that the backtracking share will further decrease with the rising use of AJAX for within-page navigation.

4.2 Navigation Support for the Dynamic Web

The interaction with HTML forms stands in direct relation to the use of interactive Web applications, since every form submission has to be processed by the server and results in a dynamically created Web page.

Overall, the most popular Web application of our participants was the Google search engine; over 15% of all page requests were sent to one of the national or international Google search sites, and 43% of all form submissions related to Google search.[7] Further important interactive sites were eBay, an online dictionary, and several ecommerce sites, for example, online shops and travel agencies.

The increased share of form submission events is not the only indicator for the prominence of dynamic Web sites. *Links* played a major role in interactive Web services, too: 40.7% of all selected links leading to a page transition had parameters encoded in the query component of the target's URI.[8] Consequently, 44.1% of all page requests used GET or POST parameters, suggesting that the resulting Web page included dynamically generated content. We expect that the average rate of dynamic Web pages was even higher, as many popular content management systems with frequently updated pages embed parameters in path or filename of the URI to be indexed by search engines. Furthermore, our logging system did not record "cookie" data, even though personalized Web content often relies on the use of cookies—for instance the product suggestions pages of the Amazon online store.

To get a more precise picture of the dynamic nature of the current Web, the logging system took fingerprints of all requested Web pages. This fingerprint was calculated as hash code of the page contents, so even minor document alterations were considered. For short term revisits within one hour, the content of 26.2% of all documents did change,[9] a rate much lower than the number of page requests involving parameters (44.1%). However, for revisits after one day or later, 69.0% of all revisited pages did experience a change, a rate that stays nearly constant even for longer revisitation periods (Figure 2). These results are considerably higher than those of studies based on random samples of Web pages: Cho and Garcia-Molina [2000] reported an average change rate for Web

[7]Other global search services like Yahoo or MSN were hardly used at all (below 1% of all form submissions).

[8]We did not consider link clicks that triggered JavaScript programs, if these scripts did not lead to a 'navigation action', i.e., the loading of a new Web page, as such events do not create new entries in the browser's history stack. Accordingly, updating only parts of a page by AJAX techniques was excluded as well as in-page interaction with JavaScript applications.

[9]The browser cache was disabled during the study for html documents by the intermediary system, so all Web pages were reloaded on every revisit.

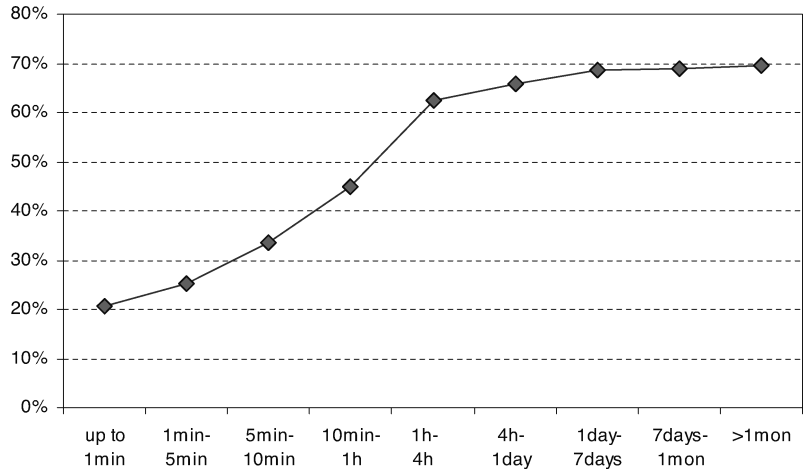

Fig. 2. Proportion of changed pages for different revisitation periods.

pages of 23% per day and 40% per week, and Fetterly et al. [2003] documented a share of 35% of updated pages after one week. However, we do not believe that the proportion of dynamic Web pages had nearly doubled from 2003 to our study in 2004–2005; it seems more reasonable that our participants accessed dynamic pages to a larger extent than static pages. This assumption is supported by the character of the most popular sites of our participants: Most of them provide online applications or frequently updated content.

Dynamic, interactive pages pose new challenges to the browsers' history mechanisms. Users often do not return to the page visited before, but to a similar or even completely different page. Furthermore, if the server forces a page update with every visit, browsers usually cannot restore their old *soft state* [Meschkat and Mittleman 2007], that is, return to the last scrolling position or show recently filled-in form values. Though it might be desirable to get frequently updated information, users might as well require some data accessed before, which happens not to be available anymore—for example a certain Google search result[10] or a deleted article. A solution could be complementary history and backtracking functions that record page contents as well and allows reaccessing any information displayed before.

Pages based on form data that was transmitted using the POST command cause even more severe problems for short-term revisits. *Backtracking* is impeded, as the data has to be submitted again, and the user accordingly gets an—often confusing—warning message. Moreover, many interactive Web sites that use POST submissions regularly do not support backtracking at all: some services try to disable the back button by opening a pop-up window without navigation toolbars for sequences of interactive forms, or they explicitly advise users not to use backtracking. If users still try to backtrack, for

[10]Google result pages also tend to change frequently. According to our data over 98% of the result pages were updated after seven or more days. For revisits on the same day less than 12% of the result pages did change, probably mainly caused by new advertisements.

example by mouse gesture or by using keyboard shortcuts, they receive an error message.

In the context of Web applications, navigation tools like the back button acquire a new meaning: if a user presses it to correct errors or to provide alternative input, it bears more similarity to an *undo button*. While undo functionality is a must in office applications, and generally considered a key factor for controllability [ISO 9241/110 2006], it is still rarely found in Web applications.

The *long-term* history tools of current browsers even completely fail with many online applications, as these tools do not store any POST data or save the login procedure to access a desired document—an issue that emerged in many interactive sites in our study: excluding Google search, about 47% of all form submissions used the POST method. Consequently, neither history nor bookmarks provide the means to revisit these documents and the user is not even notified about these problems; once the browser window is closed, the page is lost. If users want to store such a volatile document for future reference—for instance, a travel plan or an order confirmation—they have to print or save the page. However, the document is then no longer accessible through the browser interface any more—one reason why our participants hardly used archival commands like printing and saving.

First steps to solve this problem have been taken: the (rather outdated) Internet Explorer 5 for Mac OS X featured the Scrapbook, an integrated interface for storing an exact copy of the Web page as it appears in the browser window. An extension with the same name is available for Firefox, which also works with AJAX-enabled pages and allows for annotating the stored documents.[11] With the growing importance of service-oriented sites, similar functionality might be required by most Web users and should be a basic functionality of contemporary browsers. Travel plans, flight reservations, invoices, and bank statements 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.

In conclusion, common browser interfaces lack several functions for service-oriented sites, although these sites play an increasingly dominant role in contemporary Web use. 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.

4.3 Merits and Drawbacks of Multiple Browser Windows and Tabs

As mentioned before, several of our participants made regular use of multiple windows and tabs when browsing the Web. The correlation between the frequent opening of new windows and a low back button rate indicated that opening link targets in a new browser area was a strategy to circumvent the need for backtracking.

To gain a deeper understanding of the individual navigation strategies, we analyzed the number of concurrently used browser windows and tabs. Our participants had on average 2.1 windows or tabs opened when they accessed a

[11]The Firefox Scrapbook extension can be found at: http://amb.vis.ne.jp/mozilla/scrapbook/

new page,[12] suggesting that the use of multiple windows is not an exception, but the rule. However, the individual average differed from 1.07 to 8.19 concurrently opened documents, which shows that this practice was not followed by all users. About one third used mainly one window while the remaining participants used windows and tabs to a different extent. The number of people in our study using multiple windows was higher than the number of those using tabs—only six of the fifteen Firefox users (40%) regularly opened browser tabs.

During the interviews, our participants reported several advantages of opening link targets in a new browser area. First, it allowed them to keep the source page opened, in order to explore more relevant hyperlinks—especially if they were not certain that the first selected link would yield a satisfying result. This was commonly used on search result and overview pages. By keeping several pages opened in different windows, users were able to "compare them side by side." A further reason was the impression that they would save time, as "pages could be loaded in the background" while they were free to continue other navigation activities. Keeping search results and resulting navigation trails in separate windows also reduced the risk of losing the path back to a decisive page.

The tab and window management strategies ranged from the incidental to the carefully planned; one participant explained that he used "new tabs for closely related tasks and new windows for parallel tasks".

This navigation strategy provides new challenges to the often criticized [Cockburn et al. 2002; Greenberg and Cockburn 1999; Kaasten and Greenberg 2001], yet frequently used back button. Opening a link target in new windows or tabs disrupts the principal function of the back button that should always allow for returning to recently visited pages. If users split their navigation trails into multiple browser windows, the recent visit history is separated into several stacks, with no temporal relation, and each individual stack does not include any information from the originating window. Hence, users need to remember what actions they performed in which window in order to regain certain information. This places a further cognitive burden on the user, in addition to the already demanding task of keeping track of their location in the Web [Conklin 1987].

A temporally ordered list of all recently visited pages—such as provided by the hardly used browser history—would do more justice to the sequential character of the parallel trails related to one task. Similar to Tauscher and Greenberg [1997], we found that, on average, the list of 15 most recently visited pages covered about 88% of all revisits, whereas pages listed in the popup menu of the browser's back button covered in mean only 52% of all revisits. However, such a temporally ordered list could be confusing, if trails in different windows belonged to *independent tasks* that need to remain separated. We see an alternative solution in a *branching history* that shows trails in temporal order, but separates activities carried out in different windows [Herder 2006].

[12]Framesets were considered as one window.

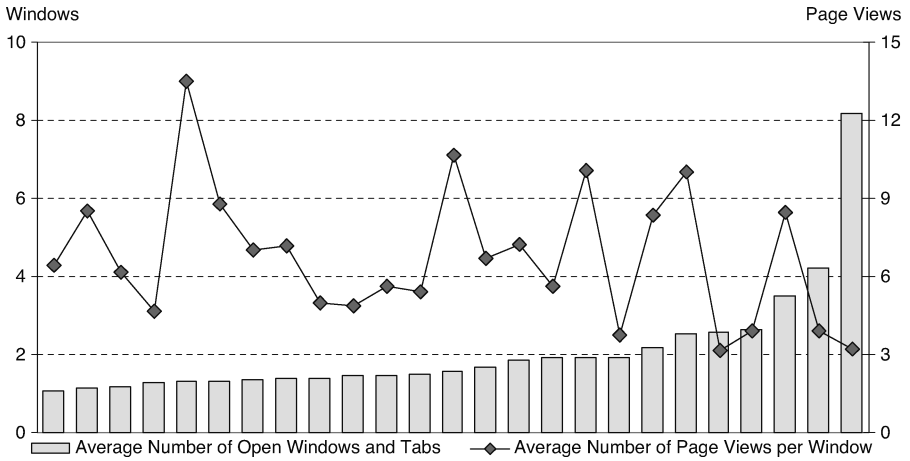

Fig. 3. Average individual number of open windows vs. path lengths.

Browser Windows as News Feeds, Task Reminders, and Appliances. Since many of our participants used windows and tabs as an alternative to backtracking, we expected that users who commonly opened only a single window would exhibit longer navigation trails in this browser window, while opening link targets in new windows would lead to many, but shorter trails. However, our data did not confirm such a correlation (Figure 3). Instead, we identified another habit that strongly influenced the mean number of concurrently opened windows: nearly half of our participants used to keep one or several windows opened in the background of their computer desktop for extended periods.

In the interviews, our participants provided different explanations for this habit. Several users wanted to have immediate access to a certain page they *monitored* for *frequently updated information*, such as a news site or a bulletin board [Kellar et al. 2007]. Although software like RSS readers could have replaced the browser for this purpose, none of our participants made regular use of the according tools.

Some users explained that occasionally these windows served as reminders for *unfinished tasks*, allowing them to continue some work in progress at a later point. The already mentioned missing history functions for dynamic pages obliged them to use such windows as *temporal bookmarks*.

A third reason for maintaining open browser windows in the background was the frequent utilization of *Web-based tools* in combination with some other desktop application. In particular some German participants often accessed an English-German online dictionary while working with their word processor. Support for an easier access to many Web applications has recently been complemented by special appliances, like Apple's Dashboard or Yahoo! Widgets. A flexible and direct integration in common office applications, for example, of online dictionaries in word processors, is, however, still not commonly supported.

Orientation Support for Multiple Windows. Handling multiple windows in information systems was already reported to cause disorientation in pre-Web

studies [Halasz 1988]. Our participants mentioned similar problems: several said that they found many open Web documents hard to manage, in particular because the page titles displayed in task bar and tabs were often not helpful. Firefox extensions like PageStyle2Tab[13] and Tab Catalog[14] already address this problem for tabs, by using page style and colors for the tab bar or showing a thumbnail overview of all currently opened documents. Still, the combination of multiple windows and tabs lacks decent support, and further research is required to minimize such problems.

4.4 The Speed of Web Navigation

An evaluation of the speed of user interaction with the Web promises to yield a better understanding of the requirements for Web page and browser design. We used the time our participants stayed on a Web page as an estimate for the time they took to read the page and think about available options before deciding on their next action. Although the time between page requests can be gained from server logs [Pitkow 1998] and the browser history [Cockburn and McKenzie 2001], the data of our client-side logging software was more precise. We 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 also distinguished 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. However, we could not identify if users actively used a Web page, which implies that the attention times per page were definitely shorter than the stay times we recorded.

Our data confirms the rapid interaction behavior with heavy tailed distributions already reported in previous studies [Catledge and Pitkow 1995; Cockburn and McKenzie 2001; Cunha et al. 1995]: 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 of the average stay times is located between 2 and 3 seconds; these stay times contribute 8.6% of all visits.

We first assumed that most of these short stay times represented *revisits*, for example, visits to pages that had been seen recently, or that were used frequently and therefore well known. To analyze how much time users take to read new Web documents, all revisited pages were excluded from the statistics and only first time page visits were considered. The effect on the distribution of

[13]PageStyle2Tab for Firefox can be found at: https://addons.mozilla.org/firefox/addon/1523
[14]The Tab Catalog extension can be found at: https://addons.mozilla.org/firefox/addon/1937

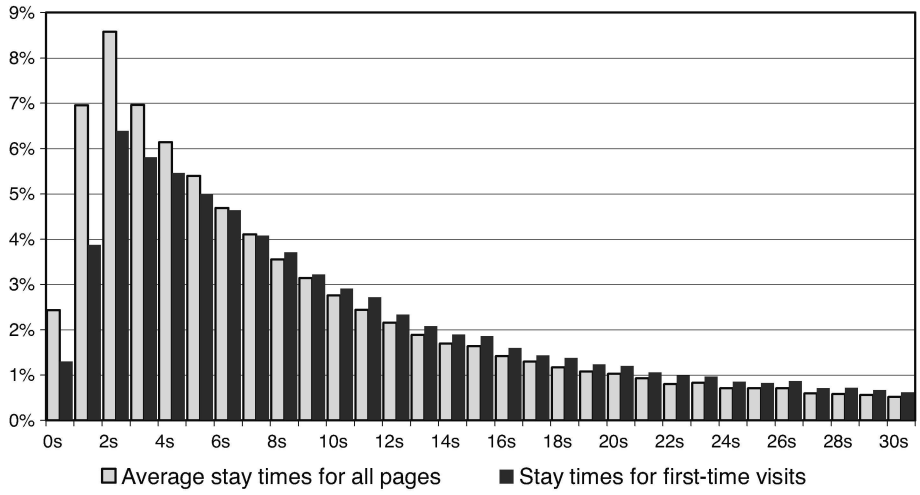

Fig. 4. Distribution of stay times for all participants.

stay times was not as strong as expected (Figure 4, blue bars): more than 17% of all new pages were still visited for less than 4 seconds, nearly 50% were shown for less than 12 seconds and 11.6% were 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.

Although Google search result pages usually show only ten items, the stay time distribution of these pages was similar to the average distribution. Again, most frequently, stay times were within 2 and 3 seconds (10%), but long stay times were more uncommon than on other pages: the median stay time was 8.0s (as opposed to 12.4s), and more than two-thirds of all pages were displayed for less than 12 seconds.

The tendency for very short page visits on the Web might have two reasons: either it expresses a cursory 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 (measured using only displayed text, not considering any markup code or any embedded objects or graphics) was 551 words ($\sigma = 811$).[15] The page stay times were dependent on the page size, but less than expected: pages visited for less than 12 seconds (which 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 5 (upper graph) illustrates the average number of words per

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

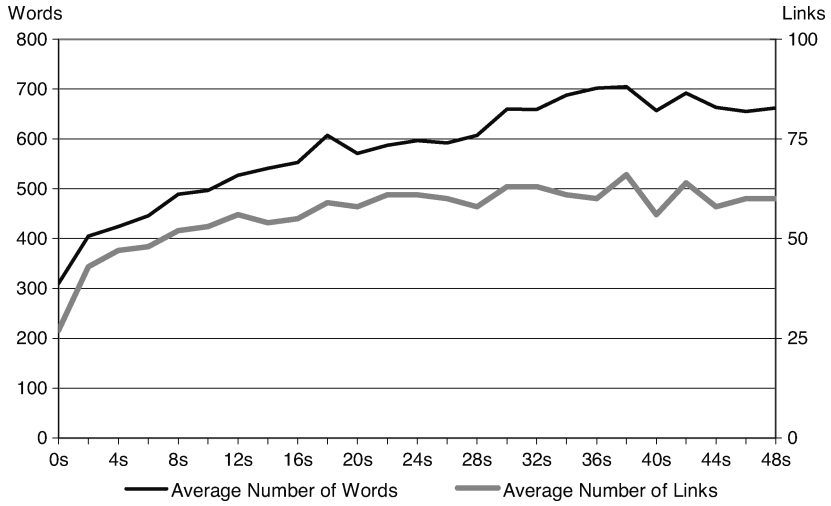

Fig. 5. Correlation between stay time, page contents and link count.

page grouped in intervals of 2s stay time. A similar difference was found for the number of navigational options per page (Figure 5, lower graph). On average, all visited pages had 53 hyperlinks[16] ($\sigma = 58$). For pages with a stay time of less than 12s, the average number of links was 46, and thus 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 the 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 users was already reported by Morkes and Nielsen [1997] and Spool et al. [1998] who observed it at several controlled Web usability studies. Although we also found differences between the habits of different users[17] 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 also next section). We expect an increase of the average page stay times with the emergent use of AJAX-enriched Web applications—as they allow system interaction without leaving a Web page—however, such techniques will hardly decrease the overall interaction speed with the browser. AJAX can even facilitate the fast interaction habits of many Web users by reducing system response times.

4.5 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 a usability challenge, as it can result in disorientation: the reader may lose track of the context as

[16]Again, outliers were removed. All visited pages had an average of 61 hyperlinks ($\sigma = 122$).
[17]While all participants exhibited a long-tailed distribution for their stay times with a maximum below 5s, five participants stayed significantly longer on many pages and had a median stay time of over 12s (up to 26s).

Table II. Location of Selected Links

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

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 out of the window, short pages allow comparing all available options at one glance. Therefore, guidelines recommend especially for *entry* and *navigation pages* that the whole document should fit on one screen and show all options immediately [Lynch and Horton 2002; Nielsen 1997].

On the other hand, for *content pages* longer documents bear advantages, since they can be read and printed without the need to flip pages [Baker 2003]. Still, it is generally advised to avoid wide pages that require horizontal scrolling for any kind of Web document, since users may need to move their viewport in two dimensions, wide text lines have a decreased readability, and the printout of wide pages may be cropped [Lynch and Horton 2002].

The short stay times measured for Web pages in our study seem to support the guidelines calling for short, nonscrolling pages; they also pose the question whether users take the time to scroll at all. As problems with scrolling are particularly considered critical for navigation pages, our analysis focused on pages that were used to navigate to another page. We examined 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 showed that most links (76.5%) were selected within the region visible on load time (Table II). Although our participants hardly scrolled horizontally to select a link (altogether 0.4%), over 23% of 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. Considering the short median stay time on Web pages, the number of followed links that appeared 'below the fold' is higher than we expected.

Part of the explanation for this finding is that the position of selected links is influenced not only 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, we created maps of link clicks by grouping the page space in sectors of 40 by 40 pixels, counting clicks within these sectors. The number of clicks was categorized to identify areas of high navigation activity.

We analyzed Google result pages separately from and all other pages, as Google search was by far the most common activity of our participants (Google's search result pages made up for 21% of the 27,000 recorded link clicks), users selected links outside the initially visible area with 38% significantly more often than on average (23%), and these pages have a layout that is clearly different from most other Web sites, that provide many graphical elements and navigation areas with structural links [Haas and Grams 1998].

The resulting click map of all non-Google pages is shown in (Figure 6, left diagram). The displayed area of 1040 $\times$ 1600 pixels covers 92% of the recorded

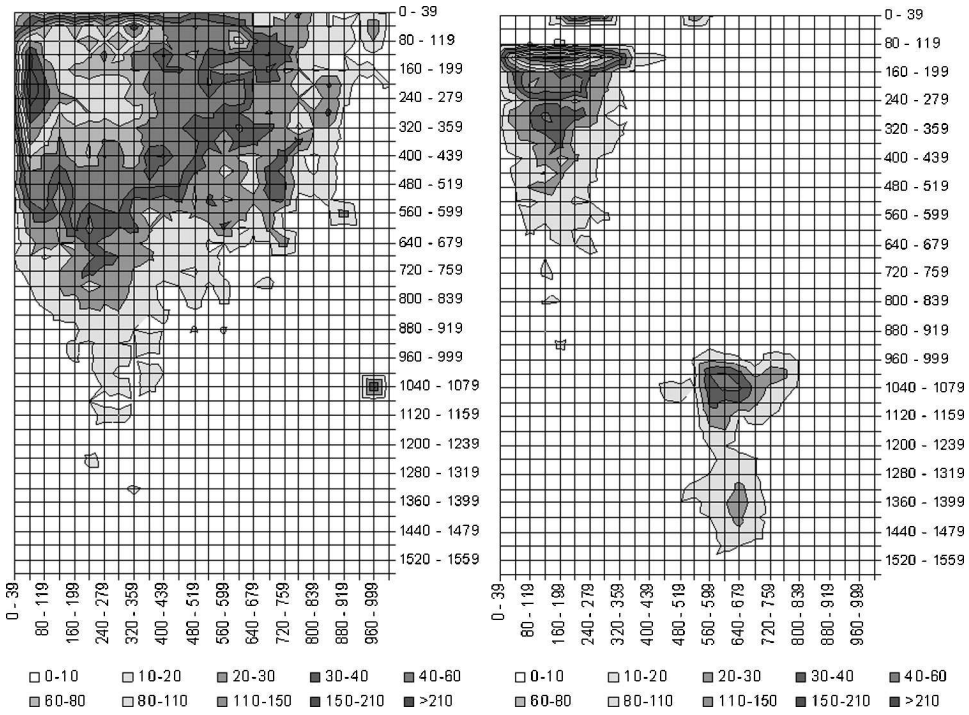

Fig. 6. Link activity areas of Web users on 'normal' (left) and on Google search result pages (right).

link clicks, the remaining clicks were done mostly below this area. The most actively selected region is located in the upper left corner of the map: about 45% of all user clicks occurred in a part of the browser window that is slightly larger than the upper left quadrant of the initially visible page area (520×400 pixels). Furthermore, the vertical navigation bars that provide the main structural hyperlinks on many sites become visible as well as several horizontal 'hot spots' in the upper part of the page.

The second click map created only from Google search result pages is displayed in Figure 6 (right diagram). Nearly all recorded link clicks (98%) are within the displayed map. It shows a quite different distribution where two active areas stick out: the upper left region stands for the first search results that were selected with an outstanding frequency, the following results drop rapidly in popularity. The second active area in the lower part of the documents—at about 600×1000 pixel and below—is mainly caused by the 'next page' link. The placement of this link required scrolling for all our participants; even without taking transfer time into account, it seems quite inefficient to require users to both scroll and flip through pages [Bernard et al. 2002].

The results also demonstrate the crucial importance of a high page rank in Google for user attention. Results below the first half of the page—usually outside the immediately visible area—got even less clicks than the "next page" button. Although this might also be an indication for the high quality of Google ranking mechanisms for many queries, the high share of page flips

shows how often users did not find the desired document among the first results.

The click maps show an apparent similarity to 'heat maps' gained from eye-tracking studies, depicting the regions of user attention on the screen, as reported by Outing and Ruel [2004], Nielsen [2006a], and Shrestha and Lenz [2007]. This likeness indicates a correlation between user attention and the selected links—not completely surprising if the short stay times of users are considered. Although this does not necessarily mean that these areas are generally suited best to place links, pages with a matching layout will meet the expectations of many users and increase the consistency of the Web.

It should be mentioned that these click maps depend on the user population. For instance, all of our participants were using the Latin alphabet, reading from left to right and top to bottom. Arabic or Hebrew Web pages—written from right to left—would most likely result in vertically mirrored diagrams.

4.6 The Actual Browser Real Estate

Web page optimization for the users' rapid navigation behavior includes incorporating scannability and providing all relevant information at a single glimpse. However, creating Web pages that do not require scrolling depend on knowledge of the available screen real estate. Although Web content should in principle be accessible for everyone regardless of different abilities and available hardware [Chisholm et al. 1999], for an aesthetic appearance device-specific designs are often inevitable. In practice, style sheets and graphical elements (e.g., bitmap graphics) are based on assumptions on the available screen resolution and browser space. 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 [TheCounter 2007]. Until 2004, many experts recommended basing 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 recommended more often.[18] We wanted to find out whether the full resolution is really available for Web pages, or whether the browser's viewport is limited for technical or personal reasons. 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 excluded, as their sub-pages would bias the statistics. All our participants used a screen resolution of at least 1024×768 pixels, twelve had an even 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 1) took some of the vertical and horizontal space

[18]Experts who support 'wide screen designs' are [Ottaway 2004; Peterson 2005] and for intranets [Nielsen 2005].

for most of them. Furthermore, some participants preferred to maximize the browser manually, leaving a border of several pixels around the window unused.

Our nine users with smaller browser windows (two of them with a screen width of 1024 pixels) left on the 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 users with a screen resolution of 1024 pixels, the average available document width was only about 890 pixels. They preferred the windowed mode for several reasons: it permitted them to select other windows directly on the screen making the organization of several applications more comfortable. Furthermore, a narrower window would improve the readability of many documents, as the line length is 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.

4.7 The Impact of User Habits and Visited Sites

The previous sections already addressed several individual differences of our participants. Although the comparison of the average shares of navigation actions with previous studies did expose severe changes in Web use, we want to emphasize the risk of drawing too extensive interpretations solely from average numbers, and the necessity to consider individual differences as well. The variety in personal navigation habits between our participants suggests that one has to be careful to speak of *the average user of the Web*.

To start with, the intensity of Web use differed strongly between our participants (see Table III); the mean number of pages visited per active day[19] ranged from 25 to 284—a span of more than one magnitude. Furthermore, they used the Web for different purposes, worked on various tasks, and consequently visited many different sites. This already became apparent in the variation of Google search use: although Google was one of four most frequently visited sites for each user, the share of visits to Google search pages varied between 4% and 39% (Table III). Similarly, some users had a couple of sites that they visited extremely often—visits to the top four sites covered up to 79% of all page requests for some users—whereas others did not have any site that accounted for more than 10% of her visits. Various individual motives for extensive use of certain sites could be identified: we observed regular work-related tasks (project sites, B2B systems), personal interests (news, bulletin boards, auctions), and also some temporal objectives (travel planning, Christmas shopping).

Another often quoted value of Web use that varied strongly between participants was the recurrence rate. Our average rate of 45.6% is composed of individual rates ranging from 19% to 65% (Table III). Consequently, an estimation of a

[19]For our calculations we considered only "active days"—days in which at least one event was logged—as some participants accessed the Web only on certain days of the week and others went on trips during the study.

Table III. Some Individual Differences of This Web Browsing Study

	Average	Std. Dev.	Min	Max
No. of visits	5491	6565	912	30756
Visits per day	89.8	63.4	24.9	283.6
Google share	16.6%	8.6%	3.9%	39.3%
Recurrence rate	45.6%	10.3%	18.6%	64.8%
Number of open Windows	2.10	1.48	1.07	8.19
Actions per window	6.7	2.6	3.2	13.5
Link	39.5%	8.6%	18.8%	56.8%
Back	13.7%	6.6%	3.3%	29.2%
Submit	15.2%	6.3%	7.0%	29.8%
New window	13.6%	7.3%	3.8%	30.8%
Direct access	9.9%	4.9%	3.8%	19.4%
Reload	1.4%	1.5%	0.0%	5.9%
Forward	0.5%	0.4%	0.0%	1.7%
Other	3.3%	3.1%	0.2%	10.7%

common recurrence rate can be misleading as the user requirements regarding revisitation support might differ. Furthermore, even the highest personal recurrence rate of our participants (65%) was about 16% lower than the average rate reported by Cockburn and McKenzie [2001] five years earlier (81%). This can be partly explained by the fact that Cockburn and McKenzie truncated some URIs as they did not consider HTTP query parameters for search result pages [McKenzie and Cockburn 2001, p. 503). We advocate a definition that considers both GET and POST parameters, as they determine the page contents like any other part of the URI[20] (see also Obendorf [2007], p. 601). These different interpretations show how carefully such values have to be analyzed.

The differences in activities and individual habits are also reflected in the personal shares of navigation actions (Table III, lower part). While links were on average the most frequently used navigation element, they made up less than 19% of the actions for one participant; by contrast, her form submissions amounted more than twice the average (26%). From the interviews we know that her job required her to search spare parts in different online catalogues and to order products in several B2B shops. As none of the shop systems supported backtracking well, her back button use was even below 4%.

We also found major differences in how users employed *direct access* actions for returning to frequently visited pages. The detailed Firefox log as well as the interviews revealed diverse stable personal habits: some mainly used the bookmark menu, others solely preferred the bookmark toolbar, and a few had the habit of typing the URI of their favorite pages into the address bar, using its auto-completion function when available. These differences show the importance of customization for the interaction with the browser. On the other hand they also suggest that none of the current revisitation tools does fulfill all needs [Abrams et al. 1998; Jones et al. 2001].

In order to explain differences in user behaviors, we tried to identify different user groups in our population. However, we did *not* find any supportable effects

[20]Apart from the 'fragment identifier' that only results in page scrolling.

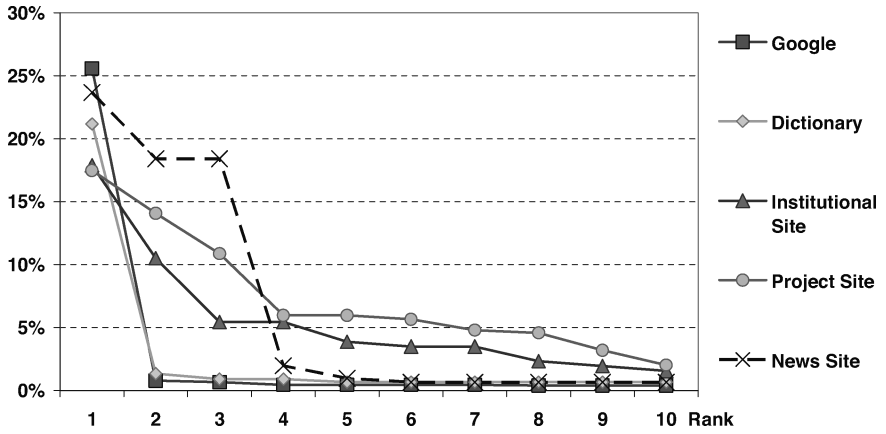

Fig. 7. Distribution of page revisits for different categories of Web sites.

of profession, gender or nationality.[21] Individual differences were mainly caused by user tasks—which also differed significantly between members of the same department or firm—personal habits, private interests and, accordingly, the sites visited.

In fact, the character of the visited sites is a factor that had strong influence on user behavior and page revisitation. As a general rule, if sites were visited more frequently, more pages were visited within the site as well ($r = 0.903$, $p < 0.01$). However, we observed a distinction between sites in which only a small number of pages are visited frequently and sites in which users continue to visit new pages.

Figure 7 illustrates this effect: search engines and dictionaries provide a single portal page as access point; from this page a query is issued, which leads to various result pages. Hence, by their very nature, these sites have only one "popular" page and a long tail of pages that are visited only a few times. By contrast, institutional and project Web sites often have a portal page which is visited quite often, but also a range of other pages that are revisited regularly; these pages may offer information on a certain topic or department, or may provide applications which are used on a regular basis. Finally, several news sites provided a few frequently visited pages; they relate to overview pages of certain news categories the user was interested in.

From the above it can be concluded that reasoning only on the average values of quantitative Web usage studies is often not advisable as it may lead to an overly simplified model of user interaction with the Web; rather than designing for the average, browser and Web site designers should take the different requirements into account and provide *adaptable systems* and *alternative interaction possibilities*.

[21]One reason might be the limited sample size of our study, as we could observe some tendencies. For instance, none of the Dutch participants made regular use of online dictionaries and daily Web users tended to open more browser windows in parallel than infrequent users. However, the effects of personal preferences and visited sites outweighed demographical influences.

Furthermore, we believe it is often worthwhile to consider the influence of use contexts even more than we were able to do in this study. Requirements caused by user tasks and the materials employed in conjunction with the Web client otherwise often stay buried. One source could be the logging of other office applications as well, like word processors and email clients. Within the context of this extensive long-term study on Web usage, we were not able to analyze many contextual aspects of Web use. Although the interviews at the beginning and at the end of the study revealed many substantial user problems, they could hardly deliver all details to understand all aspects of user actions.

5. CONCLUSIONS

This paper presents results of a naturalistic long-term field study that captured the Web browsing behavior of 25 participants with diverse backgrounds and tasks. In addition to confirmations of some results from former studies, 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 high number of new window events suggests that interaction with the Web client has changed for many users from single-window hypertext navigation to a new mode where several paths are followed in parallel and Web browsers are kept open in the background to observe pages or are dedicated to special online services.

Since navigation support of current browsers has changed little from 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 user is often unable to retrieve the contents of a previously viewed page if that page was dynamically generated or required a certain workflow. The concept of the back button fails for users of multiple windows or tabs, as every single viewing area creates its own history stack. Furthermore, with dynamic pages from http post data and AJAX applications we face problems similar to frame pages in the mid-nineties, when the back button did not return to a useful browser state but reloaded the page contents of the last visited URL. In consequence, users are now faced with new cognitive demands when they browse the Web.

Our results confirm that browsing is a rapidly interactive activity. Even new pages with plentiful information and many links are regularly viewed only for a brief period—an interesting background for Web designers, who could focus on offering concise pages that load fast. The analysis of link click positions shows that users scroll regularly—even on navigation pages. Still, about 45% of 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. Interface standards for the Web that consider such findings would help to make navigation on unfamiliar sites easier and quicker (see also Hoffman [1997]; Nielsen [2004]).

Web designers need to consider the limitations of browser real estate even for users with a high screen resolution. Our data shows that an increase in screen resolution does not necessarily mean that users employ the whole screen for

Web browsing for various reasons—working with several programs in parallel, one-click access to toolbars, and instant messengers, among others. These results give fuel to those who advocate Web designs that allow for flexible and stretchable screen designs [Nielsen 2006b].

Finally, the many personal differences and user habits indicate the need for future browsers to become even more adaptable. An extension concept as offered by the open-source browser Firefox does already provide a high degree of flexibility, and several available extensions address some of the listed problems. However, for most of our participants, the task to select and install such extensions was too demanding—only 3 of our 15 Firefox users had installed extensions by themselves. As long as new solutions are not part of the browser distribution, their usefulness is limited as they stay difficult to choose and employ for many people. It seems that an *overall concept* for browser development is missing, especially for the interaction with Web applications and dynamic, transient pages. Browsers need to become more flexible and should be able to adapt to the type of Web site, the habits of the users, and their tasks. Further research has to target these challenges so that browsers can catch up with the growing demands of the dynamic Web.

REFERENCES

ABRAMS, D., BAECKER, R., AND CHIGNELL, M. 1998. Information archiving with bookmarks: personal Web space construction and organization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'98)*, Los Angeles, CA, ACM Press, 41–48.

BAKER, K. R. 2003. The impact of paging vs. scrolling on reading online text passages. *Usability News* 5, 1. http://psychology.wichita.edu/surl/usabilitynews/51/paging_scrolling.htm.

BARRETT, R., MAGLIO, P. P., AND KELLEM, D. C. 1997. How to personalize the Web. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '97)*, Atlanta, GA: ACM Press, 75–82.

BERNARD, M., BAKER, K. R., AND FERNANDEZ, M. 2002. Paging vs. scrolling: Looking for the best way to present search results. *Usability News* 4, 1. http://psychology.wichita.edu/surl/usabilitynews/41/paging.htm.

CAPRA, R. AND PÉREZ-QUIÑONES, M. 2003. Re-finding found things: An exploratory study of how users re-find information. Tech. Rep. cs.HC/0310011, Computing Research Repository (CoRR).

CATLEDGE, L. D. AND PITKOW, J. E. 1995. Characterizing browsing strategies in the World-Wide Web. *Comput. Netw. ISDN Syst. 27*, 6, 1065–1073.

CHISHOLM, W., VANDERHEIDEN, G., AND JACOBS, I. 1999. Web content accessibility guidelines 1.0. World Wide Web Consortium Recommendation. http://www.w3.org/TR/WAI-WEBCONTENT/.

CHO, J. AND GARCIA-MOLINA, H. 2000. The evolution of the Web and implications for an incremental crawler. In *Proceedings of 26th International Conference on Very Large Databases*. Cairo, Egypt. 200–209. http://citeseer.ist.psu.edu/cho00evolution.html.

CHOO, C. W., DETLOR, B., AND TURNBULL, D. 2000. Information seeking on the Web: An integrated model of browsing and searching. *First Monday Online J. 5*, 2. http://firstmonday.org/issues/issue5_2/ choo/index.html.

COCKBURN, A. AND MCKENZIE, B. 2001. What do Web users do? An empirical analysis of Web use. *Inter. J. Hum.-Comput. Stud. 54*, 6, 903–922.

COCKBURN, A., MCKENZIE, B., AND JASONSMITH, M. 2002. Pushing back: Evaluating a new behavior for the back and forward buttons in Web browsers. *Inter. J. Hum.-Comput. Stud. 57*, 5, 397–414.

CONKLIN, J. 1987. Hypertext: An introduction and survey. *IEEE Comput. 20*, 9, 17–41.

CUNHA, C. R., BESTAVROS, A., AND CROVELLA, M. E. 1995. Characteristics of WWW client-based traces. Tech. Rep. BU-CS-95-010, Computer Science Departments, Boston University. http://citeseer.ist.psu.edu/44710.

ECMA-262. 1999. ECMAScript language definition. ECMA International. http://www.ecma-international.org/publications/files/ecma-st/ECMA-262.pdf

FETTERLY, D., MANASSE, D., NAJORK, D., AND WIENER, J. 2003. A large-scale study of the evolution of Web pages. In *Proceedings of the 12th International World Wide Web Conference*. Budapest, Hungary, 669–678.

GREENBERG, S. AND COCKBURN, A. 1999. Getting back to back: Alternate behaviors for a Web browser's back button. In *Proceedings of 5th Human Factors and the Web Conference*, Austin, TX.

HAAS, S. W. AND GRAMS, E. S. 1998. A link taxonomy for Web pages. In *Proceedings of the 61st Annual Meeting of the Information Access in the Global Information Economy (ASIS'98)*. Information Today, Medford, NJ, 485–495.

HALASZ, F. G. 1988. Reflections on notecards: Seven issues for the next generation of hypertext systems. *Commun. ACM 31*, 7, 836–852.

HERDER, E. 2006. *Forward, back and home again—analyzing user behavior on the web*. Ph.D. Thesis, University of Twente, The Netherlands.

HOFFMAN, M. 1997. Enabling extremely rapid navigation in your web or document, 2nd Version, (Dec.) http://www.hypertextnavigation.com/infoaxcs.htm.

HYAMS, J. AND SELLEN, A. 2003. Gathering and sharing web-based information: Implications for ePerson concepts. HP Labs: Tech. rep. HPL-2003-19.

HYATT, D. 2001. XML user interface language (XUL) 1.0. technical specification, Mozilla project. http://www.mozilla.org/projects/xul/xul.html.

ISO 9241/110. 2006. Ergonomics of human-system interaction—Part 110: Dialogue principles. International Standards Organization.

JANSEN, B. J. AND POOCH, U. W. 2000. A review of Web searching studies and a framework for future research. *J. ASIS 52*, 3, 235–246.

JANSEN, B. J., SPINK, A., AND SARACEVIC, T. 2000. Real life, real users and real needs: A study and analysis of user queries on the Web. *Inform. Proc. Manag. 36*, 2, 207–227.

JANSEN, B. J. AND SPINK, A. 2005. How are we searching the World Wide Web? A comparison of nine search engine transaction logs. *Inform. Proc. Manag. 42*, 248–263.

JONES, W., BRUCE, H., AND DUMAIS, S. 2001. Keeping found things found on the Web. In *Proceedings of 10th International Conference Information and Knowledge Management (CIKM)*, 119–134.

KAASTEN, S. AND GREENBERG, S. 2001. Integrating back, history and bookmarks in Web browsers. In *Extended Abstracts of Conference on Human Factors in Computing Systems (CHI'01)*, Seattle, WA. 379–380.

KEHOE, C., PITKOW, J., SUTTON, K., AGGARWAL, G., AND ROGERS, J. 1998. 10th GVU WWW user survey. http:// www.cc.gatech.edu/gvu/user_surveys/survey-1998-10.

KELLAR, M., WATTERS, C., AND INKPEN, K. M. 2007. An exploration of Web-based monitoring: Implications for design. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI'07)*, San Jose, CA, 377–386

KNIGHT, A., PYRZAK, G., AND GREE, C. 2007. When two methods are better than one: Combining user study with cognitive modeling. In *(CHI'07) Experience Reports*, ACM Press, 1783–1788

LYNCH, P. J. AND HORTON, S. 2002. *Web Style Guide: Basic Design Principles for Creating Web Sites*. Yale University Press, Yale.

MCKENZIE, B. AND COCKBURN, A. 2001. An empirical analysis of Web page revisitation. In *Proceedings of Howai International Conference on System Sciences (HICSS'01)*, Maui, Hawaii. IEEE Computer Society Press, 501–509.

MESCHKAT, S. AND MITTLEMAN, J. 2007. Keeping the Web in Web 2.0: An HCI approach to designing Web applications. *CHI 2007 Course*, http://ajaxchi.googlepages.com/.

MILIC-FRAILING, N., JONES, R., RODDER, K., SMYTH, G., BLACKWELL, A., AND SOMMERER, R. 2004. Smartback: Supporting users in back navigation. In *Proceedings of 13th International World Wide Web Conference*. ACM Press, 63–71. http://www2004.org/proceedings/docs/1p63.pdf.

MORKES, J. AND NIELSEN, J. 1997. Concise, SCANNABLE, and objective: How to write for the web, Sun Microsystems. http://www.useit.com/papers/webwriting/writing.html.

MOYLE, M. AND COCKBURN, A. 2003. The design and evaluation of a flick gesture for "back" and "forward" in Web browsers. In *Proceedings of the 4th Australasian User Interface Conference*. Adelaide, Australia. 39–46.

NICHOLAS, D., HUNTINGTON, P., JAMALI, H. R. AND TENOPIR, C. 2006. What deep log analysis tells us about the impact of big deals: Case study OhioLINK, *J. Document. 62*, 4, 482–508

NIELSEN, J. 1997. The tyranny of the page: Continued lack of decent navigation support in version 4 browsers. Useit Alertbox, Nielsen Norman Group. http://www.useit.com/ alertbox/9711a.html.

NIELSEN, J. 2004. The need for Web design standards. Useit Alertbox, Nielsen Norman Group. http://www.useit.com/alertbox/20040913.html.

NIELSEN, J. 2005. Ten best intranets of 2005. Useit Alertbox, Nielsen Norman Group. http://www.useit.com/alertbox/20050228.html.

NIELSEN, J. 2006. F-shaped pattern for reading web content. Useit Alertbox, Nielsen Norman Group. http://www.useit.com/alertbox/reading_pattern.html.

NIELSEN, J. 2006b. Screen resolution and page layout. Useit Alertbox, Nielsen Norman Group. http://www.useit.com/alertbox/screen_resolution.html.

OBENDORF, H., WEINREICH, H., AND HASS, T. 2004. Automatic support for Web user studies with SCONE and TEA. In *Extended Abstracts of CHI'04*. Vienna, Austria. 1135–1138.

OBENDORF, H., WEINREICH, H., HERDER, E., AND MAYER, M. 2007. Web page revisitation revisited: Implications of a long-term click-stream study of browser usage. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'07)*. San Jose, CA. 597–606.

OTTAWAY, M. 2004. Nielsen/Netratings data support TVNZ's widescreen Web decision. Nielsen/Netratings. http://www.nielsen-netratings.com/pr/pr_040504_nz.pdf.

OUTING, S. AND RUEL, L. 2004. The best of Eyetrack III: What we saw when we looked through their eyes. Eyetrack III, http://www.poynterextra.org/eyetrack2004/main.htm.

PETERSON, E. 2004. Which screen resolution should you be using? *Jupiterresearch Analyst Weblogs*. http://weblogs.jupiterresearch.com/analysts/peterson/archives/003965.html.

PITKOW, J. E. 1998. Summary of WWW characterizations. The *Web J. 2*, 1, 3–13.

ROSE, D., AND LEVINSON, D. 2004. Understanding user goals in Web search. *In Proceedings of 13th International Conference on World Wide Web*. ACM Press. 13–19.

SCHMIDT-MAENZ, N. AND KOCH, M. 2006. A general classification of (search) queries and terms. In *Proceedings of 3rd International Conference on Information Technology: New Generations*. Las Vegas, NV. IEEE Xplore, 375–381.

SELLEN, A. J., MURPHY, R. M., AND SHAW, K. 2002. How knowledge workers use the Web. In *Proceedings of SIGCHI Conference on Human Factors in Computing Systems (CHI'02)*, Minneapolis, MN. ACM Press, 227–234.

SHRESTHA, S. AND LENZ, K. 2007. Eye gaze patterns while searching vs. browsing a Website. *Usability News 9*, 1, January 2007, Software Usability Research Lab, Wichita State University.

SPILIOPOULOU, M., MOBASHER, B., BERENDT, B., AND NAKAGAWA, M. 2003. A framework for the evaluation of session reconstruction heuristics in Web usage analysis. *INFORMS J. Comput. 15*, 2, 171–190.

SPOOL, J., SCANLON, T., SCHROEDER, W., SNYDER, C., AND DEANGELO, T. 1998. *Web Site Usability: A Designer's Guide.* Morgan Kaufmann Publishers, CA.

TAUSCHER, L. AND GREENBERG, S. 1997. How people revisit Web pages: Empirical findings and implications for the design of history systems. *Inter. J. Hum. Comput. Studies 47*, 1, 97–138.

TEEVAN, J., ALVARADO, C., ACHERMAN, M., AND KARGER, D. 2004. The perfect search engine is not enough: A study of orienteering behavior in directed search. In *Proceedings of SIGCHI Conference on Human Factors in Computing Systems (CHI'04)*. Vienna, Austria. ACM Press, 415–422.

TEEVAN, J., ADAR, E., JONES, R., AND POTTS, M. 2006. History repeats itself: Repeat queries in Yahoo's logs. In *Proceedings of the 29th Annual International ACM SIGIR Conference (SIGIR'06)*. (Poster Sessions), Seattle, WA, 703–704.

THE COUNTER, 2007. *Global Screen Resolution Statistics*. Jupitermedia Corporation.

UNIVERSITY OF SOUTHERN CALIFORNIA. 2007. Digital Future Report. Annenberg School, Center for the Digital Future. http://www.digitalcenter.org/pages/current_report.asp?intGlobalId=19.

WEINREICH, H. AND LAMERSDORF, W. 2003. Eine Umfrage zu Link- und Objekt-Attributen im Web. In *Proceedings of Mensch und Computer.* Stuttgart, Germany. Teubner Verlag, 387–390. http://vsis-www.informatik.uni-hamburg.de/getDoc.php/publications/121/weinreich-umfrage-2003.pdf.

WEINREICH, H., OBENDORF, H., AND HERDER, E. 2006. Data cleaning methods for client and proxy logs. In *Proceedings of the Workshop of the World Wide Web Conference*. Edinburgh, UK. http://torch.cs.dal.ca/~www2006/Workshop_submissions.html.

WHITE, R. W. AND DRUCKER, S. M. 2007. Investigating behavioral variability in Web search. In *Proceedings of 16th International World Wide Web Conference 2007*. Banff, Canada. ACM Press, 21–30.

Received September 2006; revised July 2007; accepted August 2007