Co-located Collaborative Web Search: Understanding Status Quo Practices

Saleema Amershi
University of Washington
Seattle, WA, USA
samershi@cs.washington.edu

Meredith Ringel Morris
Microsoft Research
Redmond, WA, USA
merrie@microsoft.com

CHI 2009, April 4–9, 2009, Boston, Massachusetts, USA.
ACM 978-1-60558-247-4/09/04.

Abstract
Co-located collaborative Web search is a surprisingly common activity, despite the fact that Web browsers and search engines are not designed to support collaboration. We report the findings of two studies (a diary study and an observational study) that provide insights regarding the frequency of co-located collaborative searching, the strategies participants use, and the pros and cons of these strategies. We then articulate design implications for next-generation tools that could enhance the experience of co-located collaborative search.

ACM Classification Keywords
H5.3 [Information interfaces and presentation]: Group and Organization Interfaces - Computer-supported cooperative work.

General Terms
Design, Experimentation

Keywords
Collaborative search, Web search, search interfaces.

Introduction
There are many situations in which people want to search the Web together. Students often want to

collaboratively search the Web to complete homework assignments for the pedagogical value that a shared context and face-to-face communication provides [1]. Friends and families want to collaboratively search for the social experience of planning activities together [3]. In the workplace, colleagues want to collaboratively search the Web to conduct joint research [3]. Unfortunately, Web browsers and search engines are designed for single users, working alone.

Here, we present a diary study and in-lab observations that we conducted to learn about status-quo co-located collaborative search practices. Our findings from these studies suggest implications for designing next-generation Web search tools that can support co-located collaboration.

Investigations of Status-Quo Practices

People use a variety of strategies to collaboratively search the Web with mainstream technologies. In order to better understand these strategies, we investigated current co-located collaborative search practices of people at work and at home through a diary study. From this, we identified two of the most common co-located collaborative search methods, shared- and parallel-computer search. In the shared scenario, a group of people gathers around a single computer to conduct a search. In the parallel scenario, groups search using individual computers side-by-side. We then investigated the shared and parallel search strategies further through an observational study.

Diary Study

We conducted a diary study involving twenty Microsoft employees (seven female), wherein participants kept track of their co-located collaborative searches at work and at home for one week. Our participants ranged in age from 19 to 50 years old and had a variety of jobs, including both technical (*e.g.,* software developer) and non-technical (*e.g.,* administrative assistant) roles. During the study period, each participant sent us daily, form-based records about any co-located collaborative searches they performed at work and at home. For each search, they noted the search topic, the search type [2] (i.e., *navigational*, where the target information exists at a particular, known website; *transactional*, where the goal is to navigate to a particular website to carry out some activity; or *informational*, where information may span multiple websites), whether the search was planned or spontaneous, the number and relationship of the collaborators, the duration (i.e., *a few minutes*, *an hour,* or *several hours*), and the resources used.

In total, our participants engaged in 38 co-located collaborative Web searches within the one-week study period, with 90% reporting at least one and up to five such searches. 55.3% of these occurred at work and 44.7% at home.

Of the co-located collaborative searches that took place at work, most (52.4%) were *informational* (*e.g.*, searches for work-related documentation such as network protocols or design principles) or *navigational* searches (38.1%) (*e.g.,* navigating to the company's website). A few (9.5%) were *transactional* (*e.g.,* applying for a parking permit). The majority were spontaneous (61.9%) and lasted only *a few minutes* (85.7%). Searches that lasted for *an hour* (14.3%) were all *informational* in nature. Searches mostly occurred in pairs (85.7%) or small groups of three or four colleagues or friends (9.5%).

Despite the abundance of computing resources available to employees of this large technology company, most groups (76.2%) used a single workstation to collaboratively search the Web at work. On three occasions (14.3% of the at-work searches), participants reported searching on separate, side-by-side workstations with their collaborators.

Of the co-located collaborative searches that occurred at home, most (58.8%) were *informational* in nature such as planning weekend activities or trips. Some were also *transactional* (29.4%) (*e.g.,* buying theatre tickets) or *navigational* (11.8%) (*e.g.*, navigating to a transit website). Again, the majority of collaborative searches that occurred at home were spontaneous (70.6%) and lasted only *a few minutes* (64.7%); however, several also lasted for *an hour* or *several hours* (35.3%). The searches usually occurred in pairs (70.6%) or groups of three or four family members or friends (29.4%). Occasionally, groups searched using multiple devices side-by-side (two computers, or one computer plus internet-enabled mobile phones) (17.6%). However, most searches took place around a single, shared workstation (76.5%).

Observations of Status-Quo Practices

Based on our diary study of current co-located collaborative search practices in the home and at work, it appears that the two most common strategies that people use to perform co-located collaborative Web search are shared- and parallel-computer searching. This is consistent with previous research on the co-located collaborative search strategies that people use in libraries and schools [1]. To better understand the benefits and limitations of these strategies, we conducted a study in which two researchers observed 12 three-person groups (with pre-existing relationships) collaboratively searching the Web using each strategy. None of these participants had participated in our diary study. Each group performed two tasks per strategy (with a maximum of seven minutes per task), one that was pre-defined and that reflected a typical informational search (*e.g.*, determine the U.S. state with the lowest driving age) and one that was self-defined (*e.g.*, planning a family vacation). This study is discussed in greater detail in [1], where we report on findings that contrast the status-quo conditions to use of the CoSearch system; here, we report additional, qualitative observations and anecdotes that focus on the pros and cons of the two status-quo search scenarios.

SHARED-COMPUTER SEARCH

In general, both observers noted high levels of communication between participants sharing a single computer. In some cases, the communication amongst groups appeared to be effective. For example, in some groups, the person controlling the mouse and keyboard of the shared computer would democratically lead the search by asking group members for input (*e.g.,* "any suggestions on what to search for?"). This strategy appeared to be most effective when the group members had similar personalities (*e.g.,* vocal or quiet) or technical skills. However, even in these cases, arguments occurred and suggestions went ignored. For example, in one instance where a group was shopping for purses, one group member suggested to go to the 'Louis Vuitton' website, while another suggested searching for 'handbags.' The third group member attempted to mediate by saying "which one will it be?" and then elected to try the 'handbags' query. After some time looking through the results of this query, the

third group member asked the group "is there another one you want to look at?" apparently forgetting the alternative suggestion of 'Louis Vuitton' which had to be requested again by the original group member.

When group members had varying personalities or search skills, their communication was less effective and often resulted in frustrations. For example, in several instances, the person controlling the shared computer's input devices would enter queries, navigate to pages, or scroll without considering the verbalized suggestions of the other group members. While quieter group members usually remained silent after being ignored, more vocal members became more aggressive, sometimes demanding that their suggestions be executed and sometimes using other tactics such as pointing at a page or reading the page aloud in efforts to prevent another group member from scrolling or navigating away from that page. Several participants commented that this lack of control and freedom was frustrating (*e.g.*, "one person controls, others watch"). Participants also commented on the inefficiency of this search strategy (*e.g.,* "only one search query available at a time").

During the shared-computer searches, we also observed benefits of focused collaboration and shared awareness. For example, when one particular group was attempting to search for an answer to a question we posed as a search task, the person controlling the input devices navigated to a blog. Another group member then commented on the credibility of the site, stating that "this is just blogs, so it's not a fact" resulting in the group navigating to a different page. Participants also commented positively on the shared context achieved from the shared-computer searches (*e.g.,* "Everyone focusing on one thing" and "I knew what everyone was doing").

PARALLEL-COMPUTER SEARCH

Both researchers observed a decline in communication levels in all groups when they used the parallel-computer strategy to collaboratively search the Web. In several cases, all members of a group would begin searching on their own computers without communicating with each other at all. In others, one member of a group would attempt to initiate the search direction by suggesting a keyword or topic, but this would usually be followed by silence and individual searching. And although participants commented that they liked the freedom and distribution of control that the parallel computer search strategy enabled (*e.g.*, "could go off on tangents to gather info" and "I could look at my own content at my own pace"), this lack of communication and awareness amongst the groups often resulted in redundant work. For example, in at least five of the twelve groups, the researchers observed that two or all three of the group members would unknowingly navigate to the same web page, sometimes even at the same time. In one group, all the participants announced the queries they were going to try before starting the search in an attempt to avoid redundancies. As their search continued, they occasionally updated the group with new queries or sites they were visiting, but redundant queries and web pages were still observed.

When communication did occur, it mostly occurred pairwise between group members seated next to each other (rather than between group members on the outside computers) or between the more vocal or technically-experienced subset of each group. In cases

where communication occurred across the group, it was often only to communicate information about results of a search (*e.g.*, expressing opinions of the utility of a page, such as "well that's bad" or "that doesn't help") rather than the process of the search (which would have helped to reduce redundancies). Sharing information about the process of a search appeared to be difficult and frustrating for many of the groups. For example, some participants would turn their monitors towards the others in order to show part of a webpage, or read aloud from a page (often with no response from others). In other cases, participants tried to navigate to the same page as other group members so as to have their own view of the page of interest. This would require directions from others (e.g., queries to execute or descriptions of links followed) which were not always successfully followed. For example, in one group, a participant indicated that she wanted to look at the same site as another by stating "I don't know what site you're looking at." The other group member briefly looked over and attempted to direct her by stating "it's the site of the X hotel" and then went back to his search. She ended up at a different page altogether.

Design Implications

The findings from our diary study and observations suggest implications for designing useful and usable systems to support co-located collaborative search:

Start-up costs should be proportional to the duration of the search. Our diary study showed that, while some collaborative searches (particularly on informational tasks) lasted an hour to several hours, most lasted only a few minutes. Many proposed collaborative search systems, such as CoSearch [1] and SearchTogether [4], have high start-up costs (*e.g.*, connecting mobile phones over Bluetooth to the shared computer or logging into a special service). While such systems may be appropriate for search in specialized domains (*i.e.,* in schools, where students conduct informational searches to prepare in-depth reports [1]), developing lower-start-up-time techniques for enhancing collaborative search experiences is important for facilitating brief-duration collaborations.

Maintain a queue or history of suggestions. For any given search, collaborators may come up with several ideas about queries to try or web pages to visit. Furthermore, searches often involve following different leads and so maintaining a history of which leads have been followed (and by whom) can reduce undesired redundancy and preserve group members' ideas for future avenues of investigation.

Enable distribution of control. Providing a means for shared-computer searchers to interact simultaneously (through auxiliary devices like individual mice [5] or mobile phones [1]) could reduce the frustrations we observed that were caused by the lack of control over the input devices of a shared computer.

Include consensus facilities. In a shared-computer setting, in conjunction with a queue/history approach, group members must have a way to specify which queued item (query term, web page, etc.) should be the group's current focus. The use of voting mechanisms or a designated master input device can help ensure a smooth workflow, particularly when some group members may be overly domineering. Periodic switching of the master device's identity can help ensure participation by all group members in the search and decision-making process.

Provide group awareness mechanisms. For shared-computer scenarios, providing awareness of each group member's contributions (*i.e.,* keywords, links, notes) could be beneficial for ensuring that users not in control of the main mouse and keyboard (but perhaps contributing through the use of auxiliary devices) can feel that they are actively participating. For parallel-computer searches, providing awareness of which leads other group members have already followed could be valuable in reducing undesired redundancy.

Provide a shared context. While people liked the freedom of parallel searching, it was often difficult to share information and to know what collaborators were doing. Enhancing the parallel-computer search with a shared and individual view may be useful to remedy this, and also for enabling less experienced searchers to learn about more effective search strategies from their more experienced counterparts.

Make it easy to take away relevant information. A search is often part of a larger extended task (*e.g.,* a work-related project) and so collaborators often need to revisit information that they found during a search. Even if a search is independent of any other task, people may want to revisit interesting information found during a shared search at a later time. In the common shared-computer scenario, only the computer's owner has continued access to the group's activity (*i.e.*, via Web browser history), but other group members could benefit from having a "take away" artifact of their collaboration. This can be achieved through the ability to save, email, or download information onto individual devices (such as individual computers or mobile phones). Even for short searches, this can be beneficial and has low costs. For parallel searching, a way to allow group members to merge their separate findings into a final, shared product would provide a way to "take away" the information discovered on another group member's computer.

Conclusions and Future Research

We have presented the results of two studies of co-located collaborative Web search. Our diary study contributes quantitative data regarding co-located search at work and in the home, such as frequency, duration, group configuration, and task types. Our observational study contributes qualitative data regarding pros and cons of the most common co-located search configurations: shared- and parallel-computer setups. From these investigations, we distilled several implications for designing co-located collaborative Web search systems. Future studies of different sized groups and of groups of people in different settings could help refine these implications and verify their applicability to co-located collaborative Web search in a variety of contexts.

References

[1] Amershi, S. & Morris, M.R. CoSearch: A System for Co-located Collaborative Web Search. In *Proc. CHI 2008*, 1647-1656.

[2] Broder, A. A Taxonomy of Web Search. *ACM SIGIR Forum 36*, 2 (2002), 3-10.

[3] Morris, M.R. A Survey of Collaborative Web Search Practices. In *Proc. CHI 2008,* 1657-1660.

[4] Morris, M.R. and Horvitz, E. SearchTogether: An Interface for Collaborative Web Search. In *Proc. UIST 2007*, 3-12.

[5] Pawar, U., Pal, J., & Toyoma, K. Multiple Mice for Computers in Education in Developing Countries. In *Proc. ICTD 2006*, 64-71.