

Facilitating Opportunistic Communication by Tracking the Documents People Use

Jay Budzik, Xiaobin Fu, and Kristian J. Hammond

Intelligent Information Laboratory

Northwestern University

1890 Maple Ave.

Evanston, IL 60201 USA

+1 847 467 1265

{budzik, fu, hammond}@infolab.nwu.edu

ABSTRACT

Many of our conversations are not planned. Instead, they arise opportunistically as we become aware of the common interests and goals of others. We believe communication tools can be built that leverage knowledge of a user's goals and activities to make users aware of others relevant to them and their work, with the aim of facilitating this kind of opportunistic interaction. We describe a prototype system, I2I, that helps people establish communication while they are manipulating related documents by embedding awareness and communication facilities directly into everyday document manipulation applications. I2I automatically clusters the documents users are manipulating based on their content, grouping related documents into a single conceptual space. It then makes the common work contexts of users visible by displaying only those users who are manipulating related documents. I2I also allows users to initiate conversations asynchronously through a facility we call calling cards. In these ways, I2I provides users with opportunities to communicate within the context of the activities they are performing in their primary application, without requiring they manually orchestrate the communication themselves. Once users become aware of each other, they can use the system to communicate using a variety of interactive modalities.

Keywords

computer mediated communication, awareness, WWW, matchmaking, communityware, context

INTRODUCTION

Many efforts in building collaborative systems have focused on developing techniques to support awareness among group members engaged in a specific task (e.g., [7, 15]). While tools that support group awareness in the context of inherently collaborative work (e.g., meeting scheduling, collaborative design, etc.) are becoming more common, less attention has been given to supporting the kind of awareness necessary to give rise to collaboration in the first place—the kind of informal collaboration and communication that commonly occur in physically-collocated settings organized around common goals (as

Kraut, et al. [9], Kristoffersen and Ljungberg [10], and others have observed). Our focus in this work is to make the kind of opportunities for informal collaboration that are currently available in the physical world available in the electronic realm.

Handling this kind of interaction is a challenge for computer-mediated communication (CMC) systems. The challenge is to design systems that support interactions among people that are opportunistic and are often based on a establishing a shared context for the interaction. If CMC systems are designed to notice opportunities for collaboration by tracking the work people do in everyday applications, they can make users aware of common interests and goals, so that people can easily establish connections in order to cooperate and share with each other. Combined with standard collaboration and communication tools, such a system can transform traditionally solitary activities into collaborative ones by providing users with frictionless access to potential collaborators.

In this paper, we describe the design and implementation of I2I, a prototype system that represents our first attempt to address this challenge. I2I is currently embedded in Microsoft Word, Microsoft Internet Explorer, and Netscape Navigator, all of which are popular document manipulation applications that run on the Windows platform. I2I automatically clusters the documents users are manipulating based on their content, grouping related documents into a single conceptual space, allowing users to:

1. Establish one-on-one communication with others who are manipulating related documents.
2. Initiate conversations asynchronously through a facility we call calling cards.
3. Browse related information items automatically recommended by the system.
4. Join or start public chat rooms associated with the content they are viewing.

I2I manages the early stages of initiating informal collaboration by providing its users with opportunities to become aware of the activities of others that share common

interests, as represented by the documents they interact with. I2I attempts to build communities of common interest on the fly, allowing users engaged in traditionally solitary activities to discover common goals and collaborate with each other, while reducing the overhead of orchestrating the collaboration.

RELATED WORK

Tools that allow distributed users to collaborate around common electronic artifacts have had a long history, although much recent work has focused on collaboration around documents. Anchored Conversations [3], for example, allow collaborators to easily distribute shared documents and situate conversations within the context of specific places in a shared document. Ensuring collaborators share the same artifact makes collaborative activities that depend strongly on artifacts (such as collaborative writing) easier.

Tools of this nature are aimed at supporting collaboration among users who *already know each other* and have a prior goal to collaborate. Our work intends to provide opportunities for users who may not know each other to collaborate informally by making opportunities for collaboration visible, and by automating the early stages of establishing collaboration (e.g., knowing who to talk to, and how to talk with them).

Other systems have examined the role of distributed, public artifacts (e.g., Web pages) as shared contexts, allowing users who are manipulating or viewing the same object from distributed locations to communicate with each other, usually using text-based chat (e.g., [4, 6, 11]).

This vein of work is the closest to the work we describe here. The main difference is that the above systems require objects tagged by unique identifiers (in the case of the Web, the page's URL), and also that users manipulate the same object at the same time in order to collaborate. These requirements limit the opportunities for collaboration the system can make available, due to the sheer size of the Web and typical patterns of access (Web proxy log analyses have shown it is unlikely two users will be on the same page at the same time except at the most popular sites [1]). The exclusion of unpublished electronic documents (e.g., the paper I'm working on right now), as well as the fact that similar content can be found on many different URLs, also limits such systems.

I2I overcomes these problems by analyzing the content of the objects being manipulated by users in order to automatically cluster similar content. In effect, I2I builds a separate conceptual space, organized by the content of documents being manipulated, and then situates users and other information items in this space. In addition, I2I provides tools for asynchronous communication, allowing users to notice opportunities for collaboration across time.

I2I also expands the boundaries of matchmaking systems (e.g., [5, 8]), which introduce users with common interests

to each other with the goal of building online communities and fostering community awareness. Work on matchmaking systems has generally focused on introducing users based on their long-term interests. In the Yenta system [5], for example, users submit a collection of documents to their personal agent, which builds a profile from those documents and executes a kind of distributed hill-climbing algorithm to match profiles. When a matching profile is found the agent arranges to introduce the two users.

In contrast, our work on I2I focuses on introducing users based on their immediate (and perhaps short-lived) interests that arise from the tasks they in which they are currently engaged. Instead of requiring the users define a profile for themselves using documents or keywords, I2I automatically builds a lexical representation of the user's current activity (as represented by the document the user is manipulating) and uses this representation to determine what the user can see.

I2I: DESIGN AND IMPLEMENTATION

I2I integrates with Windows applications through COM, Microsoft's object-oriented inter-process communication facility. Each application has its own application adapter, which is responsible for communicating user actions and document content to a central broker. The broker is responsible for persistent information such as the user's profile (e.g., their name, password, etc.), as well as ephemeral information, such as how to contact their machine, and the representation of the document they are currently manipulating.

The Central Broker

Each application adapter is responsible for sending the broker a message when the document has changed in an attached application (e.g., the document is edited significantly, or the user opens or navigates to a new one). The broker computes a term vector, representing the current document under the vector space model [12], commonly used in information retrieval systems. The similarity of two documents is computed as the cosine of the angle formed by the two vectors that represent them. Vectors are weighted using the *tfidf* heuristic [14]. Associated with each vector is the title of the document it represents, the URL (if the document is a Web page), a list of users manipulating that document, as well as a list of calling cards associated with the document.

The broker computes a similarity matrix for documents that are currently being manipulated. The similarity matrix determines what a user sees from the vantage point of a particular document (see Figure 1). Objects associated with documents D whose similarity is above a threshold θ with respect to d , the current document, are visible from d (currently θ is set heuristically to 0.75, but we are exploring interfaces that allow users to manipulate this parameter interactively).




Figure 1: I2I builds a conceptual space based on the content of documents, regardless of their actual location.

The chance of two people reading exactly the same document at the same time may be slim. By grouping conceptually similar documents together, I2I makes it more likely that people will see each other and start a conversation. It also allows unpublished documents (e.g., a paper in progress) to serve as the entry point into the system.

Secondary objects can also be associated with a document in the space I2I has built. The simplest of these objects is people: users who are viewing a particular document are associated with that document's point in the space (see Figure 1). Currently, I2I also indexes chat rooms and calling cards (a facility for asynchronous communication) in the same way. Users who access a document, then, can see the items associated with it and other documents close to it in the conceptual space (see Figure 1).

Note that the techniques we use for clustering documents have been shown to be effective (for example, similar techniques produce huge improvements in information retrieval [13]), but that unintuitive associations among users may occur. It is also sometimes the case that the user's current document does not provide a very good window onto her goals (e.g., a single document can have multiple purposes).

However, it is important to realize that the system does not require users to collaborate. It provides users with opportunities for collaboration by automatically recommending potential collaborators. Users can make decisions about whether to collaborate with each other based on their current needs and by inspecting the documents others are viewing (if they are available online), or by viewing background information about the user (should users make such information available though the system).

It is also important to note that the novelty of this system does not lie in the document similarity measures we use—




Figure 2: I2I embedded in Microsoft Internet Explorer.

they are well known and have been shown to be successful in other settings. Instead, it lies in the use of these techniques in conjunction with the system's ability observe the user's behavior in an application that together allow it to associate similar documents with each other, and also with objects like chat rooms, people, and calling cards, while requiring minimal effort on the part of its users.

Application-embedded Client Interface

I2I tracks a user's current task context (represented by the document they are manipulating) so it can provide potentially useful resources to users in the context of a specific editing or browsing session. It embeds an interface for displaying this information directly into applications, where it is supported (see Figure 2) to allow the user to easily correspond the information I2I displays with the document it is associated with. In other cases, information is displayed in an associated window that can be "hooked" on to the main window of the application, to maintain visual correspondence. This allows users to easily keep track of their activity in several conceptual spaces at the same time.

Details of the embedded interface are shown in Figures 3 and 4. Information is grouped into tabs and includes (from left to right):

1. *System activity.* Users can see how many I2I users are online both in and outside of the conceptual space defined by their document. Other activity information includes how many people are chatting, and how many related pages other I2I users are reading.
2. *Who is online.* Users can see the login names of the people reading or writing related documents, and pointers to the documents they are viewing, if they are available on the Web (see Figure 3). Users can contact each other directly via double-blind email, or by using videoconferencing software, depending on the software and hardware available on their machine.
3. *Related pages.* I2I displays related pages from other Web sites people are currently browsing. In addition, I2I displays recommendations generated by the Watson system [2]. The Watson system recommends related documents by automatically querying Internet search engines.




Figure 3. People visible from a page about Nixon include users who are viewing different documents on the same topic.

4. *Active chat information.* I2I displays a list of chat topics created by users within the conceptual space defined by the current document. Users can also chat in a default room associated with this place in the content space.
5. *Calling cards.* I2I displays a list of calling cards that other users have left in the past while viewing the current or related documents (see Figure 4). A calling card is a note indicating that a user would like to talk about a particular topic.

Calling cards

Users can leave calling cards associated with the content area represented by their document in order to indicate they would like to discuss a particular topic with other users. Figure 4 shows calling cards associated with a page about Richard Nixon.

Leaving a calling card allows users to make their goal to chat about a particular topic visible to other users who also view pages in that topic area. If a user is eager to open a discussion channel with somebody else, but no one is available or has shown interest, the user can leave a message to invite people to talk at a later date. After leaving a calling card, the user can continue to work, or even *destroy the original document it was associated with*. Calling cards are indexed by a content vector representing the document at which they were created. This means access to the document is not necessary for other users to see them when they are browsing or writing in related areas.

For example, one user could leave a calling card at the document in Figure 3, which discusses Nixon and his presidency. That document could then be taken off of the Web. At this point, other users would still be able to see the calling card when they accessed other documents about Nixon, for example, the page in Figure 4.

This kind of indexing also nicely accommodates documents that have frequently-updated content (like the front page of




Figure 4: Calling cards visible from a page about Nixon.

a news site), because even though the content of a page might change, the system associates the calling card with the original context in which it was created. This ensures that the calling cards retrieved are actually relevant in the context of the document being viewed.

When another user sees the calling card, she can find out whether the owner of the calling card is online or not. If the owner is online, she simply drops a line to the user to say that she is interested in discussing the topic. If the owner is not online, she can find out if the owner has a public email address and send an email to the owner (if the user who left the note has specified others can contact her via email).

A calling card persists for a time period specified when it is created (currently the system imposes a limit of 30 days). When a calling card expires, the owner is notified via the global interface. The user can then choose to delete the card, or extend the time period in which it is available.

Managing Privacy

Users may be uncomfortable, at times, having a system track what they write or view. I2I allows people to manage the privacy of their work by being highly visible when it is on (see Figures 2-4), and by allowing users to shut it off at any time (using the close button in the interface). I2I also does not expose the details of offline (non-Web) content to any third party. In addition, it does not expose a user's email address or online identity directly. Instead email is sent through a mediating server that automatically makes message sender and recipient anonymous. This allows users to disclose their real email addresses at their own discretion.

Global User Interface

I2I has a global interface that allows the user to control whether or not she is available for conversation, as well as edit her profile and preferences. Calling cards are also

managed using this interface (owners of calling cards can edit or delete them at any time).

ONGOING AND FUTURE WORK

We have distributed I2I to several researchers in our department for limited testing. For the most part, the feedback was positive. Users said they like the sense of being in a community and enjoy the kind of ready connectivity brought by I2I. We were encouraged by this initial test, but are currently working on more thorough evaluations of both the system's general usability, as well as how it affects the performance of users in particular contexts of use. In particular, we are planning to test the system with students in a large freshman survey course. We anticipate that those who use I2I will be more productive, and have a greater sense of community when compared with those who don't. In the near future, we intend to deploy the system on our campus, to test scalability, study usage patterns, and gather user feedback.

We are also working on improved techniques for clustering users and other objects. Efforts in this area involve filtering potential collaborators by profiles built from their long-term history of interacting with documents (so that people with similar backgrounds are preferred), and incorporating task models, where possible. We are also investigating interfaces that allow the user to narrow or broaden the conceptual spaces in which they are situated.

CONCLUSION

The growing importance of the Internet is changing the way people access information and the dynamics of how they interact with each other online. Although there are a wide variety of tools that support communication over the Web and Internet, they tend not to support the kind of informal and opportunistic communication that arises out of an awareness of common goals and interests. I2I attempts to address this issue by embedding communication facilities in the user's everyday applications so that users that share interests can be aware of each other and communicate freely in an informal environment, even though they may have never met nor discussed the interests they share.

ACKNOWLEDGEMENTS

Andrei Scheinkman contributed to the development of early I2I prototypes. Mason Warner developed a robust Java version of the system. This work has benefited from the comments of anonymous reviewers, Larry Birnbaum, Louis Gomez, Peter Dinda, and Sparky Wilson.

REFERENCES

1. Breslau, L., Cao, P., Fan, L., Phillips, G., and Shenker, S., "On the Implications of Zipf's Law for Web Caching," in *Proceedings of IEEE INFOCOM '99*, (New York, USA), IEEE Press, 1999.
2. Budzik, J., and Hammond, K. J., "User Interactions with Everyday Applications as Context for Just-in-time Information Access," in *Proc. IUI 2000* (New Orleans, Louisiana, USA), ACM Press, 2000.
3. Churchill, E. F., Trevor, J., Bly, S., Nelson, L., and Cubranic, D., "Anchored Conversations: Chatting in the Context of a Document," in *Proceedings of CHI 2000*, (The Hague, The Netherlands), ACM Press, 2000.
4. Donath, J. S., and Robertson, N., "The Sociable Web," in *Proceedings of the Second International WWW Conference*, (Chicago, IL), Elsevier, 1994.
5. Foner, L., "Yenta: A Multi-Agent, Referral Based Matchmaking System," in *Proceedings of Agents 97*, (Marina del Rey, CA USA)1997.
6. Gooley, Gooley, Available at <http://www.getgooley.com/>.
7. Gutwin, C., and Greenberg, S., "The Effects of Workspace Awareness Support on the Usability of Real-Time Distributed Groupware," *ACM Transactions on Computer-Human Interaction*, 6(3), 243-281, 2000.
8. Hattori, F., Ohguro, T., Yokoo, M., Matsubara, S., and Yoshida, S., "Socialware: Multiagent Systems for Supporting Network Communities," *Communications of the ACM*, 42(3), , 1999.
9. Kraut, R., Egido, C., "Patterns of Contact and Communication in Scientific Research Collaboration," in *Proceedings of CSCW 88*, (Portland, OR USA), ACM Press, 1988.
10. Kristoffersen, S., Ljungberg, F., "An Empirical Study of How People Establish Interaction: Implications for CSCW Session Management Models," in *Proceedings of CHI 99*, (Pittsburgh, PA USA), ACM Press, 1999.
11. Palfreyman, K., and Rodden, T., "A Protocol for User Awareness on the World Wide Web," in *Proceedings of CSCW 96*, (Cambridge, MA USA), ACM Press, 1996.
12. Salton, G., Wong, A., and Yang, C. S., "A vector space model for automatic indexing," *Communications of the ACM*, 18(11), 613-620, 1971.
13. Salton, G., and Buckley, C., "Improving Retrieval Performance by Relevance Feedback," in: Spark Jones, K., and Willett, P., ed., *Readings in Information Retrieval*. Can Francisco, CA: Morgan Kauffman, 1997.
14. Salton, G., and Buckley, C., "Term-Weighting Approaches in Automatic Text Retrieval," in: Spark Jones, K., and Willett, P., ed., *Readings in Information Retrieval*. San Francisco, CA: Morgan Kaufmann, 1997.
15. Whittaker, S., Swanson, J., Kucan, J., and Sidner, C., "Telenotes: Managing Lightweight Interactions in the Desktop," *ACM Transactions on Computer-Human Interaction*, 4(2), 137-168, 1997.