

A Study of Tabbed Browsing Among Mozilla Firefox Users

Patrick Dubroy, Ravin Balakrishnan

Department of Computer Science

University of Toronto

{dubroy, ravin}@dgp.toronto.edu

ABSTRACT

We present a study which investigated how and why users of Mozilla Firefox use multiple tabs and windows during web browsing. The detailed web browsing usage of 21 participants was logged over a period of 13 to 21 days each, and was supplemented by qualitative data from diary entries and interviews. Through an examination of several measures of their tab usage, we show that our participants had a strong preference for the use of tabs rather than multiple windows. We report the reasons they cited for using tabs, and the advantages over multiple windows. We identify several common tab usage patterns which browsers could explicitly support. Finally, we look at how tab usage affects web page revisitation. Most of our participants switched tabs more often than they used the back button, making tab switching the second most important navigation mechanism in the browser, after link clicking.

Author Keywords

WWW, Web browsing, Hypertext, Tabs, Tabbed document interfaces, Web browser interfaces.

ACM Classification Keywords

H5.4. Hypertext/Hypermedia: User issues, navigation.

General Terms

Human Factors

INTRODUCTION

Despite the fact that both web browser usage and the web itself have changed immensely since inception, there have been very few significant changes to the user interface of mainstream web browsers. A user familiar with the browsers of the mid-1990s would likely have no trouble using the latest versions of today's most popular browsers, Microsoft Internet Explorer and Mozilla Firefox. The main navigation mechanisms—hyperlinks, the back and forward buttons, the URL bar, bookmarks, and the history—have remained almost unchanged for more than 14 years. One exception is the inclusion of tabbed browsing—which allows multiple web pages to be open concurrently under a single top-level browser window, each page in its own “tab”—as a standard part of the browser interface.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2010, April 10–15, 2010, Atlanta, Georgia, USA.

Copyright 2010 ACM 978-1-60558-929-9/10/04...\$10.00.

According to estimates [1, 3], 50% of people now use a browser that supports tabbed browsing. This figure is expected to continue to increase, as Internet Explorer users upgrade to the latest version, and alternatives such as Firefox and Safari continue to gain market share.

During the same time period, there has been a rise in the usage of web-based applications—such as Google's web-based office suite, and web-based email clients—which replace conventional desktop applications. As a result, web *browsing* is ceasing to become the only task that people perform using a browser. Rather, the browser is becoming an alternative interface for many traditional computing tasks. The growing popularity of tabbed browsing may reflect the desire of browser users to have better ways of managing the increasing complexity of their activities on the web.

Problems with Tabs

While tabs are intended to make it easier for users to organize and manage web browsing sessions, they also introduce some potential problems. First, tabs are another organization and task management system that is completely separate from the mechanisms provided by the operating system. For example, in Microsoft Windows, one can use the taskbar to switch between open applications, but only the currently active tab in a browser window will show up in the taskbar. Similarly, on Mac OS X, the Exposé feature allows the user to quickly see a thumbnail image of all open windows, but this will not include any inactive tabs in web browser windows.

Another potential problem with tabbed browsing is that it interferes with the use of the back button. Studies have consistently found that the back button is the second most commonly used navigation mechanism [6, 14, 16], behind clicking on hyperlinks. (Note that when we refer to use of the back button, we also include all its keyboard shortcuts.) The use of multiple tabs complicates the use of the back button, because each tab has its own history stack. Even if a tab was opened by clicking on a link, the back button will not take you back to the originating page, because it is in another tab.

Finally, tabbed browsing adds another level of complexity to the web browser interface. Tabs add more visual clutter to the user interface, and like managing applications and windows, tab management becomes another necessary task that distracts from the ultimate goal at hand.

Motivation

Despite these issues, tabs seem to be a very popular feature among web browser users. During preliminary investigations, we found that many people find tabs to be an indispensable feature. Perhaps the issues we described do not actually create significant problems for users, or maybe the advantages of using tabs simply outweigh the negatives. Given the increasing reliance on the web browser as the interface to more of our computing applications, it is important that we understand how and why people use tabs, in order to guide the design of future web browser interfaces by appropriate usage data. Our study provides data in this regard, as well as insight into what purposes tabs serve, and how tabs relate to other features of the web browser and the overall computing experience.

RELATED WORK

Long-term Click-Stream Logging

Several studies have used client-side “click-stream” logs to study web browsing behaviour. In one of the earliest, Catledge and Pitkow [6] used an instrumented version of XMosaic to record the actions of 107 participants over the course of 21 days. A few years later, Tauscher and Greenberg [16] performed a similar study, analyzing the revisitation patterns of 23 people over the course of 5–6 weeks each. Their study defined the *recurrence rate* as the probability that any URL visited is a repeat of a previous visit, and found a mean recurrence rate of 58% across their participants, and 61% from 55 of the subjects from the Catledge and Pitkow study. Based on these findings, they concluded that the web can be characterized as a “recurrent system.”

Both these studies examined the percentage of navigation actions caused by the various navigation mechanisms of the browser, and found that clicking on hyperlinks is the most frequent cause of navigation actions, followed by backtracking (using the back button or one of its shortcuts). Catledge and Pitkow [6] found that clicking on links accounted for 52% of navigation actions, and backtracking 41%. Tauscher and Greenberg [16] found both of these numbers to be lower, at ~40% and ~30%, respectively.

These two studies were performed in the mid-1990s, when the web was still in its infancy. Given the evolution since then of both the web and web browser user interfaces, it is questionable how applicable the results are today. In 2004 and 2005, Weinreich et al. [17, 18] conducted a long-term click-stream study of 25 participants over a period of 52–195 days. They re-confirmed link clicking and the back button to be the most common causes of navigation actions, but found that backtracking accounted for a much smaller percentage, only 14.3%. Other recent studies have found this number to be slightly higher: Milic-Frayling, et al. [13](23%); and Kellar et al. [12] (19%).

Qualitative Studies of User Behaviour on the Web

While low-level click-stream studies like the ones mentioned above can provide a great deal of insight, they

do not capture the whole picture. Without knowing *why* people use the web the way they do, it is difficult to strongly infer implications for design. Hence, other studies have attempted to supplement click-stream data with qualitative information provided by the user.

Byrne et al. [4] videotaped 8 participants who used a verbal protocol during web browsing to describe what they were doing. From this data, the researchers extracted a taxonomy of web browsing tasks at a similar level as previous click-stream studies, but capturing the user's goals rather than just the observed behaviour.

Choo et al. [7] performed a higher-level categorization of information seeking tasks. They studied the web browsing behaviour of 34 knowledge workers over a period of two weeks. In follow-up interviews, participants were asked to recall and discuss “critical events”. Through analysis of these interviews, the researchers identified four high-level categories of information seeking. The approach taken in this study was quite similar to that used in our study. However, the click-stream logs were only used to support their qualitative analysis, whereas in our study, the quantitative and qualitative results are complementary.

Sellen et al. [15] observed 24 knowledge workers for two days each. At the end of each day, participants were interviewed in front of their browser history, and asked to describe each of the activities that they performed. The analysis divided the activities into 6 main categories of behaviour. Like [7], they focused on the qualitative results.

Kellar et al. [12] performed a study that used click-stream data while also asking participants to categorize their current task into one of five categories. Although the goals of their study were different from ours, as they sought to understand how tasks influenced user behaviour, their approach was quite similar, as both the quantitative and qualitative results were presented with equal importance.

STUDY BACKGROUND

Broadly speaking, tabs serve two main functions. First, using multiple tabs allows you to put a page aside and return to it later. In a way, this can also be accomplished using the back button, bookmarks, and the history. With tabs, however, you can have many pages open concurrently and switch between them in any order. The same thing can be accomplished using multiple windows. In most browsers, new windows and new tabs can be created in very similar ways.

The second major function of tabs is the ability to “fork the chain of navigation”. A link can be opened in a new tab without disturbing the current page. For example, on a search results page, one might open several links in new tabs, allowing the ability to switch between the tabs to compare the various results.

It is also possible to open a link in a new window. The main difference between opening a link in a new tab versus a

new window is that new tabs can be opened in the background without disturbing the current page. You can switch to the new page right away, or stay on the current page and view the new page later. New windows, on the other hand, generally open in the foreground, taking focus away from the current window.

Thus, the major functionality of tabs—with the exception of being able to open links in the background—can be duplicated by using multiple windows. And a major part of that functionality can also be accomplished, albeit with a few key differences, by using the traditional revisitation mechanisms: the back button, the history, and bookmarks.

Given that most tab functionality is not fundamentally new or different, then the popularity of tabs indicates that they likely have some more subtle advantages over using multiple windows and the traditional browser revisitation mechanisms. Our study seeks to discover these advantages.

METHODOLOGY

Participants

21 people (13 female, 8 male; 15 aged 18-29, 4 in their 30s, and 2 in their 50s) were recruited via email through the extended network of the researchers, and by posters on bulletin boards around campus. The study was advertised as “a research study exploring how people use web browsers”, looking for participants who “use Mozilla Firefox for several hours a day, and often use multiple tabs or windows.” The goal of our study was to learn how and why tabs are used by active tab users. In order to learn the most about different kinds of tab usage, we felt that there would be little use in studying people who rarely use tabs. Also, while our main goal was to examine tab usage, we thought we could gain some insight by also having participants who use multiple windows rather than tabs. However, only one participant showed a preference for multiple windows over tabs.

The selection of only Firefox users was pragmatic, since Firefox has a powerful extension mechanism that made it relatively easy to write our click-stream logger. Previous studies have had participants who were either entirely from Computer Science backgrounds [6, 9, 16], or mostly so [18]. In contrast, participants in our study came from more varied backgrounds: only 5 participants had a computer science or engineering background, while others had studied education, environmental science, business, and psychology. 6 of the participants were full-time students, and 15 were working in office environments where they spent most of their time on the computer.

Instruments and Data Collection

Four types of data were collected from study participants. Our Firefox extension gathered click-stream logs and free-form diary entries. The researchers also interviewed each participant 2–4 times over the course of the study, and some basic demographic data was collected via a pre-study questionnaire.

Click-stream Data

Click-stream data was collected using a custom extension for Firefox that was compatible with Firefox versions 1.5 through 3.0 so that it could be installed without upgrading or significantly changing the user's browsing environment. Participants were asked to install the extension on their primary computer. The extension logged low-level user interface and network events, and linked related events together at a level of granularity similar to [12].

In order to collect results that accurately represented the natural web usage of our participants, we felt that it was important to take several steps to protect their privacy. First, the actual URLs and sites that participants visited were not made available to the researchers. However, in order to determine the relationships between the various pages that each participant viewed over the course of the study, our extension associated a unique code for each site and URL that the participant visited. In this way, we were able to determine (for example) whether a particular page had been previously visited, or whether a link led to a new site or to another page on the same site.

The second step that we took to protect the privacy of our participants was to ensure that our extension did not automatically submit any information to the researchers. All of the click-stream data was stored in a human-readable format, which the participants could review before manually submitting to the researchers via email.

Diary Entries and Interviews

Qualitative data was obtained from the participants in several ways. First, using random sampling, our extension periodically prompted the user to make a short diary entry describing what windows and tabs they had open, and why they had them open. These prompts were designed to be unobtrusive, so that our extension wasn't overly distracting or annoying. The prompts appeared as a banner at the top of the current page. The user could choose to ignore any of the prompts without taking any explicit action.

Due to the short nature of the diary entries, there was little information that could be directly extracted from them. Overall, the diary entries served less as a data-collection mechanism, and more as the structure around which to conduct the interviews. Many participants commented that the prompts caused them to reflect on their tab usage in a way that they previously hadn't.

Before each interview, the researchers reviewed the participant's logs, including his or her diary entries, and identified any interesting aspects to discuss during the interview, including any unusual or anomalous events in the logs. For example, if a participant didn't usually use multiple browser windows but did on one specific occasion, the researchers asked the participant for further details.

Procedure

After completing the initial questionnaire and installing the extension, each participant was instructed by one of the

researchers on how the diary entries would work, and then the researchers scheduled a time for the first interview. Participants were asked to email their logs to the researchers at a point prior to the next scheduled interview. Once the logs were received by the researchers, they were examined for any events that required follow-up. This was determined on a case-by-case basis by the researchers, looking for diary entries that were unclear, or any anomalous or exceptional events that required description. Interviews were conducted either in person or by phone, and varied in length from 15–45 minutes.

At the end of the two week period, participants were given instructions on how to remove the Firefox extension and all associated logs. For participants who had only installed the extension on their work computer, an effort was made to extend the study to a full 14 days of use. After the extension was uninstalled, participants were paid \$50 for their participation in the study. One participant (participant 8) opted to end participation part way through the study.

Data Analysis

The two main sources of data analyzed were the click-stream logs (for quantitative results) and the interviews. The quantitative data analysis was done through a set of custom scripts written in Python. Interviews were audio recorded and transcribed by the researchers and a research assistant. For the qualitative data analysis, analysis codes were developed in a first pass through the interview transcripts using open coding. A research assistant performed an independent open coding pass on a subset of the data, and the final set of codes was agreed upon by the two coders. Finally, all of the data was coded independently by the two coders, and all discrepancies were discussed until agreement was reached.

RESULTS

Rather than present the results of our quantitative and qualitative analysis separately, we take a more holistic approach, presenting the results organized under three themes. First, we demonstrate that participants in our study showed a marked preference for the use of tabs over multiple windows, and we discuss our participants' stated reasons for this preference. Second, we take a closer look at the tab usage of our participants. We examine several measures of tab usage from the click-stream logs, and identify a subset of users (“tab power users”) who used tabs much more frequently than the others. We discuss our participants' reported reasons for using tabs, and identify a key difference between the tab power users and the others: the tab power users reported habitually opening links in new tabs. Third, we examine in greater depth one aspect of tab usage which we believe to be particularly significant: the use of tabs as a revisitation mechanism.

Analysis Methods

In our call for participation, we asked for Mozilla Firefox users who “often use multiple tabs or windows.” This is a subjective description that needs to be quantified. What

does it mean to “use” multiple tabs or windows, and how should it be measured? There are several possible answers.

Number of tabs or windows created. New tabs or windows may be created by the user, or automatically by a web site; we counted both. We did not count windows or tabs created automatically on browser startup.

Number of focus switches. Since a new tab must be switched to in order to use it, we consider tab switches to be a similar and better measure of tab usage than tab creation.

Number of concurrent tabs and windows. We measured the number of windows and tabs that were open when a navigation event occurred, as in Weinreich et al. [18]. For simplicity, we ignored all navigation events caused by Firefox's session restore feature in calculating this measure.

In order to facilitate comparisons between participants, in some cases we normalized the measures by dividing by the total number of navigation actions (excluding session restore) for the participant. We refer to the normalized measures as the *window creation rate*, *tab creation rate*, and *tab switch rate*.

Use of Multiple Windows vs. Tabs

Click-Stream Analysis

Figure 1 shows the window and tab creation rates of all participants as calculated from the click-stream logs.

Only one participant (P10) had a window creation rate higher than the tab usage rate—she was the only one who was more likely to create a new window than an additional tab in an existing window. She created a new window more than three times more often than a tab, and created a total of 338 windows over the course of the study, more than any other participant. However, only 19% of the windows were created intentionally. This means that she frequently uses sites (or one particular site) where links open in a new window.

Apart from P10, only three other participants used windows nearly as often as they used tabs: P15, P11, and P18. The remaining 17 participants were at least twice as likely to create an additional tab rather than an additional window. So, despite our call for participation asking for people who “frequently use multiple tabs or windows,” there was a strong bias towards the use of tabs over windows.

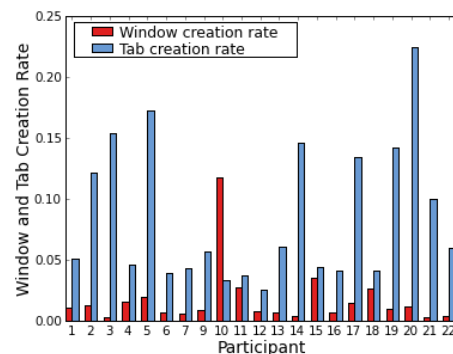


Figure 1: Window and tab creation rates

Qualitative Comparisons

Since using tabs is functionally very similar to using multiple windows, one of the first questions we asked participants in our study was, “Before you used a browser that supported tabs, did you use multiple windows in the same way that you use tabs now?” 17 of the 21 participants said that they use tabs more frequently than they used multiple windows in the past. Many of the participants, including all four who reported using multiple windows heavily in the past, expressed a preference for tabs.

These subjective reports reinforce the results presented above that most participants in our study showed a strong preference for multiple tabs rather than multiple windows. Our participants cited many different reasons why using multiple tabs is preferable to using multiple windows. The most common reason (cited by 10 of 21 people) was that using tabs is cleaner, more organized, or less cluttered.

“It was more messy, I think. Like when you have all these multiple windows open [...] it was more confusing.” (P13)

This result was a bit surprising, because it's not clear how exactly tabs are more organized than multiple windows. Most of our participants used Microsoft Windows, where multiple browser windows would be accessed via the task bar along the bottom, quite similar to how tabs are lined up across the top of the browser window.

Five participants said that it was helpful that tabs kept their web browsing “all in one place,” whereas multiple browser windows would be interspersed with other application windows. This could be one reason why tabs are seen to be cleaner or more organized. Two participants said that combining windows and tabs allows multiple levels of organization, but the low window usage we observed makes it unlikely that this is a key advantage.

The second most cited reason (by 7 of 21 participants) for using multiple tabs instead of multiple windows was that tabs are easier or more convenient to access or switch between. One possible reason for this is suggested by Weinreich et al. [18], who found that 45% of all link clicks occurred in the upper left quadrant of the browser window. If a person is much more likely to click on a link near the top of the page (for whatever reason), it might be easier or more convenient to click on a tab that is nearby, rather than on the taskbar at the very bottom of the screen.

Seven participants reported that the way Windows XP and Windows Vista group windows in the taskbar (an option which can easily be turned off) was a problem, and said that tabs provide a better solution:

“I used to have about 15 windows in [Internet] Explorer. They were all smooshed together on the task bar. [...] It was a big pain in the [expletive deleted].” (P12)

Only three people cited resource usage or performance reasons for using tabs instead of multiple windows. Some other advantages were each mentioned by a single

participant: the order of tabs in the tab bar can be rearranged, and new tabs can be opened in the background, whereas new windows open in the foreground.

Mixing tabs with multiple windows

Although our click-stream analysis showed that 17 participants showed a very strong preference towards the use of multiple tabs over multiple windows, in the interviews, several people reported using multiple windows from time to time. This was verified in the click-stream logs: 10 participants purposely created an additional window at least once during the course of the study. In the interviews, more than half of our participants cited some reason when the use of multiple windows can be useful. One of the most commonly cited reasons (by 8 people) was that using many tabs within separate windows can be beneficial to create a higher-level task grouping. For example:

“I generally don't like mixing and if I am doing two things at once I usually have one window with all the tabs for one thing, and one window with tabs—or even just one tab—for something else.” (P11)

4 of the 8 participants said that they sometimes use windows to separate work- or school-related tabs from personal or “fun” tabs.

The other highly cited reason for using multiple windows (also cited by 8 people) was that it makes it possible to view two pages at the same time. This can be useful when comparing two pages, for example, or when watching a video while browsing some other sites.

Characterizing Tab Usage

To get a better idea of how the tab creation rates are distributed among our participants, Figure 2 plots the tab creation rates as a histogram. There is a very clear bi-modal distribution. Over half our participants (13/21) are clustered around the 0.04 mark. In other words, these people created about 4 tabs for every 100 navigation actions. The rest of the participants are loosely centered around 0.14, meaning that they created about 3 times as many tabs as the other people. Two participants are even higher, creating (respectively) 17 and 22 tabs per 100 navigation actions.

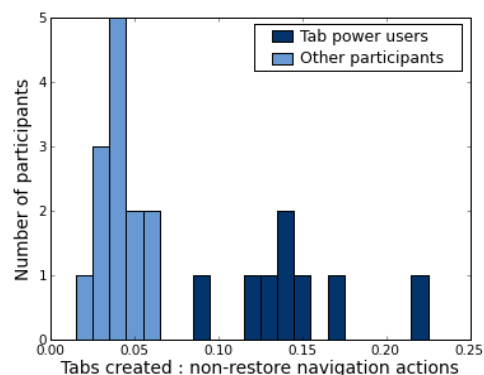


Figure 2: Tab creation rate

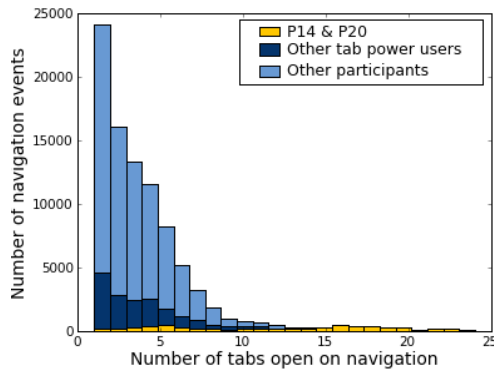


Figure 3: Number of tabs open on navigation

We refer to the 8 users in the higher part of the distribution as “tab power users.” Unsurprisingly, the four highest tab creation rates belong to the four participants with Computer Science and/or programming backgrounds (P3, P5, P14, and P20). However, the other four participants in the high part of the distribution are not stereotypical “power users”: one is a civil engineer, and the others have backgrounds in communications, marketing, and the humanities.

Figure 3 shows the number of tabs open when a navigation action occurred. The most common number of tabs to have open is one, with a steady descent down to 9. There is a second peak at 16 tabs, but this was almost entirely due to just two of the participants, P14 and P20.

For further insight, we examine the median and maximum number of tabs that each participant had open (Figure 4). Participant 8’s data is omitted as he dropped out part way through the study.

Participants 14, 19 and 20 definitely stand out. Participant 14 had *by far* the highest median number of tabs open with 17, while no other participant had a median higher than 6. Participant 20 had the highest number of tabs open at once, with 42. The tab bar would be scrolling two or three times over at this point. In one of his diary entries, he actually commented about having so many tabs open, saying: “Now I am opening tabs from Digg [a social news site] and they are appearing at the end of my massive list. This is truly a bad way to browse.”

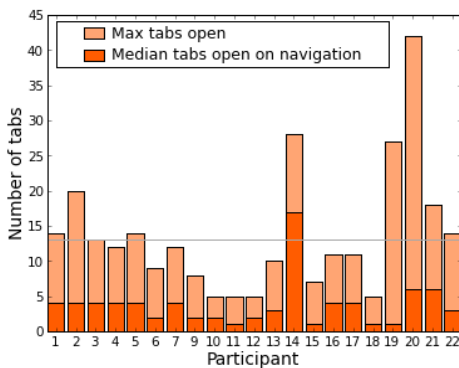


Figure 4: Median and maximum number of open tabs

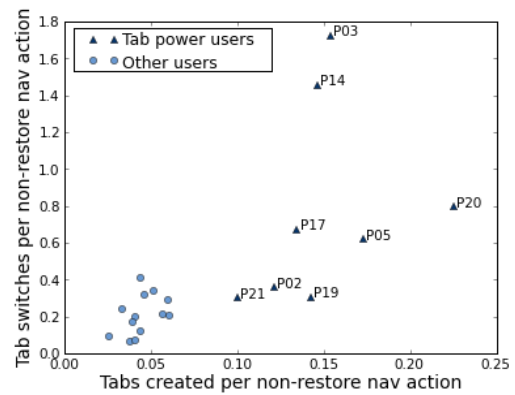


Figure 5: Tab creation rate vs. tab switch rate

P19 is interesting—he had a median of only one tab open, but a max of 27. Participant 2 is similar, but not nearly as extreme: a median of 4 and a max of 20. Even if these two preferred not to have many tabs open most of the time, they weren’t afraid of opening lots of tabs when needed.

Another tab usage measure that’s interesting to look at is the frequency of tab switching. Two users may create tabs at a similar rate, but one might switch between tabs much more often, indicating an important distinction in usage patterns.

In Figure 5, the tab creation rate is plotted against the tab switch rate. As expected, we see a rough correlation between the number of tabs created and the number of tab switches that occur. Intuitively, this makes sense: if more tabs are created, more tab switches must occur in order to make use of those tabs. There are two very clear outliers: participants 3 and 14, two of the people identified as “tab power users”. They switched tabs 2–3 times as often as we would expect based on their tab creation rate. Even among participants who use tabs to a similar degree, there is a considerable variation in the number of tab switches. This is likely partially caused by individual differences, but it may also be due to differences in the type of work being done.

Reasons for using tabs

In addition to quantifying tab usage, one of the main goals of our study was to determine *why* people use tabs in the way that they do. What purposes do tabs serve for them, and what are the advantages of using tabs? We explore these questions via our interview and diary entry data.

Reminders

17 participants reported that they often use tabs as a reminder to do something. This can be related to multitasking; a person might begin a task and then leave the tab open to remind them to complete.

“Having the tab open is a reminder to me. [...] if it’s at the end of the day or lunch time while I am cracking a sandwich or something and I’ll say, ‘Oh yeah, I want to go back and look at that link’ because I see the tab sitting there.” (P16)

Opening links in background

14 participants—including all of the tab power users—said that the ability to open tabs in the background was useful because it allowed them to open several links at once. On a page with many links—such as a search results page (reported by 10 participants) or a news site—these users would first open up all the links of interest, and then go on to read the pages they had opened.

“If I search for something on Google, I just go ‘right-click, rick-click, right-click’—you know, opening all the tabs. And then I would look at them later.” (P2)

P20 reported that *“most links that I click [...] like when I do a Google search, I hold on ctrl and click on links.”* He also reported doing this on Digg, a social news site.

A similar behaviour reported by 6 of the same 14 participants was opening a link in the background while continuing to read the current page. P19, who reported using tabs heavily on Wikipedia, said:

“Usually I find interesting links halfway through the article, and that's why I like tabs.”

Although being able to open links in the background instead of the foreground may seem like a subtle difference, these examples show that it can have a significant effect on the way that people browse.

In fact, this behaviour of opening links in new tabs is likely one of the major differences between the tab power users and the other participants. Every single one of the tab power users—and none of the other participants—reported that opening links in new tabs had become habitual:

“Usually I have this sort of reflex of right-clicking and saying ‘open in a new tab’.” (P2)

“Most of the time, over 70% of the time, opening a new link means to be opening a new tab.” (P20)

P21 had even installed a Firefox extension which opened every link in a new tab. When asked how she would use Google without tabs, she said, “I couldn't imagine it.”

A common behaviour that seemed to coincide with opening many links at once can be termed a *process of elimination*. 11 participants reported having a set of tabs open for a task, and pruning it to a smaller set of the most useful tabs. This could be especially useful during an open-ended search task, when the ultimate goal of the search is not known.

Multitasking

Another highly-cited reason for using tabs (cited by 11 of 21 participants) was to facilitate multitasking: putting the current task on hold and switching to a tab to perform a new task. In many cases, the participants indicated that the current task was being put on hold for only a short period. For example:

“Just a quick little side task, when, you know, I want to come right back to what I was working on.” (P1)

Going “back and forth”

10 participants reported using tabs for comparing between two pages, and 13 reported using tabs to switch back and forth between two or more pages. In many cases, the back button could have accomplished the same task, but participants generally reported that using tabs was easier.

Frequently used pages

4 participants mentioned keeping a page open in a tab when it was being used frequently. For example, one participant was a graduate student and said that he found it convenient to keep open in a tab the web page for a course he was teaching. Interestingly, he said that it was easier to look for the tab that was already open—even if it meant scrolling through the tab bar—rather than clicking on the button in his bookmarks toolbar: “it might not be actually faster *physically*, but it's easier.”

Short-term bookmarks

4 participants reported using tabs as a kind of short-term bookmark. For example:

“If it is something I plan to refer back to in the near future, like later that day or the next day, [...] sometimes I'll leave the tab open.” (P16)

Similarly, 11 participants mentioned using multiple tabs when they were planning to return to a page at some unspecified point in the future. 10 participants also mentioned leaving a page open in a tab even if they weren't entirely sure they'd actually return to the page.

“I will often lose interest in something, and I think I might go back to it, so I will leave the tab open and open a new one. I don't necessarily always go back to them, but the opportunity is there if I want to.” (P12)

This is interesting, because it shows that many people don't associate a high cost with leaving a tab open in the browser; it is worth keeping around even if they are not sure that they are going to return to it. This was reinforced by the 12 participants who reported that they generally didn't make an effort to close tabs when they were finished with them.

The Relationship Between Tabs and Revisitation

The use of tabs can be seen as a kind of revisitation mechanism, providing a shortcut for returning to previously-visited pages. However, using tabs is distinctly different from the traditional revisitation mechanisms—the back button, the history, and bookmarks—because the current page and its navigation history can be left as-is and returned to later, without reloading or re-rendering the page.

Numerous studies have examined and characterized web page revisitation. However, most of those studies took place before tabbed browsing was widely used, with the exception of recent studies by Adar et al. [2] and Weinreich et al. [18].

Decreased Use of the Back Button

Previous studies on revisitation have noted that a relatively large portion of navigation actions are caused by the use of

the back button. However, this seems to be following a downward trend. In Catledge and Pitkow's study [6], the back button accounted for 41% of all navigation actions. A few years later, Tauscher and Greenberg [16] found that it only accounted for about 30%. In one of the more recent studies, Weinreich et al. [14] found it to be only 14%.

Because tabs offer a kind of revisitation, one would expect that the use of other browser revisitation mechanisms would be significantly lower among people who use tabs (or multiple windows) frequently. Indeed, Weinreich et al. [17] found that this was true: for participants who used tabs or multiple windows frequently, the back button accounted for about 10% of navigation actions, compared to 14% for their entire study population.

In our study, the back button accounted for a median of 7.1% of navigation actions. Among the group of tab power users, the median was 5.8%. The participant with the lowest tab creation rate, P12, was also the highest user of the back button. This data supports the hypothesis put forth by Weinreich et al. that increased use of multiple windows and tabs results in decreased use of the back button.

Measuring Tab Revisitation

One of the events logged by our extension was the *tab select* event. At any time, there is exactly one selected tab in every open window. For any given tab, it can be selected zero or more times over its lifetime. A tab can be selected zero times if it is opened in the background (e.g. by Ctrl-clicking on a link) and either the tab, the window, or the browser is closed without the user having selected the tab. Figure 6 shows a histogram showing how often tabs are selected over the entire lifetime of the tab.

Across all of our participants, we found that 45% of tabs were selected exactly once. As we saw in Figure 3, the largest portion of navigation events during the study happened with only one tab open, and Figure 6 showed that 7 participants had a median of only one or two tabs open. From this, we can posit that it's common to have a window that never has more than one tab in it. This case might account for many of the tabs that were selected exactly once. Another possibility is that it's related to one kind of multitasking described in the previous section—a short-term sub-task or “side task.” Many tabs may be opened with a particular purpose and closed as soon as the task is complete. Surprisingly, 25% of tabs were *never* selected. That is, they were opened but not switched to before the tab, or the entire window, was closed. The remaining 30% of tabs were switched to multiple times. So, more than two thirds of tabs are used once or less.

However, if we examine the number of tab-switch revisitations compared to the total number of tab switches, it paints a different picture. A median of 73.3% of tab switches were revisitations. This varied between participants, from 45% to 91%, but the aggregate total across all participants was 77.7%, quite close to the median.

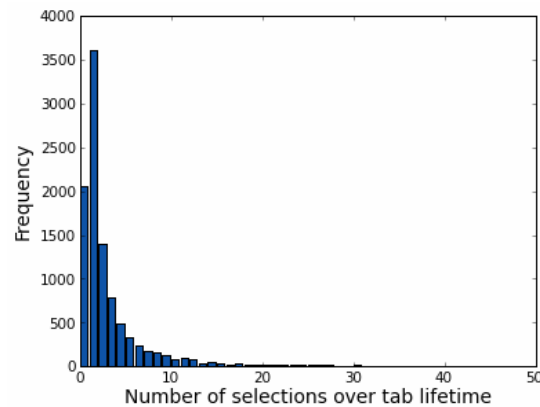


Figure 6: Frequency of tab selection

This shows that a large percentage (and for all but two participants, a majority) of tab switches are revisitations. But how does it compare to conventional revisitation? Figure 7 shows the rate of tab revisitation compared to conventional revisitation (i.e. page reloads).

Participants are sorted from lowest tab revisitation rate (on the left) to the highest. On the right side, we see 5 participants who have a tab revisitation rate higher than the rate of conventional revisitation. That is, if they were viewing a page previously viewed at any point during the study, it was more likely that they switched to a tab containing that page, rather than reloading the page in some way. Two of these participants actually had more tab-switch revisitations than total navigation actions (excluding session restore).

5 more participants had a tab revisitation rate that was at least 75% of the conventional revisitation rate, while the remaining 11 participants had a significantly higher rate of conventional revisitation than tab revisitation.

One more interesting thing to note is that for 17 of the 21 participants, tab switching was more frequent than back navigation. This is especially significant because previous studies have consistently shown that the back button (and its shortcuts) is the second most frequently used navigation method, behind clicking on hyperlinks.

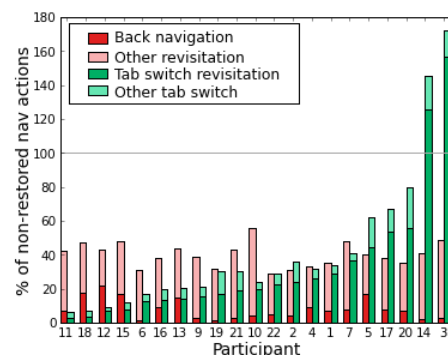


Figure 7: Tab switch and conventional (page reload) revisitation

Tabs compared to other revisitation mechanisms

We have shown that tabs are often used as a kind of revisitation mechanism. But web browsers already provide three special mechanisms for revisitation: the back button, the history, and bookmarks. Why would people instead choose to use tabs for revisitation?

The major difference between using tabs for revisitation is that switching to a tab does not cause the page to be reloaded again, whereas using the back button, history, or a bookmark does. Reloading the page might be desirable, because it allows you to see any information that has changed. However, reloading the page might also mean that you lose some of the state on the page. Six participants mentioned keeping a page open in a tab when they wanted it to stay “as-is”, and three people said that they liked the fact that when they switched back to a tab, they could return to the exact place in the page as when they left. This is not true of bookmarks or the history, and although the back button in Firefox does restore the position, it does not do so until after the page has reloaded, and is not as reliable.

7 participants said that tabs are quicker or more efficient for returning to a previously-viewed page than using the back button. 6 said that it is generally better, easier, or more convenient to use tabs instead of the back button.

“I don’t have to remember exactly where I was. I can just go back to the tab instead of [pressing] the back button [...]. It is right there.” (P12)

4 participants said they liked that tabs would keep their spot on the page. 6 people said that the back button can require too many or an unknown number of clicks to get back to a desired page, whereas returning to a tab is more predictable:

“You might have to click the back button like six times to get back to where you were. Or maybe never even find it again [...] whereas the tab just kind of stays at the originating page.” (P11)

4 people mentioned the visual or spatial aspect of tabs being an advantage over the back button:

“[Using tabs] really gives me the big picture of where everything is at. So spatially, kind of, my head knows which ones I’ve clicked on and which ones I’ve read.” (P10)

As previously mentioned, several people reported using a tab like a short-term bookmark. No one mentioned any specific advantages that tabs had for that purpose, but several reported that they do not use bookmarks very often.

These reports correspond to the relatively low use of bookmarks reported in previous studies. Those studies have shown the use of the browser history to be even lower than bookmarks, usually accounting for a fraction of a percent of all navigation actions. None of our participants claimed to make regular use of the browser history, so they were unable to make any comparisons with their usage of tabs.

IMPLICATIONS FOR DESIGN

Our participants’ comments point to several aspects of the browser user interface that are perceived to be lacking. While tabbed browsing provides a possible solution, an alternative is to address the core problem. Our study suggests three main areas of the browser user interface that could be improved.

Advantages of Using Windows

Although our participants generally preferred to use multiple tabs, many of them reported reasons for using multiple windows. The two most common reasons for using multiple windows were for viewing multiple pages at the same time, and for creating a higher-level task grouping. This suggests that it could be helpful if web browsers supported these features with multiple tabs. For example, a simple mechanism for selecting two tabs and allowing their contents viewed side-by-side in the browser window could be very useful. Regarding the organizational benefits of using multiple windows, several proposals have been made for “web workspaces” (e.g. [5, 11]) that could help organize browsing tasks. Perhaps these ideas could be revisited within the context of tabbed browsing.

Support for Common Tab Usage Patterns

We also discussed several common usage patterns that our participants reported in the interviews. It could be useful for future web browsers to include explicit support for these usage patterns, e.g. for using tabs as reminders or short-term bookmarks. Another common use of tabs was to open links in the background. It might be helpful for the web browser to show the relationships between tabs, and distinguish new, unread tabs from others.

Better Revisitation Mechanisms

Our results also illustrate a strong relationship between tabs and web page revisitation. We identified different ways in which tabs serve similar purposes as the conventional revisitation mechanisms of the browser, and reported what our participants cited as advantages of using tab-based revisitation instead of the back button. This indicates a need to improve the revisitation mechanisms in browsers. Indeed, studies have pointed out the problems with the stack-based behaviour of the back button, and have suggested alternatives [9, 16]. There have been attempts to design new and improved revisitation mechanisms for the web browser, from simplified back/forward behaviour [9, 13] to systems which combine the back button, browser history, and bookmarks in a single interface [8].

One aspect of the revisitation problem that previous studies have not addressed is the problem of maintaining the current page state. In our study, we found that, even if users were confident that they could easily return to the current page (e.g. by using the back button), they still expressed concern that they would lose important state. Future browsers might alleviate this problem by modifying the back button to work like tabs—i.e., when the user returns to the page, it is in *exactly* the same state as they left it, including the position of the scroll bar, etc.

STUDY LIMITATIONS

The study population was limited to Firefox users who often used multiple windows or tabs during browsing. Our choice of *active* tab and window users was intentional, in order to learn the most about how and why these people use tabs. Given these constraints, the population is probably skewed towards intermediate to advanced users, which limits the generalizability of our results. However, although Internet Explorer is the most popular browser, it did not support tabbed browsing by default until version 7.0 (the most recent at the time of the study), whereas all versions of Firefox support tabbed browsing, so Firefox users likely comprise a significant portion of people who use tabs often.

Unlike many other studies of this nature, we conducted the study without seeing the actual sites and pages that the participants visited. We felt that this privacy protection was important to obtain a true representative sample of people's browsing habits. However, it is still possible that the participants modified their habits simply because they knew they were being observed. As with any diary study, in order to obtain the most relevant data, it was necessary to inform the participants as to what aspects we were most interested in. Asking them to think about and describe their own behaviour may have caused them to change it. However, by conducting the study for two weeks, we hoped that these kinds of observational effects would occur early and diminish over the course of the study.

CONCLUSIONS AND FUTURE WORK

In this study, we gained insight into how Firefox users use multiple tabs and windows during web browsing. Through interviews, diary entries, and analysis of click-stream logs, we have learned not only what they do, but why they do it.

We demonstrated our participants' preference for the use of multiple tabs rather than multiple windows, and discussed the reasons they cited for this preference, as well as why they still use multiple windows sometimes. We presented several quantitative measures of our participants' tab usage, and discussed the reasons they cited for using multiple tabs. We identified a subset of participants who used tabs much more frequently than the others, and who reported that they habitually open links in new tabs. Finally, we discussed the relationship between tab usage and web page revisitation. We found a far lower rate of back events than previous studies (e.g. [14]), which may be related to the use of tabs.

As web applications become more popular, the browser is quickly becoming "the other desktop." We are interested in pursuing research into higher-level task management behaviour in the browser, and how the task management mechanisms in the browser can be better integrated into the mechanisms provided by the operating system.

REFERENCES

1. Mozilla's Firefox global usage share is still growing according to onestat.com.
http://www.onestat.com/html/aboutus_pressbox57-firefox-mozilla-ie-browser-market-share.html.
2. E. Adar, J. Teevan, & S. Dumais. (2008). Large scale analysis of web revisitation patterns. *CHI*. 1197–1206.
3. Browser market share.
<http://marketshare.hitslink.com/report.aspx?qprid=0&qptmeiframe=M&qqsp=119>. Retrieved April 18, 2009.
4. M. Byrne, B. John, N. Wehrle, & D. Crow. (1999). The tangled web we wove: a taskonomy of www use. *CHI*. 544–551.
5. S. Card, G. Robertson, & W. York. (1996). The webbook and the web forager: an information workspace for the world-wide web. *CHI*. 111–119.
6. L. Catledge & J. Pitkow. (1995). Characterizing browsing strategies in the world-wide web. *Comput. Netw. ISDN Syst.*, 27(6). 1065–1073.
7. C. Choo, B. Detlor, & D. Turnbull. (2000). Information seeking on the web: An integrated model of browsing and searching. *First Monday*, 5(2).
http://www.firstmonday.org/ISSUES/issue5_2/choo/index.html.
8. A. Cockburn, S. Greenberg, S. Jones, B. McKenzie, & M. Moyle. (2003). Improving web page revisitation: Analysis, design and evaluation. *Information Technology and Society*, 1(3). 151–183.
9. A. Cockburn, B. McKenzie, & M. Smith. (2002). Pushing back: evaluating a new behaviour for the back and forward buttons in web browsers. *Int. J. Human-Computer Studies*, 57(5). 397–414.
10. A. Dotzler. (2005). Tabbed browsing history.
<http://weblogs.mozillazine.org/asa/archives/008433.html>. Retrieved April 2, 2009.
11. N. Jhaveri & K. R  ih  . (2005). The advantages of a cross-session web workspace. *CHI*. 1949–1952.
12. M. Kellar, C. Watters, & M. Shepherd. (2006). The impact of task on the usage of web browser navigation mechanisms. *Graphics Interface*. 235–242.
13. N. Milic-Frayling, R. Jones, K. Rodden, G. Smyth, A. Blackwell, & R. Sommerer. (2004). Smartback: supporting users in back navigation. *WWW Conference*. 63–71.
14. H. Obendorf, H. Weinreich, E. Herder, & M. Mayer. (2007). Web page revisitation revisited: implications of a long-term click-stream study of browser usage. *CHI*. 597–606.
15. A. Sellen, R. Murphy, & K. Shaw. (2002). How knowledge workers use the web. *CHI*. 227–234.
16. L. Tauscher & S. Greenberg. (1997). How people revisit web pages: empirical findings and implications for the design of history systems. *Int. J. Human-Computer Studies*, 47(1). 97–137.
17. H. Weinreich, H. Obendorf, E. Herder, & M. Mayer. (2006). Off the beaten tracks: exploring three aspects of web navigation. *WWW Conference*. 133–142.
18. H. Weinreich, H. Obendorf, E. Herder, & M. Mayer. (2008). Not quite the average: An empirical study of web use. *ACM Trans. Web*, 2(1). 1–31.