

“Ask Before You Search”

Peer Support and Community Building with ReachOut

Amnon Ribak, Michal Jacovi, Vladimir Soroka

IBM Research Lab in Haifa

Haifa University Campus

Haifa – 31905, Israel

+972-4-8296408

{ribak, jacovi, vladi}@il.ibm.com

ABSTRACT

This paper presents ReachOut, a chat-based tool for peer support, collaboration, and community building. We describe the philosophy behind the tool and explain how posting questions in the open directly benefits the creation, distribution, and use of organizational knowledge, in addition to enhancing the cohesion of the community involved. ReachOut proposes new methods of handling problems that include locating, selecting, and approaching the right set of potential advisers. We discuss the advantages of public discussions over private, one-on-one sessions, and how this is enhanced by our unique combination of synchronous and asynchronous communication. We present and analyze results from a pilot of ReachOut and conclude with plans for future research and development.

Keywords

Peer support, organizational knowledge, community building, instant messaging, expertise location

INTRODUCTION

Information Technology (IT) allows people to work together while situated physically remote from each other. Although convenient to the individual, this separation has its drawbacks. By replacing face-to-face communication with telephone, e-mail, and instant messaging, we have also forfeited overhearing hallway conversations and the constant subconscious awareness of the state of our teams and work environment. In a study conducted by Teasley et al. [31], development teams that were seated in ‘war rooms’ and worked in ‘radical collocation’ exhibited twice as much productivity and one-third less time-to-market, as compared to the company’s baseline. Technology has deprived us of these benefits once granted to collocated teams; a good system for Computer Supported Cooperative Work (CSCW) should at least partially compensate for this loss of unnoticed, casual knowledge distribution.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page.

To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CSCW’02, November 16-20, 2002, New Orleans, Louisiana, USA.

Copyright 2002 ACM 1-58113-560-2/02/0011...\$5.00

IT has also introduced some cultural obstacles to knowledge sharing. The abundance of content on the Internet, and the ease with which new content can be added, have created an image of a land rich in layers upon layers of knowledge, just waiting to be mined and excavated. The dominating paradigm is that of a lonely explorer or navigator, and the implied motto is “Search before you ask”. Organization structure and culture can also play key roles in discouraging people from seeking the help of others. In many organizations, knowledge is viewed as a competitive asset – an advantage that people will not easily give away [15].

Communities, characterized by voluntary interactions, mutual reciprocity, and a strong sense of accountability, usually provide a natural environment for knowledge sharing. Communities are often seen as an informal network that exists in parallel to the formal organizational structure, providing a sense of democracy and equality that the organization cannot offer. This same organization will often provide the facilities and resources needed for its communities to flourish, thereby acknowledging their importance and contribution to the productivity and well being of its employees. A good CSCW system for knowledge sharing must therefore nourish community building [22].

Although the amount and quality of accessible information have dramatically increased with IT, not every answer can be found online, and not all knowledge is explicitly documented. Recently, more and more organizations recognize the importance of undocumented knowledge, and the key role it plays as part of the organizational memory [1,4,9]. This organizational implicit knowledge includes shared heritage, intuitions, and values, and is the access point to documented and undocumented answers on who knows what, what people are doing, who to turn to for advice, etc. By definition, this knowledge is hard to verbalize and cannot be reduced to rules and recipes; in addition, using it relies on communicating with people. Even when knowledge *is* documented, written text is not always the best source for others to acquire that knowledge and make it their own. In many cases, talking to people may serve as an instant “readers digest”, helping you figure out

the important bits and absorb the information, molding it into your own knowledge.

Knowledge workers, who perform missions which are not only complicated but are constantly changing, must therefore rely on others and on the knowledge that others possess. In many organizations, the employee's key asset is their network of contacts and those people they can approach for advice or help [3]. Information technology allows employees to access the knowledge and experience of people way beyond their immediate network. While some initiatives focus on finding content experts [3,14,16,20], others center on getting the support of peers [12]. The success of the UseNet newsgroups and discussion forums is a well-known example of this use of IT [11].

However, discussion forums are not fully adapted to changes and demands in the corporate world. Less people today actively enter newsgroups and corporate forums in order to assist others [24]. This is probably due to two factors: information overload that made attention the scarcest resource of all, as well as the need to revisit a forum in order to stay updated. New interaction techniques, and especially Instant Messaging (IM), have introduced new, more attractive rhythms and ambiance with which asynchronous forums cannot compete.

We developed ReachOut to incorporate IM and its benefits with the power of focused discussion forums. ReachOut is a methodology and a chat-based tool for peer support and community building. It provides a unified interface for both advisers and askers, allowing them to participate in existing discussions and to create new ones. To start a discussion, a user phrases the question title, supplies some background information, and selects target groups. The system invites potential advisers to join a conference chat discussion, by pushing the question to a non-intrusive client on their screen. Somewhat oversimplifying, ReachOut can be described as bridging the gap between newsgroups – which focus on particular subjects and have the ability to retain discussion history – and real-time synchronous chat – with its immediate awareness of new questions, light tone, and interactive mode. ReachOut takes the best of both, and adds push technology to portray new, by-topic awareness and mid-level persistency. In this paper, we present ReachOut and discuss its features, and how they provide for knowledge sharing, collaboration, and community building.

We start the paper with an overview of related work. We then present the ReachOut philosophy, along with the model and design of the tool. We continue with implementation details, and describe a pilot, conducted over a period of two months with 130 users. The paper concludes with plans for future research and development.

RELATED WORK

ReachOut has roots among many well-known trends, domains and tools in the field of CSCW. It draws from Bulletin Board Systems (BBS) and newsgroups, mailing

lists, recommender systems and profile matching, awareness through publish/subscribe technology, question answering, expertise location systems, virtual, on-line communities, and the use of Instant Messaging (IM) for collaboration.

Newsgroups allow people to post questions, open discussions, and seek information based on a predefined taxonomy [11]. Help is being sought 'from strangers', without addressing anyone in particular. Research has shown that strangers often offer more help than acquaintances, due to their diversity of backgrounds [35]. Newsgroups thus became a sort of "group mind" for problem solving [29]. A major drawback of newsgroups is the need to revisit them periodically to discover new questions or new answers, and the lack of high-level support for awareness, which is essential for a good CSCW system [10]. The publish/subscribe technology is a simple yet efficient solution to the awareness problem [13,26]. Mailing list servers (LISTSERV) used this approach by allowing users to subscribe to certain groups and receive email messages on their topics of interest. However, mailing lists put a heavy strain on bandwidth and storage, as every message is sent to multiple targets, and on readers, having to manage the discussion threads in their own mailboxes.

Information overload on the reader's side is, indeed, another problem faced by both newsgroups and mailing lists, which sometimes can even be viewed as spam or junk mail. To handle the huge amount of information, users were forced to adopt cognitive, economic, and social techniques for information filtering [19]. Tools such as GroupLens [27], RINGO [30], and PHOAKS [33] are some examples of 'social navigation' [7] that use recommendations of others to distinguish important from unimportant content. Recommender systems have matured into a branch of their own, using similarity of profiles to recommend products and information [28].

Question answering is another well-established means of information sharing. Both in newsgroups and in collections of Frequently Asked Questions (FAQ), people find answers to specific questions presented in natural language. FAQs, by nature, are independent of other users' good will, giving them a clear advantage, especially in disconnected mode. Yet, despite their name, FAQs tend to be supply-based rather than demand-based. Ackerman and McDonald dealt with this issue by introducing Answer Garden [1] and later Answer Garden 2 [2]. By allowing users to forward questions to an expert, both systems offered a bridge between FAQs and newsgroups.

Expertise location systems recognize the role of tacit knowledge, and the fact that, by definition, it lies only within the minds of people. These systems therefore try to harness IT to locate people and engage with them in order to benefit from their tacit knowledge. They are becoming very popular in both research [3,20] and commercial

implementations [14,16]. Expertise location systems still assume that answers should come from experts – often experts whose main job is to help others – and use previously created content to reduce their workload. This ignores the powerful concept of peer-support [12], and the advantages that sharing information needs and conducting public discussions have for the enhancement of community cohesion.

Computer-supported, virtual communities have become very popular in the context of CSCW. Successful examples of virtual communities abound, including Babble [5], and WELL [29,34], although their added value is still described in qualitative rather than quantitative terms. Many of the CSCW branches mentioned above – newsgroups, recommender systems, and expertise location systems – can be viewed as helping in community building; research is conducted to understand the characteristics of effective communities [32].

Chat, or Instant Messaging, is the most popular means of communication in virtual communities. Despite the drawbacks that the unstructured nature of chat has for archiving and navigation, chat – with its light, casual and informal tone and its interactive, instant support for negotiation of meaning [23] – has many characteristics needed for free flow and sharing of tacit knowledge, and is confidently making its way to the corporate world.

While topic-oriented chats, like Internet Relay Chat (IRC) channels, compensate for some of the chat drawbacks, they still focus on broad topics rather than particular subjects.

ReachOut tries to take the best of all these technologies, and come up with an alternative collaboration and knowledge sharing solution. It takes the structure from newsgroups, but supplies easier navigation through interest groups and discussion threads. It uses the publish/subscribe technology from listserv, but focuses on more narrow topics to prevent information overload. It also brings awareness to events from Instant Messaging. Its collaboration interface is light, like that of chat-based tools, but is more persistent and focused on particular questions or discussion subjects. In addition, ReachOut's deployment model helps build a sense of real community in organizations, bridging between technologies 'for fun' and for the workplace.

In the next sections, we describe the ReachOut tool in more detail.

REACHOUT – PHILOSOPHY AND DESIGN

By applying modern interaction techniques to well-trying knowledge sharing paradigms, ReachOut calls for a 'retro' movement in Knowledge Management, by encouraging users to pose questions and seek the advice of fellow workers. While today's 'common practice' is to first look for information on the Internet or the organization's intranet, we suggest turning directly to people, for the following reasons:

- People within one's group, department, division, or corporation may not only have the needed information, but by direct interaction and richer context, they can formulate it into relevant, action-enabling knowledge. By turning to people, one exploits the most advanced search engine that exists – the retrieval mechanism of the human mind – and in this case, the minds of many.
- Given the right culture and technology, most people are willing to share knowledge and experience, their incentive being good will, social capital, and personal prestige.
- Turning to people within the organization may not only speed up finding solutions for specific problems, it can also create a valuable social network. This network can then help foster new ideas and may connect similar projects to create new, better solutions, at reduced cost.
- Questions should be asked in the open, even if the answers to them can be found elsewhere (e.g., by search). Each question, and each answer, carries with it a lot of additional, implicit information about people, their work, their skills, and their attitudes. This information is then stored in the minds of readers, and weaved into the organizational knowledge.
- Preference should be given to open, public discussions over one-on-one sessions between the asker and each adviser. Shared effort is often the only way to reach a solution. In addition, open discussions foster a sense of community, by letting members 'follow' even if they themselves do not contribute directly to the discussion.
- Whenever possible, discussions should be handled in a real-time, synchronous manner. The informal and lightweight tone of chats usually includes greetings and gestures, which provide immediate social gratification to the adviser – gratification that is missing in asynchronous communication [23].
- Discussions should be carried out in both synchronous and asynchronous mode, over a span of days. In international organizations, people may work in different time zones and the end-of-day for one person may mean the morning for someone else. A discussion should be kept open even when there is no participant online, and users that join it should be able to see its history in terms of participants and their contributions.
- Discussions, however, should not be kept alive forever. There is no point in creating more written, explicit knowledge; it is the real collaboration and the sharing of tacit knowledge that needs to be fostered. Providing access to previous discussions will imply the traditional "Search before you ask".

The ReachOut Model

ReachOut supports an environment where peers are aware of each other's information needs, and are easily accessible to offer their knowledge and experience.

When first registering with ReachOut, users provide a profile that best describes them. The profile may include affiliation, location, role, and level in the organization, but more importantly, it indicates fields of interest. These profiles will later dictate which questions are directed to which users. Users can state their level of willingness to receive questions for each topic or group, thus controlling the number of questions they will receive in each group. They can further control it at any time by stating their 'current' motivation to answer, thus filtering out questions on busy days. By letting its users define their own profile, ReachOut distinguishes itself from other expertise location tools that automatically create such profiles [14,16,20]. Although the latter may seem to free the user from a tedious job, they also take away a sense of autonomy that is offered by ReachOut. This autonomy offers a sense of democracy and is a necessary condition for the existence of a community [6].

ReachOut users express their questions in a very simple form, and if desired, supply background information. The system analyzes the input and the question is transferred to the most appropriate advisers, based on their profiles, the question, the profile of the asker, and temporal conditions. Literature on expertise location distinguishes between 'locating' potential advisers and 'selecting' the ones to approach for assistance [3]. ReachOut not only handles these steps, but also inherently addresses the third step, that of 'approaching' the advisers. By posting questions to the system, ReachOut users avoid the initial embarrassment of requesting help from strangers.

The system directs the question to potential advisers based on a load balancing algorithm that takes into consideration factors such as: the adviser's motivation to answer questions on the topic, or to participate at this very instant, as well as the adviser's history of participation. Incoming questions are pushed to a non-intrusive client that resides on the user's screen. The questions fade in and out, until the user decides to open a question and join its discussion. The same client serves both askers and advisers providing awareness of ongoing discussions and access to peers.

Each question constitutes a virtual chat room, where all participants can have a discussion. Clarification questions may be asked and different ideas proposed, until a solution is found. Other advisors can also be invited to join ongoing discussions in the virtual chat room.

These chat rooms are kept open and available, even when they have no participants, thereby allowing users to join at a later time. ReachOut is among the few instant-messaging tools to offer this persistency [5], usually found only in asynchronous environments like newsgroups and team rooms. After some configurable inactivity period, the discussion is 'closed' and is no longer accessible to users of the system. This prevents them from relying on old content rather than seeking advice.

REACHOUT IMPLEMENTATION

ReachOut is implemented as a typical client-server application, and uses Lotus Sametime™ [17] as its backbone. Sametime provides a number of features essential for online collaboration, including instant messaging, a meeting center, storage services, etc. Sametime uses the 'place' paradigm, as a shared online space for collaboration; a space that is capable of containing data, making it readily available for those who have entered the place, and supplying 'who is here' information for all visitors. ReachOut makes extensive use of Sametime places by creating a place for each group – allowing people to subscribe to a group of their interest by entering the place – and a place for each discussion – giving it a separate area for a focused discussion.

This design decision has freed us completely from having to handle connectivity issues such as broadcasting and bandwidth. ReachOut clients, and the ReachOut manager, which handles users and discussions, use the Community Services of Sametime, to 'chat' with each other. The use of the Sametime Java API for this purpose, as well as for security and registration, has allowed us to focus on the user experience and ease of use, while creating a simple yet effective application.

The first client instance of ReachOut to be implemented was a Java applet, embedded in an e-Learning environment, where it supplied awareness and a discussion channel between users taking the same online course, or related ones, over a defined period of time. This incarnation of ReachOut is beyond the scope of this paper and is not described here.

Two instances of the ReachOut client were developed as Java applications for the pilot described in this paper. Many improvements were introduced as a direct result of the pilot, and most of the requirements and suggestions for new features were discussed among the pilot users by using the tool itself for discussion. The user experience and the user interface described here, belong to the second, improved version. Where appropriate, we describe the variation from the initial version and the discussions that led to the change.

User Experience

The main window of the ReachOut client (see Figure 1) is a rather slim line, typically the height of a window caption that floats over other applications. The window contains three areas:

- Questions area
- Awareness area
- Buttons area



Figure 1 – Main window of ReachOut

The questions area shows the questions in a fade-in-fade-out manner for minimal intrusion. The last name of the asker appears in brackets next to the question title, along with an icon that indicates the status of the question. The user can set the font and colors (background and foreground) as well as the fading speed. Since the client floats over other applications in the ‘always on top’ manner, we found it important to allow the users to customize the client. We preferred the fade-in-fade-out method to a text ticker with text scrolling from right to left, based on research that has shown this to be less time consuming and more efficient in terms of mental resources [18, 21].

The awareness area displays several gauges that indicate the number of active users and discussions in the system and in the user’s groups. The gauges are “water-filling gauges”, developed by Cohen et al. in the context of awareness on the Web [8]. Using a logarithmic scale and different shades of color, each gauge is capable of supporting a range of 0 – 10,000 values, and may be tuned for larger or smaller ranges (by altering the logarithmic base). While a fuller gauge indicates more units, a darker shade indicates a larger order of magnitude. This enables the gauge to be sensitive to small differences in small amounts, and to large differences in large amounts. Parking the mouse over a gauge opens a ‘tooltip’ with the exact value.

The buttons area supports navigation, by opening a drop down menu from which the user can choose to ask a question, configure the client, or view the help file. The most important step in the configuration, is subscribing to groups of interest, typically done by users when they first log into the system. The current implementation only employs this rather limited profile, not yet taking into account details such as affiliation, location, and other details. We plan to add these in the future, for advanced profile matching.

Subscribing to a group means expressing interest in receiving questions addressed to that group. In the current implementation, the user has to check or uncheck boxes near the names of the groups. This is a binary simplification, using simple ‘yes’ and ‘no’ values, of the ideal model for expressing willingness. The same water gauges described above, provide the user with an indication of the number of people subscribed to each group. The current implementation of ReachOut also employs a hierarchical set of containing groups (see Figure 2) – the higher up the tree, the more people are in the group. This setting allows questions to ‘propagate’ down the tree – a feature that was initially requested by the pilot users, and which, later on, was heavily debated. Our plans for the future are to arrange groups on an affinity-graph, which we believe more aptly mirrors real-life situations.

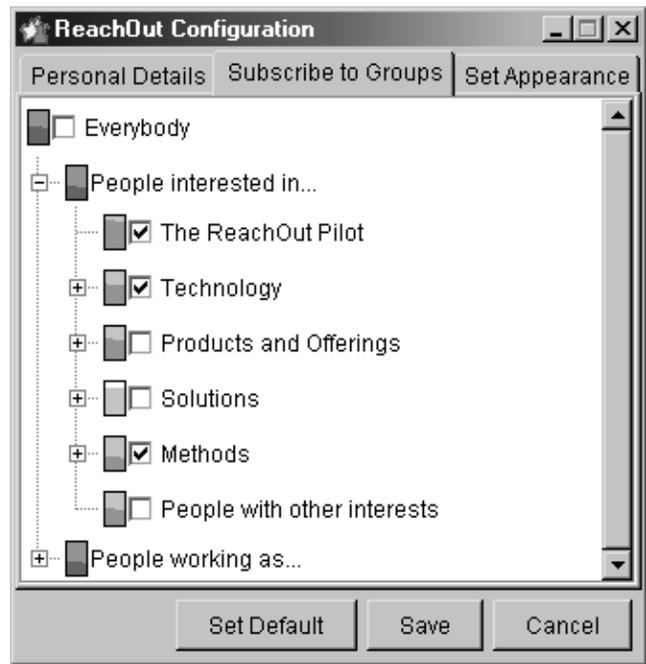


Figure 2 – Subscribing to groups by hierarchy

When users wish to ask a question, they open a ‘two-step wizard’ window that guides them through this simple process. First, they supply a title, which should deliver the essence of the question in a relatively short phrase; this title will later appear on the fader area of potential advisers. Users can optionally supply more details in free text (see Figure 3). These details form the context of the problem and may help potential advisers understand it better.

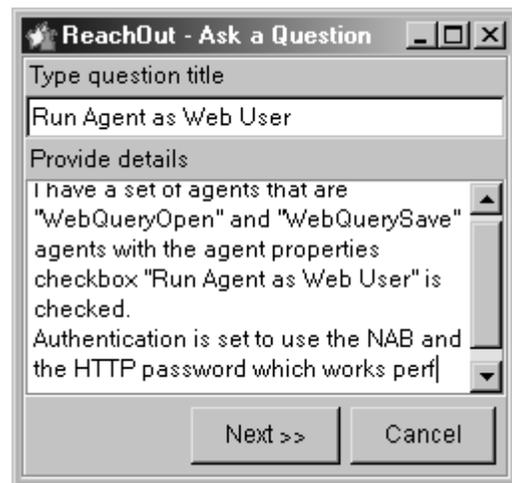


Figure 3 – Asking a question

Next, users must select the target groups for the question. This process is similar to the subscription process, where the user simply checks relevant boxes in the same hierarchy of groups (see Figure 4). The system then sends the question to the people subscribed to the target groups. In the currently deployed release, load balancing is not yet enabled and all questions posted to a group are sent to all

the users who subscribed to it, to increase critical mass in the early stages of deployment.

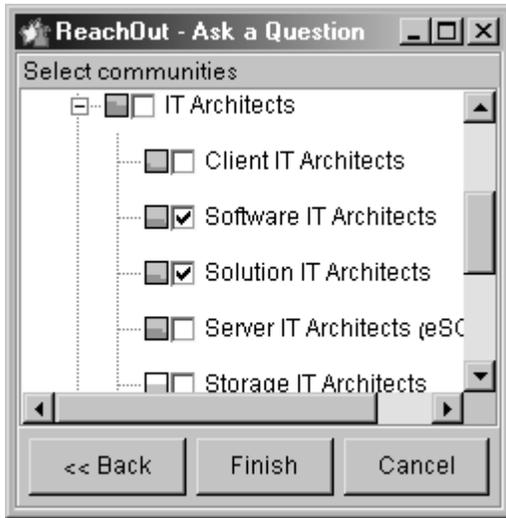


Figure 4 – Selecting target groups

Our use of Sametime™ places, and its publish/subscribe mechanism, introduces a distribution paradigm that is unique to ReachOut. Pure instant messaging systems can push questions only to those currently online; on some occasions, there may be no one online. Posting a question to a newsgroup requires users to actively go and fetch the question. With ReachOut, the question is pushed to each user the next time they go online.

On their ReachOut client application, users see the questions pushed to them, as animated fading in and out messages. Next to each question, an icon indicates the status of the question: ‘new’ (less than 15 minutes old), ‘never opened’, ‘updated’ (since last time the user opened it) or ‘fully read’. This type of grouping, with the ‘my questions’ group and an associated pop-up menu (see Figure 5), allows for quick access to the various discussions.



Figure 5 – Popup menu

This type of quick navigation developed iteratively during the pilot. It started as a request from our users, followed by a cycle of implementation, and was refined using further suggestions from users.

When selecting a discussion, the user enters the discussion’s conference chat room (see Figure 6). This is where ReachOut users, using a simple textual interface,

meet and work collaboratively on a question posed by the discussion initiator.

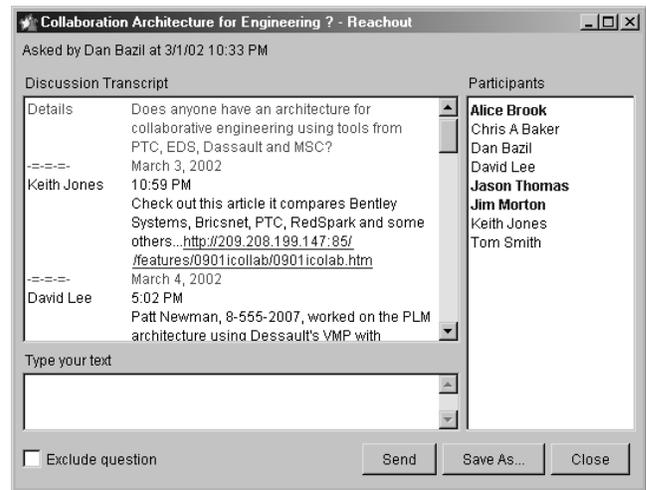


Figure 6 - Conference chat room

This room has all the regular functionality found in other chat applications, as well as some additional features required for the type of blended synchrony [5] offered by ReachOut. As in other chat room applications, users ‘see’ the other users in the room with them; ReachOut, however, also lets them see those who have visited the room, contributed to the discussion, and are no longer in the room. Likewise, the full history of the discussion is shown, regardless of the time the user joined. Timestamps, including date change where needed, are translated to reflect the viewer’s time zone. These features allow ReachOut discussions to persist over days, and ReachOut keeps them open even when there is no user in the discussion room. Discussions are closed by the system 72 hours after the last append is made (i.e., when they cease to evolve).

REACHOUT PILOT

In a pilot program, ReachOut was used by a community of IBM Sales Support people, from January 21 to March 17, 2002. During this period, 130 people – about 25% of the people invited to participate – used ReachOut at least once, and more than 100 questions were posed in different groups. We kept track of the activity on ReachOut using two parallel methods:

- The ReachOut server logged every user action on the system, including users signing in and out, entering different groups, asking questions, participating in discussions, etc. This 4 MB file was then read and parsed by a viewer built using MS-Access™. This allowed us to trace user activity, collect discussion statistics, and monitor the different groups.
- Our team manually reviewed all the chat transcripts, analyzed them, and collected information on the topics discussed.

As indicated by Chart 1, there were at least 10 users online on the system on most days, and between five and 10 users were usually logged in simultaneously.

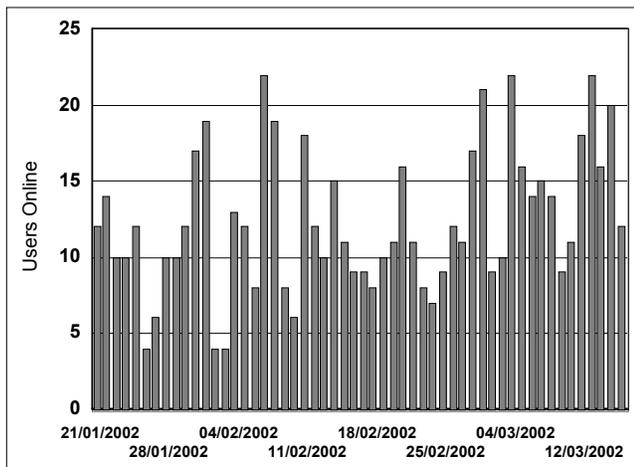


Chart 1 - Users per day

When analyzing the amount of time users spent on the system, we discover that approximately one third of the users used ReachOut for less than 30 minutes. It seems that these users did not see the benefits of the tool during their first use, and did not return. The other two thirds – 73 users – are nicely distributed across a logarithmic scale of time online (See Chart 2), with most of them logging in for four to 16 hours.

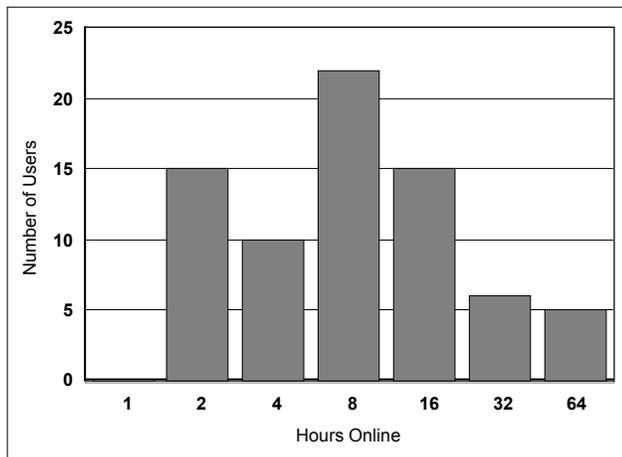


Chart 2- Users by hours online

These figures exclude, of course, the core team members, all of whom logged in for over 128 hours. These findings lead us to believe that we have a valid user base from which conclusions can be drawn regarding the correctness of the ReachOut philosophy and the usability of our tool.

Chart 3 below indicates the time elapsed from when a question was first posted, until the last response was submitted, and therefore excludes the ‘inactivity period’ before the discussion was closed. The chart does not include the 8% of discussions that did not take off at all –

either being too specific (or, in some cases, too trivial) in content, or since they were asked at periods of low activity levels and no one was available to offer answers.

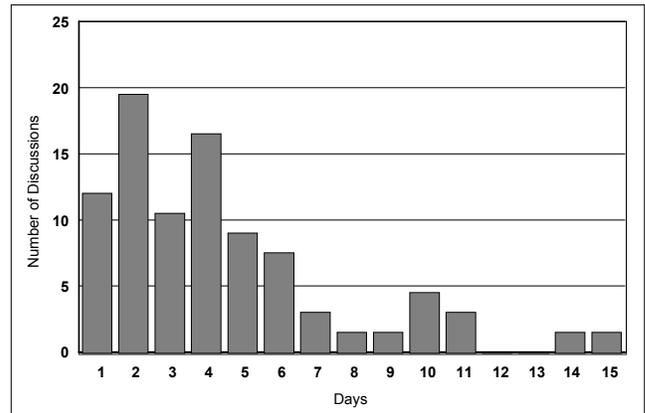


Chart 3 - Discussions by lifetime

The chart shows that 70% of the discussions lasted for more than one day, carrying on way beyond the period of other synchronous-only chats. This, and an average discussion length of 4.1 days, proves the key role that semi-persistence plays in making ReachOut a useful tool.

By analyzing the content of the discussions (see Chart 4), we saw that at the beginning of the pilot, a large percentage of questions focused on ReachOut itself (light-gray, top bars), with a drift towards more ‘business-oriented’ discussions (dark, lower bars) towards the end of the period.

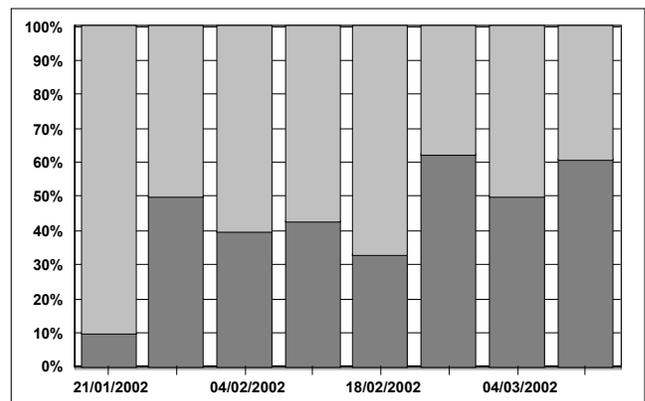


Chart 4: Discussion content – ReachOut-oriented vs. business-oriented discussions

Some of these initial discussions about ReachOut itself, turned into ‘requirement gathering’ sessions, with many new features being suggested, discussed, and refined by our users. This type of activity assisted in building the core community of ReachOut users, turning some of them into our most successful advocates and promoters. ReachOut has unexpectedly proved to be a very suitable environment for collaborative design, supporting ‘structured

brainstorming'. Most of the suggested features focused on the ease of getting and staying online in order to improve the 'stickiness' of the tool and its chances of being adopted. The following suggestions were implemented in the course of the pilot:

1. Automatic registration to the Sametime server the first time a user runs the tool.
2. Leaving the client idle on the screen, even when the network is disconnected. This increases the chances of users renewing login once reconnected.
3. Automatic start of the client at system startup.
4. Moving to a slimmer UI, to increase the chances of users keeping the tool open on-screen.
5. Enhanced audio notification, to eliminate the need for keeping all discussion windows open for update notification.
6. A pop-up menu and new definitions of question status, to support easier navigation and enhance user experience.

Corroborating ReachOut Philosophy

The results of the pilot corroborate many of the points we mentioned when discussing the ReachOut philosophy:

- The 'negotiation of meaning' process, where context is discussed and the real problem is verified, indeed occurred in many cases. Take, for example, the following transcripts¹:

A: Anyone proposing a thin-client solution to the desktop management problem?

B: Which aspects of the thick client solution caused the problem? HW upgrade? SW upgrade? App distribution? Is more control of the display required than a browser allows?

A: How can I find out what the "Big Plays" are for each industry?

After some discussion, asker comments:

A: I see everybody is getting in on the "big play" bandwagon. I think I should have asked for the Industry Value Proposition (IVP) Big Plays.

- We noted that people are indeed willing to share their knowledge, even without being materially compensated. We found that on average, each discussion had 3.8 participants, and each user participated in 5.3 discussions.
- The role of ReachOut in expanding social networks is well exemplified by 'referral' discussions, in which users act as 'connectors', suggesting others as a source of information. 18% of the business-oriented

discussions were referrals. The following transcripts are examples of this:

A: I am looking for a security reference architecture for wireless devices

B: Bob, Try Mike Aldridge in the UK. I will alert him as well.

A: Where do I find an expert on EJB's on 390?

B: Have you tried the eTZ folks in POL? I think Dan Bazil is the manager.

C: You can try Joe White out of Orlando. He may be able to assist depending on your need.

D: I do know that Diana Cole did run a JAVA competency group. She would probably be able to point you in the right direction.

E: The guy we use in the UK is Peter Benson, based in the Netherlands.

- It is quite obvious that valuable information about the asker is provided, even without reading the discussion. The following question titles carried a lot of 'extra' knowledge about the askers, their working environments, and missions:

- I am looking for a security reference architecture for wireless devices
- How to get Tivoli Software Package Editor to work on AIX
- Expertise Location Projects - looking for comparisons

- Finally, maintaining open discussions not only proved to be helpful in finding the right answer to the question, in some case discussions turned into real brainstorming sessions, with five, eight, and even 12 contributing participants, with many others 'watching'. Many ideas were triggered when users read other people's answers.

While the pilot results were encouraging, some shortcomings were evident:

- Only rarely throughout the pilot did ReachOut cross the 'critical mass' line, and its support for real-time, synchronous discussion was thus not fully exploited.
- It seems that some users tend to log in to ReachOut only occasionally, and thus get to see and ask questions, but are not back before those discussions cease to exist. If too many people follow this pattern, our semi-persistence philosophy may need to be revisited. Based on the pilot, we can assume these are not the majority of our users.
- Although ReachOut actively assisted people in getting answers and brainstorming ideas, the sense of a strong community still covered only a small circle of the very active users. It is likely that a longer period of time is needed to expand this circle.

¹ All personal names and brands were modified for privacy reasons.

CONCLUSIONS AND FUTURE WORK

We presented the methodology, design, and implementation of ReachOut, and analyzed its usage by conducting a two-month pilot with 130 participants. ReachOut provides the peer-support features of newsgroups, while overcoming the problem of information overload by utilizing push technology and providing immediate awareness to discussions. Its use of 'profile matching' directs questions only to those who have expressed interest, automatically notifying them to free the asker from addressing impartial strangers. By turning from experts to peers, it assists in expanding social networks and in building a community. The fader area, where question titles and names of askers are shown, provides constant awareness of the heartbeat of that community. The public conference-chats provide a useful stage for collaborative problem solving, and semi-persistence allows discussions to span over days, with users joining in at different times.

The pilot results confirmed most of our theoretic postulations, and the extensive usage of the tool has led to many improvements in its implementation. The pilot proved ReachOut's potential and usability in dealing with real world problems, and we found some unexpected uses for it in collaborative design.

Future research plans focus around the notion of 'people mining', in which finding the best potential advisers also takes into account the asker-to-advisers relations. We will also look into the factors behind successful and unsuccessful disseminations, to guarantee adoption and the achievement of a critical mass. Development plans include incorporating ReachOut with business processes and persistent repositories, and improvements to the client, allowing it to work on handheld devices and in intermittently connected modes.

ACKNOWLEDGMENTS

We are grateful to Sigalit Ur and Noga Meshulam for the implementation of the tool, and to Raanan Nevet for the implementation of the log analyzer. We also thank Yoelle Maarek for her support and brainstorming sessions. Dick York and Mark Palmer deserve special thanks for their enthusiasm and dedication, without which the pilot would not have been the success that it was. We also thank IBM Global Technical Sales Support for letting us run the pilot, and of course, the pilot participants themselves, for their time and effort.

REFERENCES

1. Ackerman, M. S. (1994) "Augmenting the Organizational Memory: A Field Study of Answer Garden". *Proceedings of ACM Conference on Computer Supported Cooperative Work (CSCW '94)*, New York: ACM Press, 243 - 252.
2. Ackerman, M. S. and McDonald, D.W. (1996). "Answer Garden 2: Merging Organizational Memory with Collaborative Help". *Proceedings of ACM Conference on Computer Supported Cooperative Work (CSCW '96)*, New York: ACM Press, 97-105.
3. Ackerman, M.A. and McDonald, D.W. (1998) "Just Talk to Me: A Field Study of Expertise Location", *Proceedings of Conference on Computer Supported Cooperative Work (CSCW'98)*, New York: ACM Press, 315-324.
4. Ackerman, M.S. and Halverson, C. (1998). "Considering an Organization's Memory", *Proceedings of ACM Conference on Computer Supported Cooperative Work (CSCW '98)*, New York: ACM Press, 39-48.
5. Bradner, E., Kellogg, W, and Erickson, T. "The Adoption and Use of Babble: A Field Study of Chat in the Workplace". *Proceedings of the European Computer Supported Cooperative Work (ECSCW '99)*. Dordrecht: Kluwer, 139-158.
6. Carroll, J.M. and Rosson, M.B. (2001). "Better Home Shopping or New Democracy? Evaluating Community Network Outcomes", *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI 2001)*, New York: ACM Press, 372-379.
7. Chalmers, M. and Dourish, P. (1994). "Running out of space: models of information navigation" in *Proceedings of HCI '94*, Glasgow, 1994. Available at <ftp://parcftp.xerox.com/pub/europarc/jpd/hci94-navigation.ps>.
8. Cohen, D., Jacovi, M., Maarek, Y.S. and Soroka, V. (2002). "Livemaps for Collection Awareness", *International Journal of Human-Computer Studies*, Vol. 56, No. 1, 7-23.
9. Davenport, T.H. and Prusak, L. (1998). *Working Knowledge*. Boston, MA: Harvard Business School Press.
10. Dourish, P. and Bellotti, V. 1992. "Awareness and Coordination in Shared Workspaces". *Proceedings of ACM Conference on Computer-Supported Cooperative Work (CSCW'92)*. New York: ACM Press, 107-114.
11. Google Groups. <http://groups.google.com/>
12. Greer, J.E., McCalla, G. I., Cooke, J. E., Collins, J., Kumar, V. S., Bishop A. and Vassileva, J. I., "The Intelligent Helpdesk: Supporting Peer-Help in a University Course", in *Proc. of the International Intelligent Tutoring Systems Conference (ITS'98)*, San Antonio, TX, 1998, 494-505.
13. Hall, R. W., Mathur, A., Jahanian, F., Prakash, A. and Rassmussen, C. (1996), "Corona: A Communication Service for Scalable, Reliable Group Collaboration System", *Proceedings of*

- ACM Conference on Computer Supported Cooperative Work (CSCW '96)*, New York: ACM Press, 140-149.
14. KnowledgeMail. <http://www.tacit.com/products/knowledgemail.html>
 15. Leonard, D. and Sensiper, S. (1998). "The Role of Tacit Knowledge in Group Innovation", *California Management Review*, 40 (3), 112-132.
 16. Lotus Discovery Server (LDS). <http://www.lotus.com/lds>.
 17. Lotus Sametime Server. <http://www.lotus.com/sametime/>.
 18. Maglio, P.P. and Campbell, C.S. (2000). "Tradeoffs in displaying peripheral information". *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2000)*, 241-248.
 19. Malone, T.W., Grant, K.R., Turbak, F.A., Brobst, S.A. and Cohen, M.D. (1987). "Intelligent Information Sharing Systems". *Communications of the ACM*, 30, 5, 390-402.
 20. Mattox, D., Maybury, M., and Morey, D. (1999). "Enterprise Expert and Knowledge Discovery". *Proceedings of the 8th International Conference on Human-Computer Interaction (HCI International'99)*, Munich, Germany, 303-307.
 21. McCrickard, D.S., *Maintaining Information Awareness in a Dynamic Environment: Assessing Animation as a Communication Mechanism*. PhD thesis, Georgia Institute of Technology.
 22. Mynatt, E., Adler, A. Ito, M. and O'Day, V. (1997). "Design For Network Communities". *Proceedings of the Conference on Human Factors in Computing Systems (CHI'97)*, 210-217.
 23. Nardi, B., Whittaker, S., Bradner, E. (2000). "Interaction and Outeraction: Instant Messaging in Action". *Proceedings of ACM Conference on Computer Supported Cooperative Work (CSCW '00)*, New York: ACM Press, 79-88.
 24. Pegoraro, R. (2000). "Logging On: Is Usenet Becoming Yesterday's News?" *Washington Post*, Internet edition, 4 February 2000. Downloaded from <http://www.washingtonpost.com/wp-srv/WPlate/2000-02/04/0661-020400-idx.html> on 3/21/2002.
 25. Polanyi, M. (1966), *The Tacit Dimension*. New York: Doubleday & Company Inc., 1966.
 26. Ramduny, D., Dix, A. and Rodden, T. " Exploring the design space for notification servers", *Proceedings of ACM Conference on Computer Supported Cooperative Work (CSCW '98)*, New York: ACM Press, 227-235.
 27. Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P. and Riedl, J. (1994). "GroupLens: An open architecture for collaborative filtering of netnews". *Proceedings of ACM Conference on Computer-Supported Cooperative Work (CSCW'94)*, New York: ACM Press. 175-186.
 28. Resnick, P. and Varian, H.R., "Recommender Systems", *Communications of the ACM*, 40 (3), 56-58.
 29. Rheingold, H. (2000), *The Virtual Community: Homesteading on the Electronic Frontier*. Cambridge, Mass: MIT Press.
 30. Shardanand, U. and Maes, P. (1995). "Social Information Filtering: Algorithms for Automating 'Word of Mouth'", *Proceedings of the Conference on Human Factors in Computing Systems (CHI95)*, New York: ACM Press, 210-217.
 31. Teasley, S, Covi, L., Krishnan, M. S. and Olson, J.S. (2000). "How does radical collocation help a team succeed?" *Proceedings of ACM Conference on Computer Supported Cooperative Work (CSCW'2000)*, New York: ACM Press, 339-346.
 32. Terveen, L. and Hill, W. (2001). "Human-Computer Collaboration in Recommended Systems". In Carroll, J. (ed.) *HCI In the New Millennium*, Addison-Wesley, 486-507.
 33. Terveen, L., Hill, W., Amento, B., McDonald, D., and Creter, J. (1997). "PHOAKS: A system for sharing recommendations". *Communications of the ACM*, 40(3), pp.59-62.
 34. The WELL, <http://www.well.com/>.
 35. Whittaker, S. (1996). "Talking to Strangers: An Evaluation of Factors Affecting Electronic Collaboration". *Proceedings of the 1996 ACM Conference on Computer Supported Cooperative Work (CSCW'96)*, Cambridge MA, 409-418.