SearchTogether: An Interface for Collaborative Web Search

Meredith Ringel Morris
Microsoft Research
One Microsoft Way
Redmond, WA, 98052, USA
merrie@microsoft.com

Eric Horvitz
Microsoft Research
One Microsoft Way
Redmond, WA, 98052, USA
horvitz@microsoft.com

ABSTRACT
Studies of search habits reveal that people engage in many search tasks involving collaboration with others, such as travel planning, organizing social events, or working on a homework assignment. However, current Web search tools are designed for a single user, working alone. We introduce SearchTogether, a prototype that enables groups of remote users to synchronously or asynchronously collaborate when searching the Web. We describe an example usage scenario, and discuss the ways SearchTogether facilitates collaboration by supporting *awareness*, *division of labor*, and *persistence*. We then discuss the findings of our evaluation of SearchTogether, analyzing which aspects of its design enabled successful collaboration among study participants.

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

General terms: Design**,** Human Factors

Keywords: Web search interfaces, collaborative search, persistent search, computer-supported cooperative work.

INTRODUCTION
Web search is generally considered a solitary activity; browsers and search engine homepages are designed to support single-user scenarios. However, studies of search strategies in educational settings [19, 32] and among knowledge workers [22] reveal that users often desire to collaborate on search tasks.

In this paper, we first discuss our need-finding research on people's current and desired practices surrounding collaborative Web search. We also discuss related work on studies of information retrieval habits, sensemaking, collaborative Web browsing and bookmarking, "passive" collaboration, and multi-user search. Then, we introduce SearchTogether, a prototype developed to address the findings about collaboration and search. SearchTogether enables direct collaboration among friends, family, and colleagues in information seeking and review. SearchTogether makes a contribution in extending current Web search tools by providing explicit support for small groups of people who know each other, enabling them to collaborate on both the process (*i.e.,* formulating queries, choosing results to explore) and products (*i.e.,* commenting on and rating found items, creating a shared summary) of a search. We provide a detailed overview of the SearchTogether system, including a sample usage scenario. We conclude by discussing our findings from an evaluation of SearchTogether.

Motivation
The design of SearchTogether was motivated by a survey we conducted that gathered information on the current Web search practices and needs of 204 knowledge workers [22]. The survey revealed that a majority of respondents wanted to collaborate with friends, relatives, and colleagues when searching the Web. 97.1% of respondents reported engaging in at least one of the collaborative search behaviors mentioned in the survey. For example, 30.4% of the 204 respondents reported that they had instant-messaged other people to coordinate real-time Web search, and 18.1% reported having explicitly divided up responsibilities for parts of a search task among several people, and then shared their results.

The survey respondents had employed a variety of strategies for joint searching in the absence of explicit collaborative search functionality in Web browsers and search engine sites. The most common methods reported for collaborating on search were emailing links back and forth, using instant messaging software to exchange links and query terms, and speaking with a collaborator on the phone while viewing a Web browser.

From this survey, we also learned about classes of informational goals that provoke attempts to search the Web in a collaborative manner. Planning travel, making expensive purchases, planning social events, researching medical conditions, and finding information related to a joint project or report were the most frequently mentioned tasks for multi-user searching. For instance, 30 of the 204 respondents (14.7%) said they had tried to collaborate with their co-travelers to find trip-planning information on the Web.

Based on survey respondents' descriptions of their current and desired collaborative search practices, we identified three key features for supporting collaborative Web search:

Figure 1. The SearchTogether client. (a) integrating messaging, (b) query awareness, (c) current results, (d) recommendation queue, (e)(f)(g) search buttons, (h) page-specific metadata, (i) toolbar, (j) browser

awareness, *division of labor*, and *persistence*, which we discuss in more detail in the following sections.

RELATED WORK

SearchTogether builds upon several areas of research, including studies of people's information retrieval habits, systems that support sensemaking, systems supporting multi-user Web browsing and bookmarking, "passive" collaboration systems, and systems supporting multi-user searching.

Studies of Information Retrieval Habits

Prior studies of users' information retrieval habits have informed the design of the SearchTogether system. Broder [4] and Rose and Levinson [27] developed taxonomies of Web search activities based on analyses of search logs. Both of these classification schemes identify *informational search* as a common search activity, *i.e.*, searches in which a user is seeking information rather than merely trying to locate a single target site. This class of searches, which often involves multiple refinements of query terms and can be quite exploratory [33] in nature, is the type of rich search task that would benefit from collaboration, and which SearchTogether is designed to support.

The value of providing a persistent representation of search is suggested by several studies, including a survey by Aula *et al.* [1] and a log study by Teevan *et al.* [31], which both found that people frequently find themselves re-entering previously entered queries in order to re-find useful webpages. Tauscher and Greenberg's study of Web browser logs [30] found that over half of page visits were re-visits. These results suggest that persistent representations of search activities and results could assist with the recovery of context and efficient re-access of useful pages. Jones *et al.*'s work on Keeping Found Things Found [16] reported that among the people they interviewed, "doing nothing" was the most popular strategy for recording useful information found online. These results highlight the value of providing searchers with a means for capturing activity in an implicit manner and making this information available later for solo or collaborative uses. SearchTogether's persistence features were motivated by these results on retrieval habits.

Studies of students in grade school [19] and in university [32] have observed that joint information retrieval occurs in such settings, as students often work together on team projects and assist each other with strategies for locating relevant information online or in library databases. Our own study of knowledge workers' search habits and needs [22] informed SearchTogether's design most directly, by providing data about people's current practices and needs regarding collaborative Web searching.

Sensemaking

Sensemaking [28] (*i.e.*, processing, organizing, and analyzing information) is intricately related to informational and

exploratory searches. SearchTogether supports sensemaking tasks by providing summary views of information, the ability to rate and comment on online content, and the ability to discuss content via messaging with group members. Several other systems seek to integrate the processes of search and sensemaking.

For example, InkSeine [13] is a Tablet PC search application that allows users to store a pointer to a search via a *breadcrumb* object intermixed with their handwritten notes. Dontcheva *et al.*'s system for summarizing personal Web browsing sessions [8] allows users to define patterns for extracting structured information from a set of Web pages. SearchPad [3] allows a user to explicitly flag a webpage for inclusion in a workspace in order to help him maintain context during complex search tasks. Hunter Gatherer [29] and Google's Notebook application[1] allow users to collect snippets of content from several webpages and combine them in a single document.

In his visionary article, "As We May Think," [5], Vannevar Bush described a futuristic information system that would allow users to follow "trails" of information. SearchTogether is intended to help users follow not only their own trails through an information-finding and analysis task, but also the trails of their collaborators.

Multi-User Web Browsing and Bookmarking

Several systems have explored interfaces that allow multiple users to collaboratively interact with online information, though not search per se. The Sociable Web [7], for example, allows a user to know that others were currently viewing the same webpage, and to communicate with those people. Several systems allow users to share bookmarks or favorites lists, such as WebTagger [17], Wittenburg *et al.*'s system [34], or the commercial site del.icio.us[2].

Collaborative Web browsers support synchronous remote collaboration, usually by providing "yoked" views, where one user's navigation causes other group members' browsers to navigate to the same page. This yoking can have a master/slave aspect, where only one group member's browser has the capability to lead the others'. Cabri *et al.*'s group browsing system [6] divided the browser into two frames, with one frame displaying the jointly-viewed Webpage, and the other showing the history of all pages visited by the group and displaying integrated chat. The W4 browser [10] provided follow-me navigation and scrolling, integrated chat, and remote telepointers. GroupWeb [11] also provided yoked browsing and telepointers, as well as allowing users to associate comments with jointly-viewed Webpages.

The WebSplitter system [12] and Maekawa *et al.*'s system [21] take a different approach to joint Web browsing, by providing a framework for dividing a single webpage into several portions and serving each portion of the page to a different member of the group in order to facilitate parallelized visual search of a single webpage.

Unlike these systems for shared bookmarking or browsing, SearchTogether focuses on supporting collaboration during the process of searching the Web, including formulating queries, exploring search results, and evaluating the information that has been found.

"Passive" Collaborative Search Systems

Most research in the area of "collaborative" search focuses on passive forms of collaboration (*i.e.*, using data generated by large numbers of users' interactions with a system to automatically refine system behavior). Examples from this line of work include systems that use query logs and clickthrough data to generate query substitutions or recommendations, such as [2] and [15]. VisSearch [20] uses data mining algorithms to uncover patterns in users' queries and subsequent browsing in order to generate recommendations for users with similar queries. Yoople[3] is a search engine that solicits explicit user feedback on result rankings, and uses the feedback from a large number of users to change the rankings seen by others. Implicit user feedback, such as clickthrough data from a large number of users, can also be used as a collaborative filtering mechanism to re-rank search results [25]. Unlike these systems, SearchTogether focuses on active collaboration amongst a small group of users who know one another and are working together toward a shared goal.

Multi-User Search

There have been some previous research efforts toward enabling collaboration when searching. However, most of these prior systems are designed for specialized domains or devices, rather than for general-purpose Web search. For example, TeamSearch [24] allows co-located groups of up to four people seated around a tabletop display to search through a database of images using a visual query language. C-TORI [14] is a multi-user system for querying a relational database. One member of the group is designated the coordinator, and the coordinator has the power to put other group members into a tightly-coupled mode, in which the queries made by one user are visible to the other user. MUSE [18] is a system that supports synchronous, remote collaboration between two people searching a medical database. MUSE users perform standard single-user searches, but have built-in chat and the ability to press a button that shares metadata about the current database results with the other user. CIRE [26] is a system targeted toward multi-user Web search; users perform only standard, single-user searches, but can add comments to pages they find. These comments are then visible to other members of the group who visit these same pages later. A group history can be viewed, which lists the pages visited by the entire group. S^3 [23] allows users to asynchronously share useful sites found during a Web search by representing search results in a persistent file format that can be sent to and augmented by several people.

[1] http://www.google.com/googlenotebook/overview.html

[2] http://del.icio.us/

[3] http://www.yoople.net/

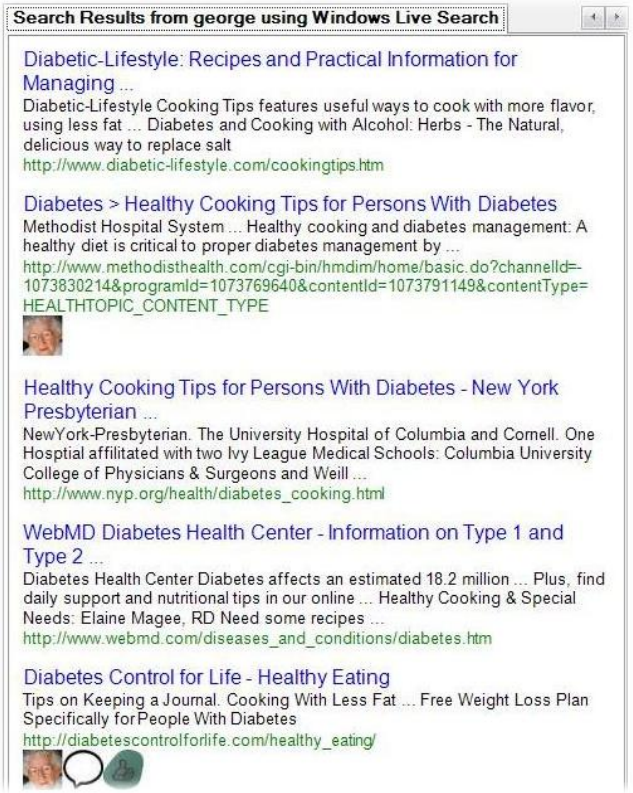

Figure 2. Split search results appear in a separate tab within the recipient's client.

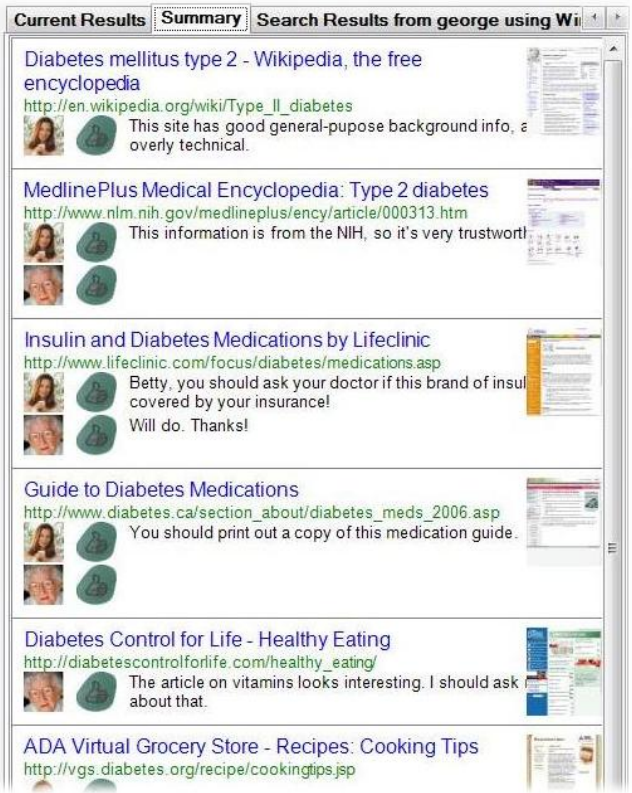

Figure 3. Automatically-generated summaries provide a quick way to access the pages group members identified as relevant.

There are also a few commercial systems that enable limited forms of collaborative Web search. The ChaCha search engine[4] pairs searchers up with a live guide, who chats with them and suggests urls to visit. This is not truly intended as a collaborative search experience, but rather as a way to use human labor to simulate the ability to process natural-language-style search queries. Windows Live Messenger[5] is an instant messaging client that contains a "search" button. If someone types a query into the IM box and hits "search", hyperlinks to the top three URLs returned by Windows Live Search in response to that query are shown as part of the chat transcript for both chat participants.

In contrast to these systems, SearchTogether's design supports two or more synchronous or asynchronous searchers using a variety of Web-based search services, and provides integrated support for awareness, division of labor, and persistence.

SEARCHTOGETHER

SearchTogether is designed to enable either synchronous or asynchronous remote collaboration (*e.g.*, each participant is in a distinct location, with his own computer). SearchTogether employs a client/server architecture. The server acts as an intermediary for sending shared state among clients, as well as being a repository for storing SearchTogether session data in order to enable session persistence.

In the following sub-sections, we illustrate the experience of using SearchTogether by describing a sample usage scenario. We then describe how the features of SearchTogether support our design goals of supporting *awareness*, *division of labor*, and *persistence*.

Example Usage Scenario

Rachel and her husband George, who reside in New York, learn that George's mother Betty, who lives in California, has recently been diagnosed with diabetes. Betty wants to learn more about the causes of and treatments for her condition, but she is not a particularly skilled Web searcher. Her son and daughter-in-law are also interested in finding out more about Betty's condition and helping her to learn about the illness. They decide to use SearchTogether to help them communicate and share their findings about diabetes.

Rachel logs onto SearchTogether and selects "Create New Session," entering "diabetes" as the session topic. Next, she is prompted to specify other users who will share the session with her. She chooses George and Betty's IDs from her buddy list. The session is created, and, for now, Rachel is the only member of the group logged in. She performs some standard searches for "type II diabetes" and "diabetes medications." Her search for medications brings up information on SuperMed, a new treatment shown to be especially effective for patients who have a combination of diabetes and high blood pressure, an ailment that also affects Betty. Rachel clicks the "Recommend" button while view-

[4] http://www.chacha.com/

[5] http://get.live.com/messenger/overview

ing the page about SuperMed, selecting Betty as the recipient and entering the comment, "Betty, you should ask your doctor if this medication is appropriate for you! You may want to check if your insurance will cover it."

Later that evening, Betty logs into SearchTogether. On the login page, she sees a list of all of the joint search sessions that she is a part of, and she chooses "diabetes" from that list. Upon logging in, she can see from the Query Awareness area that Rachel has already searched for information on "type II diabetes" and "diabetes medications." Betty sees a webpage thumbnail in her recommendations queue and clicks on it, which loads the SuperMed webpage into the browser, displaying Rachel's comments above it. Betty prints the page to bring to her doctor's office. An instant message from George ("Hi, Mom!") appears in the SearchTogether chat window, heralding that George has just logged into the session.

Betty sends George an instant message: "My doctor told me that I should eat a low-sugar diet. Can you help me find recipes?" George enters the query "diabetes cooking tips" into his SearchTogether client, and chooses the Split Search option. George's half of the search results appear in his current results tab, and a tab labeled "Live.com search results from George" appears in Betty's client, containing half of the results from George's query. Betty clicks on the tab, and begins exploring the search results.

Meanwhile, George also views several pages, giving thumbs-up ratings to several low-sugar recipes he thinks his mother will enjoy. He can see from the results list that one of the pages also came up in response to one of Rachel's earlier queries, and that she has given it a poor rating, so he skips that one.

Several months later, Betty's neighbor, Angela, is diagnosed with the same form of diabetes. Angela mentions to Betty that she is having trouble preparing healthy meals that follow her doctor's guidelines. Betty logs onto SearchTogether, finds the "diabetes" session in her list of joint searches, and opens it up. She adds Angela as a member of the session. Now, Angela can log on and view the session's summary tab, which contains a list of all the pages that Betty, George, and Rachel gave positive ratings to when they conducted their joint search.

Design Goals

SearchTogether's design was guided by the findings of our survey on Web search habits and needs [22]. Based on the current and desired collaborative search practices of the survey respondents, we hypothesized that an effective interface for remote, collaborative Web search should support awareness, division of labor, and persistence. In this section, we explain how SearchTogether satisfies each of these criteria.

Awareness

Facilitating awareness of group members' search processes and findings can enable lightweight collaboration by reducing overhead involved in explicitly asking other group members to provide this information. Awareness of other group members' activities also has the potential to reduce undesired duplication of effort, and to promote learning of search techniques through awareness of other group members' search strategies.

Per-user query histories are one of SearchTogether's awareness mechanisms. Each group member's screen name and photo is shown in the "Query Awareness" region of the client (Figure 1b); these items are rendered in gray if that user is offline. Each time a user executes a search, the query terms used are added to a list underneath that user's photo; this history is synchronized across all group members' clients. This history is also interactive – any user can click on any of the query terms in order to view the results it produced. In order to preserve shared context for all group members, clicking a term from the Query Awareness region does not re-execute the query, but instead shows the results from the query's initial execution, which were stored on the SearchTogether server.

We provide query awareness information for several reasons. One frustration of users attempting to collaborate with current search tools is unnecessary duplication of effort (*i.e.*, when multiple people unknowingly enter the same queries, rather than exploring different sub-parts of a topic) [22]. One aim of providing query awareness information is to facilitate awareness of different aspects of a shared search task being pursued by other group members, in order to avoid this unnecessary duplication of effort. Prior studies of search behavior [1] [31] report that users often reissue the same queries multiple times in order to re-find previously visited pages. The persistence of the query history associated with a search topic in SearchTogether also aims to reduce the wasted effort of query repetition by providing users with a shortcut to revisit prior results.

As formulating Boolean-style queries is notoriously difficult for non-technical users [9], query awareness information may also be a valuable learning mechanism. Less experienced searchers can view the syntax used by their teammates, and perhaps be inspired to reformulate their queries using different query terms or advanced syntax that they observe others using.

Page-specific metadata is another SearchTogether feature for supporting awareness among collaborators. SearchTogether associates three types of metadata with each webpage that users visit: visitation information, ratings, and comments.

Whenever a user views a webpage in SearchTogether's browser (Figure 1j), the date, time, and identity of the viewing user are recorded and associated with that page. People viewing pages in the browser can see if any other group members have already visited that same page, since their photos are shown in the region immediately above the browser window (Figure 1h). Hovering over these photos with a mouse presents a tooltip that provides the detailed history of all the dates and times the corresponding user has visited the page. This visitation information is also revealed in search result lists – if a user executes a query whose results include pages already viewed by other group members, then those users' photos are shown below the snippet

for those pages (Figure 1c). This visitation awareness information can help a user choose to avoid a page already visited by other group members in order to spread out search efforts, or perhaps to choose to specifically visit such pages, as they may signal promising leads, as indicated by the presence of comments and/or ratings.

Users can associate ratings ("thumbs up" for a positive rating and "thumbs down" for a negative one) and/or comments with a webpage by pressing buttons on the toolbar above the browser window (Figure 1i). These ratings and comments are visible to other group members visiting the same page; they are shown next to the corresponding user's photo in the region directly above the browser. Ratings and comments are also depicted below the snippets for search results that point to the corresponding pages.

Division of Labor

Many of our survey respondents described employing complex methods for managing joint search tasks in order to avoid unnecessary duplication of effort. Such methods included explicitly dividing up the space of potential keywords, search engines, or sub-tasks, and assigning a portion to each group member. To address this issue, SearchTogether includes several mechanisms for managing division of labor among collaborators.

Instant messaging is integrated into the SearchTogether client (Figure 1a), to provide a way for group members to discuss the current task and coordinate their efforts. Our survey found that people often open an IM client and a Web browser at the same time as a way to allow them to communicate with others about shared search tasks. By integrating instant messaging with our search client, we aim to reduce the cognitive overhead of application-switching between a browser and an IM application. Additionally, by making messaging a first-class citizen in SearchTogether, we are able to capture and store all IM conversations so that they are available for perusal by asynchronous collaborators or by users revisiting a search at a future date who want to remember additional context surrounding the task.

Another SearchTogether feature enabling division of labor among group members is the recommendation mechanism. If a user encounters a webpage that they would like another group member to read, she can select the "Recommend" button from the toolbar (Figure 1i). The user can then select which group members to recommend the page to, and can enter a comment. A positive rating is automatically associated with all recommended pages. When a user receives a recommendation, a thumbnail of the recommended page appears in her "Recommendation Queue" (Figure 1d). Clicking on an item from the queue displays the corresponding webpage, along with any page-specific metadata, in the browser.

"Split Search" and "Multi-Engine Search" are two options SearchTogether offers that provide *automatic* division of labor. When a user enters a query into the search box, he can choose from three different search options. "Standard Search" (Figure 1e) fetches the query's results from the user's default search engine, and displays them to that user only, in the "Current Results" tab (Figure 1c).

"Split Search" (Figure 1f) sends the query to the user's default search engine, and retrieves the highest-ranked results. These results are then divided up among all online group members in a round-robin fashion (so that one user doesn't receive all of the highest-ranked results). The user who issued the query sees his portion of the results in his Current Results tab. For each of the other online group members, a new tab containing his portion of the results appears in the results area. This new tab is labeled with the name of the group member who executed the split search (Figure 2). We placed the split results in a separate tab to avoid interrupting the recipients, allowing them to decide on their own when to task-switch to view the results. Split searching allows multiple group members to all tackle the same sub-task without duplicating efforts, and facilitates parallelization of evaluating a set of query results.

"Multi-Engine Search" (Figure 1g) takes the user's query and sends it to *n* different search engines, where *n* is the number of online group members. The user initiating the multi-engine search is shown a dialog box pairing each online group member with a search source (several general-purpose, travel, health, and shopping search sources are provided). The user may change the assignment of sources to group members via drop-down menus, or accept the default pairings. As with split search, the user executing the multi-engine search sees the results in his Current Results tab, while the other online group members see a separate tab appear, labeled with the sender's name and the source of the results. The multi-engine search is a technique that a group can use to increase their coverage of a topic area while parallelizing their search efforts.

In addition to facilitating parallelization and automatic division of labor, split search and multi-engine search can serve as tools for facilitating involvement in the search process by users who are less technically skilled, and who may be inexperienced at formulating queries but quite capable of evaluating the quality of returned results. Such a technique makes SearchTogether interesting as an educational tool, or as a mechanism that might allow less-experienced searchers to slowly progress from an observational role to more active participation in a shared search task.

Persistence

Storing a shared search session is necessary for enabling asynchronous collaboration. All aspects of a SearchTogether session are persistent, including instant message conversations, query histories, recommendation queues, and page-specific metadata. Thus, when a user logs in, she can see what other group members accomplished while she was offline, as well as have access to her search context to serve as a reminder of what she had been doing when she last logged in.

In addition to automatically saving and restoring all session state, SearchTogether also supports persistence by automatically creating a shared artifact that summarizes the find-

ings of a collaborative search. By default, the summary shows the title, URL, thumbnail, and page-specific metadata for any page that any group member has assigned a positive rating (Figure 3). However, using the options menu, users can change the rules used to generate the summary view. For example, they may choose to include any page that has received a comment, or to include any page viewed by the group as long as it has not received a negative rating. Clicking any entry in the summary causes the corresponding webpage to be displayed in the browser, along with any metadata associated with that page. Viewing the summary is a quick way for users to get up to speed on the current status of a shared task when logging in asynchronously or after an intervening gap in time. The summary also serves as an interactive mechanism for allowing users to quickly re-find useful information, rather than needing to repeat a search.

EVALUATION

We conducted an initial user study of the SearchTogether system in order to answer the following questions:

- Would users find SearchTogether helpful when collaboratively searching the Web? How does it compare to their current collaboration strategies?
- How much does support for awareness, division of labor, and persistence impact the utility of SearchTogether?
- How can the design of SearchTogether be improved?

Methodology

Fourteen subjects (seven male/seven female) completed our study in seven pairs of two users each. All but one of the subjects were employees at Microsoft, in a variety of job roles (*e.g.*, software engineer, program manager, researcher). Each pair of subjects consisted of two people who had a pre-existing relationship. Four pairs were romantically involved couples, two pairs were friends who socialized outside of work, and one pair was a son and his mother (the non-employee).

Each pair of subjects was shown a tutorial in which the experimenter demonstrated the features of the SearchTogether system. After the tutorial, the experimenter asked the pair to choose a topic of mutual interest that they could explore using SearchTogether. Subjects were instructed to select a task that the pair was already planning to investigate. Topics subjects in our study chose to explore included the joint purchase of home-décor related items (two pairs), planning upcoming joint travel (four pairs), and learning more about local opportunities to participate in a shared hobby (one pair).

Subjects were seated in a room containing two computers, each facing opposite walls (*i.e.*, the subjects were facing away from each other). Subjects were instructed to pretend that they were not in the same room, in order to simulate remote collaboration; consequently, they were asked to refrain from speaking to each other during the study task.

Each pair then had 20 minutes to use SearchTogether to jointly search on their chosen topic. Logs of their activity were recorded by the SearchTogether software, and the experimenter took notes on user actions and comments during the session. At the end of 20 minutes, the experimenter distributed a questionnaire to each participant eliciting subjective feedback on the SearchTogether experience. Subjects were asked to fill out their questionnaires individually, without conferring with their partners. The questionnaire asked subjects to list their three favorite and three least favorite features of SearchTogether, to rate how well SearchTogether helped them accomplish their chosen task, to describe their typical strategies for collaborating on tasks involving Web search, and to compare the SearchTogether experience to their normal collaboration strategies.

RESULTS

We now analyze the log, observation, and questionnaire data from our study. We first discuss findings related to the overall utility of SearchTogether, and then discuss more detailed usage patterns, analyzing the utility of the system's awareness, division of labor, and persistence features.

Overall Utility

Our study re-confirmed the survey findings [22] regarding both the topics that people were interested in exploring via collaborative Web search (travel, large purchases, and social planning) and the status quo methods subjects reported using to carry out these collaborative tasks (using the phone or IM in conjunction with a Web browser, or emailing lists of links back and forth).

The SearchTogether system proved to be an effective tool for facilitating these types of synchronous joint search tasks between remote partners. Subjects rated whether SearchTogether helped them achieve their joint goal, showing overall agreement with a mean rating of 3.9 on a five-point Likert scale ($\sigma = 0.62$). Subjects also agreed that SearchTogether was a more effective way to accomplish their task than the methods they would typically use, with a mean rating of 4.1 ($\sigma = 0.62$). Four of the fourteen participants requested that we send them a copy of their SearchTogether summary, highlighting both the realism of their chosen tasks as well as the utility of the system.

Usage Patterns and Feedback

Awareness

The history of their own and their partners' queries was considered a highly useful feature, with eight participants citing query awareness as one of their three favorite aspects of the system. Users frequently clicked on their own and others' query histories in order to re-examine results lists. In total, search results lists were viewed due to an actual search on 189 occasions, and were viewed due to clicking on a prior query on 88 occasions; thus, 31.8% of all search result lists viewed during the study were the consequence of using the interactive query histories. These histories also served their intended purpose of increasing awareness of a partner's activities; for instance, one user commented that this information "saved me some typing" when he noticed that his partner had already entered a query on a topic he had been about to explore.

The association of metadata (visitation history, ratings, and comments) with webpages was also popular, receiving sev-

en mentions in the "three favorite things about SearchTogether" lists. Although the ratings were a popular feature, there was unbalanced use of the positive and negative ratings, with the positive ("thumbs up"), being much more heavily used. A total of 70 thumbs up ratings were given during the study (36 via the "recommend" button and 34 via the "thumbs up" button), while only 9 negative ratings were given. One possible cause for this imbalance is that people found the positive ratings more useful since they not only expressed an opinion, but also impacted inclusion in the summary. We suspect that this bias might shift as people become more familiar with collaborative search tools, and learn to help their partners save time by bypassing pages that appear promising but turn out to have little informational payoff per their shared goals.

The comments on pages fell into two main categories: comments summarizing the content of a page (*"could be a good site for various day trips"*) and comments reiterating specific information found within a page (*"it appears to take three hours to drive between Disneyworld and Fort Myers"*). This latter class of comments suggests that allowing users to highlight specific portions of pages, and making this highlighting visible to other group members on the page itself and in the summary view, might be a valuable capability to add to SearchTogether.

Several users felt that SearchTogether would have been even more valuable had it provided an additional form of awareness information – the ability to see what page another group member was currently viewing. Five users cited the lack of ability to know what their partner is currently looking at as one of their least favorite things about the system. Providing an option to temporarily yoke a user's browser to that of another group member's would be one way to facilitate this type of awareness.

Division of Labor

The integration of instant messaging with the browsing experience was extremely popular among subjects, with six listing it among their favorite aspects of the system. A total of 315 instant messages were sent over the course of the study sessions. Messages fell into four main categories:

- Planning search strategies and assigning sub-tasks (*e.g., "I'll take training." "OK, I'll do running routes."*)
- Discussions related to the general task topic (*e.g., "We need to pack insect repellant."*)
- Sharing a fact discovered during the search (*e.g., "Looks like there are two local airlines."*)
- Telling partners to view a recommendation, divided search, or summary (*e.g., "I just sent you a map."*)

The first two categories illustrate the value of the messaging feature for allowing users to discuss meta-issues related to the task and negotiate division of labor issues. The latter two categories suggest that participants also used the messaging to facilitate awareness. The last category, in particular, suggests that SearchTogether needs to do a better job of notifying users when they have received content from another group member; associating audio feedback with these types of events would be one means of accomplishing this.

The ability to send and receive recommendations was also highly-rated by users; nine of fourteen mentioned recommendations as one of their three favorite things about SearchTogether. System logs show that our subjects sent a total of 36 recommendations during the study; however, only 22 of these recommendations were subsequently viewed by the recipients. The popularity of recommendations, combined with the fact that several of them went unnoticed, suggests that the queue of incoming recommendations should be moved to a more prominent position in the user interface, rather than in the bottom corner of the screen.

However, four users indicated that the conflation of positive ratings and recommendations was one of their least favorite aspects of the system. Subjects indicated that sometimes they wanted the "push" functionality offered by recommend, but without the automatic attachment of a rating. Rather than providing a "recommend" option, providing a "share this" option would allow this more lightweight style of sending pages back and forth.

The automatic division of labor features (split search and multi-engine search) were not heavily used by study participants. A total of five split searches and four multi-engine searches were issued during the entire study. Opinions were divided on the utility of these features, with four users listing them among their favorite aspects of the system and three listing them among their least favorite.

We believe that there are several reasons why these two features may not have been used by our study participants. The use of these features implies that one member of the group is the leader of the activity, and is sending information to the other group members; however, all of the participants in our study were completing a task as peers exploring social topics. We hypothesize that in a business/productivity scenario, especially in cases where one member of the group is of higher rank or position than others, these features might be used more often.

The split and multi-engine search features were also intended to allow for collaboration among people of differing levels of search skill, such as by allowing a more expert searcher to compose a query while still involving other group members in exploring the results. However, all but one of the participants in our study were highly experienced searchers. The only pair for whom this was not true was the mother/son pair. The son is a software developer with a great deal of search experience, and the mother is a homemaker who said she searched the Web only a few times a week, when looking for recipes. The son, surprisingly, did not initiate any split or multi-engine searches during the study. When asked afterwards why he had not used those features, he explained that when he tried instant messaging his mother in order to let her know that he might send her information, she never replied, so he did not send her anything. (The mother was the only user in the study who did not send any instant messages; she later reported being confused by that feature.) The exploration of the util-

ity of split and multi-engine search in scenarios with different combinations of search skill merits further investigation.

Also, split and multi-engine search are quite different than the ways people currently search the web; the low use of these features during the study might simply be a product of their novelty. Perhaps as users become familiar with these capabilities, they would be inclined to use them more. A longitudinal study of SearchTogether would help in examining this possibility in more detail.

Persistence

Because our study focused on a synchronous search scenario that was limited to a single instance in time, we did not have the opportunity to evaluate the value of SearchTogether's persistent state for enabling asynchronous collaboration or for reminding users of the context associated with an interrupted search task. However, even in the brief, synchronous searches in our study, the summary feature proved to be quite useful.

Automatically-generated summaries were one of the most popular features of the system, with six of the fourteen subjects listing it as one of their three favorite things about SearchTogether. Users frequently viewed the summaries and used the summaries to return to the pages they pointed to. Several even requested copies of the summaries for use after the study ended. To further improve the value of the summary as a persistent artifact of a joint search, several users requested the ability to be able to manually edit the summary, in order to organize the pages listed into a meaningful order (rather than the chronological order in which they were first visited).

During the study, users frequently switched between the search results tabs and the summary tab, since they felt that the summary was not only the product of their shared searching efforts, but also a useful tool for communicating with their partners about good pages they had found. In light of this use of the summary as a mechanism for communication and awareness, it would make sense to redesign the interface so that the summary view would be continually visible, in order to reduce task-switching and enhance the summary's utility as a communicative device by increasing its visibility.

CONCLUSION

We presented SearchTogether, a prototype that enables remote users to synchronously or asynchronously collaborate when searching the Web. The design of SearchTogether was guided by prior research on search and sensemaking behavior, particularly by the information about collaborative search needs and behaviors gathered in a survey of 204 knowledge workers [22]. The system aims to support collaboration with several mechanisms for awareness, division of labor, and persistence.

Our user study demonstrated that SearchTogether enabled participants to successfully collaborate on realistic Web search tasks. Participants described their methods for achieving similar goals with currently-available tools, and agreed that SearchTogether allowed them to collaborate more effectively than these other techniques.

Awareness was the most valuable aspect of SearchTogether's design. Features supporting awareness (query histories, visitation histories, ratings, and comments) were among the most highly-rated and utilized aspects of the system. Additionally, features that were intended to support division of labor (such as integrated messaging) and persistence (such as summaries), were also used by participants as awareness mechanisms. Participants even requested additional awareness mechanisms, such as optional yoked browser views.

Although support for flexible division of labor (via integrated messaging and recommendation mechanisms) was appreciated by our study participants, automatic division of labor was not heavily utilized. Whether automatic division of labor would be more valuable with different user populations (*e.g.*, groups with variable search ability levels or groups with clearly identified leaders), task types (*e.g.*, work-oriented tasks), or increased experience with the tool, remains an area for future investigation.

Providing a shared product for a collaborative search, in the form of an automatically-generated summary, was also key to SearchTogether's success, with several subjects requesting copies of their summaries for use after the study. Users' desires to further edit these summaries suggests that providing rich sensemaking experiences is an important aspect of supporting collaborative Web search.

In summary, we introduced the SearchTogether system for collaborative Web search. Our survey found that existing technologies do not adequately support users' desires to engage in active, small-group collaborative searching, and we specifically designed a technical solution to support collaborative search practices and needs identified through our survey. We did so by carefully combining a set of mature technologies (IM, commenting, ratings) and novel ones (automatic division of labor, shared summaries, query awareness, visitation awareness, and persistence). Further, we contributed descriptions of collaborative searching behavior, and based on these observations we provided design knowledge on features that enhanced the collaborative search process, such as awareness mechanisms, flexible division of labor support, and automatically-generated session summaries. Our study also identified ways in which the collaborative searching experience could still be improved, thereby identifying directions for future research in this domain.

ACKNOWLEDGMENTS

We thank Susan Dumais for many insightful conversations about this work. We also thank A.J. Brush, Dan Morris, and Jaime Teevan for their feedback on drafts of this paper.

REFERENCES

1. Aula, A., Jhaveri, N., and Kaki, M. Information Search and Re-access Strategies of Experienced Web Users. *WWW* 2005, 583-592.

2. Beeferman, D. and Berger, A. Agglomerative Clustering of a Search Engine Query Log. *KDD 2000*, 407-416.

3. Bharat, K. SearchPad: Explicit Capture of Search Context to Support Web Search. *WWW 2000*, 493 - 501.

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

5. Bush, V. As We May Think. *The Atlantic Monthly*, July 1945.

6. Cabri, G., Leonardi, L., and Zambonelli, F. Supporting Cooperative WWW Browsing: A Proxy-Based Approach. *Seventh Euromicro Workshop on Parallel and Distributed Processing*, 1999, 138-145.

7. Donath, J. and Robertson, N. The Sociable Web. *Second International WWW Conference*, 1994.

8. Dontcheva, M., Drucker, S., Wade, G., Salesin, D., Cohen, M. Summarizing Personal Web Browsing Sessions. *UIST 2006*, 115-124.

9. Green, S.L., Devlin, S.J., Cannata, P.E., and Gomez, L.M. No IFs, ANDs, or ORs: A Study of Database Querying. *International Journal of Man-Machine Studies*, 32, 3 (1990), 303-326.

10. Gianoutsos, S. and Grundy, J. Collaborative Work with the World Wide Web: Adding CSCW Support to a Web Browser. *OZ-CSCW 1996*, 14-21.

11. Greenberg, S. and Roseman, M. GroupWeb: A WWW Browser as Real Time Groupware. *CHI 1996 Conference Companion*.

12. Han, R., Perrett, V., and Naghshineh, M. WebSplitter: A Unified XML Framework for Mutli-Device Collaborative Web Browsing. *CSCW 2000*, 221-230.

13. Hinckley, K., Zhao, S., Sarin, R., Baudisch, P., Cutrell, E., Shilman, M., and Tan, D. InkSeine. *CHI 2007*, 251-260.

14. Hoppe, H.U. and Zhao, J. C-TORI: An Interface for Cooperative Database Retrieval. In Karagiannis, D. (ed), *Database and Expert Systems Applications*, Springer-Verlag.

15. Jones, R., Rey, B., Madani, O., and Greiner, W. Generating Query Substitutions. *WWW 2006*, 387-396.

16. Jones, W., Bruce, H., and Dumais, S. How Do People Get Back to Information on the Web? How Can They Do it Better? *Interact 2003*, 793-796.

17. Keller, R., Wolf, S., Chen, J., Rabinowitz, J., and Mathe, N. A Bookmarking Service for Organizing and Sharing URLs. *WWW 1997*.

18. Krishnappa, R. Multi-User Search Engine: Supporting Collaborative Information Seeking and Retrieval. *Master's Thesis, University of Missouri-Rolla*, 2005.

19. Large, A., Beheshti, J., and Rahman, T. Gender Differences in Collaborative Web Searching Behavior: An Elementary School Study. *Information Processing and Management*, Vol. 38, 2002, 427-433.

20. Lee, Y-J. VisSearch: A Collaborative Web Searching Environment. *Computers & Edjucation*, 44(4), May 2005, 423-439.

21. Maekawa, T., Hara, T., and Nishio, S. A Collaborative Web Browsing System for Multiple Mobile Users. *PERCOM 2006*, 22-35.

22. Morris, M.R. Collaborating Alone and Together: Investigating Persistent and Multi-User Web Search Activities, *Microsoft Research Technical Report #MSR-TR-2007-11*, January 2007.

23. Morris, M.R. and Horvitz, E. S^3: Storable, Shareable Search. *Interact 2007*, in press.

24. Morris, M.R., Paepcke, A., and Winograd, T. TeamSearch: Comparing Techniques for Co-Present Collaborative Search of Digital Media. *IEEE Tabletop 2006*, 97-104.

25. Pujol, J.M., Sangüesa, R., and Bermúdez, J. Porqpine: A Distributed and Collaborative Search Engine. *WWW 2003 (poster)*.

26. Romano, N., Nunamaker, J., Roussinov, D., and Chen, H. Collaborative Information Retrieval Environment: Integration of Information Retrieval with Group Support Systems. *Hawaii International Conference on System Sciences*, 1999.

27. Rose, D. and Levinson, D. Understanding User Goals in Web Search. *WWW 2004*, 13-19.

28. Russell, D., Stefik, M., Pirolli, P., and Card, S. The Cost Structure of Sensemaking. *CHI 1993*, 269-276.

29. schraefel, m.c., Zhu, Y., Modjeska, D., Wigdor, D., and Zhao, S. Hunter Gatherer: Interaction Support for the Creation and Management of Within-Web-Page Collections. *WWW 2002*.

30. Tauscher, L. and Greenberg, S. Revisitation Patterns in World Wide Web Navigation. *CHI 1997*, 399-406.

31. Teevan, J., Adar, E., Jones, R., and Potts, M. History Repeats Itself: Repeat Queries in Yahoo's Query Logs. *SIGIR 2006*, 703-704.

32. Twidale, M., Nichols, D., and Paice, C. Browsing is a Collaborative Process. *Information Processing and Management*, 33(6), 1997, 761-783.

33. White, R., Kules, B., Drucker, S., and schraefel, m.c. Supporting Exploratory Search. *Communications of the ACM*, 49(4).

34. Wittenburg, K., Das, D., Hill, W., and Stead, L. Group Asynchronous Browsing on the World Wide Web. *Fourth World Wide Web Conference*, 1995.