Opportunistic Interfaces for Promoting Community Awareness

A Thesis
Presented to
The Academic Faculty

by

Qiang Alex Zhao

In Partial Fulfillment
of the Requirements for the Degree of
Doctor of Philosophy in Computer Science

Georgia Institute of Technology
August 2001

Opportunistic Interfaces for Promoting Community Awareness

Approved:

John T. Stasko, Chairman

Gregory D. Abowd

Amy S. Bruckman

Mark Guzdial

Scott Hudson

Date Approved

Acknowledgements

A number of people have influenced this work and my life as a student at Georgia Tech. I would like to use this opportunity to express my gratitude.

First, my advisor, Dr. John Stasko encouraged and guided me through the many adventures in the development of this work as well as other related projects. He pointed to me new directions to explore and allowed me to investigate wide variety of ideas while still keeping me focused on the most important aspects of the research. His inspiration and support has made this work possible, and helped me utilize my strengths and strengthen my weaknesses on my journey to become a better researcher.

I would also like to thank other faculty members with whom I had the pleasure to meet at Georgia Tech. In particular, my other committee members, Drs. Gregory Abowd, Amy Bruckman, Mark Guzdial, and Scott Hudson have all offered suggestions, critiques, reviews, and advice on this research as well as the preparation of this dissertation. In addition, Dr. Beth Mynatt provided her expertise on the design and evaluation issues in this work. Drs. Mostafa Ammar, Jessica Hodgins, and H. Venkateswaran all gave encouragements and valuable suggestions.

Many of the College of Computing support staff have helped me carry out this work as well as other activities in my student life. Dan Forsyth, Randy Carpenter, Karen Carter, Keven Haynes, Austin Hijar, Steve Park, Peter Wan, and Jonathan Shaw have made computer related problems seem trivial despite my habit of causing them. Cathy Dunnahoo,

Barbara Durham, Joan Morton, Chrissy Hendricks, and Tempo Tinch have gladly and efficiently helped me with paperwork despite my obvious clueless-ness in this area. And thanks to Leash Chappell, Linda Dillon, Carolyn Maddox, Rashmi Patel for furnishing me with information that I needed in the process of this work.

A number of my fellow students provided intriguing ideas that provoked thoughts on many subjects and diversions that kept me sane. In particular, Scott McCrickard gave me many useful advice and critical suggestions that I have been trying to learn from. Despite my initial problem with his accent, his humorous ways always brought smiles to my face. I would also like to thank people in the Software Visualization then later Information Interfaces group as well as those in our neighboring labs for discussions about almost everything, and for constantly asking me "have you finished your thesis, yet?" Thanks to students who have volunteered in College events and grad-teas for providing examples of inspiration, especially those who participated in the various projects in this research. Also thanks to my friends outside of College of Computing for the moral support.

Finally, I dedicate this dissertation to my family — to my wife, Shumei Lin for tolerating me with my extended time spent on computers and supporting me to attempt a doctorate degree; to our son, Bryan YuanSong Zhao for bringing us the joys of having him in our lives; to our sisters and brothers and our parents for helping us and supporting us throughout the years.

Qiang Alex Zhao
Atlanta, GA
August 16, 2001

Table of Contents

List of Tables

List of Figures

Summary

Community awareness is the general knowledge about community members, the state of the community, as well as the norms and roles in the community. In an academic research setting, maintaining an adequate level of community awareness helps build and sustain social networks, eases information sharing, and facilitates the collaboration in creative work. This general awareness often lacks concrete details, rarely affects people's immediate activities, and is poorly supported by current technologies that often presume established goals and existing contexts of interpersonal interactions. Opportunistic interfaces, on the other hand, convey information by chance, often unobtrusively in the periphery of attention, and help people take advantage of such opportunities. This thesis explores opportunistic interfaces as an alternative technique in fostering community awareness.

Specifically, we assessed awareness issues in a real-world academic research community and experimented with popular ways of enhancing awareness and interpersonal interaction. Based on our findings, we devised a set of objectives of community awareness applications and iteratively developed two opportunistic interfaces that demonstrated the trade-offs between small, corner-of-the-desktop interfaces and large, full-screen displays. Finally, our evaluation of these example applications helped formulate recommendations for integrating opportunistic interfaces to promote community awareness and communication.

CHAPTER 1
Introduction

Communities operate on different granularities and under different circumstances [105]. Generally speaking, a community is a group of people associated in a certain way. People can associate by living in the same residential area, attending the same school, being interested in the same topic, taking part in the same organization, or engaging in related professional activities. The size of a community can be as small as several people, or as large as the global population on this planet. Even though a person may associate with many different communities under different contexts, each particular community helps define a certain aspect of that person's identity. At the same time, members of a community collectively shape the identity of the community.

Community can also describe the level of cohesiveness that a group of people experience. For example, *community awareness* often refers to the degree that people generally know about each other, about issues that affect the community, and about social norms as well as people's different roles within the community. When members know the community well, orient toward each other and the group as a whole, and feel a strong sense of attachment to the group, the *sense of community* is achieved.

The benefits of communities are sometimes discussed through the concept of *social capital*. Similar to physical resources, the relationship between one person and other people has value. Social capital is the collective value of such relationships as well as the level

of tolerance and trust toward each other, and the inclinations that arise through generalized reciprocity [21]. Social capital accumulates primarily through interpersonal interactions such as talking to each other and doing activities together. It is also an enabler for future interactions that in turn may produce even more social capital. Conversely, social capital may dissipate over time if it is not reinvested or if it is spent in bad interactions.

A level of community awareness is gained and maintained through information exchange, and is a fundamental aspect of the social interactions that produce social capital. For example, knowledge about people's interests and what may affect them facilitates information routing, which in turn facilitates sharing and exchanging resources other than information. Knowledge of each other, combined with adequate level of trust and expectation, helps people provide emotional support to each other. Knowledge of each other's ongoing activities facilitates coordination of interdependent actions. And finally, knowing people's opinions and intentions, as well as community roles and norms, helps members overcome collective action dilemmas such as social mobilization[1].

Putnam demonstrated in his recent book *Bowling Alone* that, in high awareness and high social-capital communities, members often experience better health, better education and economic outcomes, and of particular interest to this thesis, better collaboration and sharing [93]. These benefits are not limited to those who actively participate in the cummunity. Everybody, including those who never participate and only marginally relate to the community may benefit from being a part of the community. Additionally from an individual person's point of view, even though one's relationship with the general commu-

1. The social mobilization dilemma is when no one acts on an issue even though people know that the community can benefit greatly if people act together. Individually, they all fear that others will not join and a single person's effort may not be worthwhile.

nity is usually not as strong as that with immediate colleagues and family members, maintaining the "weak-ties" with the rest of the community has many benefits such as preparing for future "strong-tie" relationships and accessing to locally unavailable but important information [91].

In this thesis, we focus specifically on the community of academic research — a group of people working in the same academic organization, often at the same locality, sharing similar interests or goals in scientific research. Since an academic research community is an environment for creative work, and creative activities in such organizations are mainly collaborative [111], a cohesive and healthy community nourishes collaborative activities and is potentially more productive and satisfying than otherwise [95].

Academic research organizations are facing several community issues today. First, the rapid growth that some communities are celebrating may bring extra burdens to the maintenance of community awareness. In a large or geographically separated organization, social as well as professional interactions with other people are relatively more difficult than in smaller, collocated communities, and therefore likely to be less frequent. The quality of interpersonal interactions, and consequently the restocking process of social capital, may suffer. Without adequate level of knowledge about the community that matches the growth, a person may have difficulties in adapting to the changes in the community and may not be able to take advantage of the expanding possibilities of collaboration.

Second, although an increasing large percentage of information about individual communities is becoming available on-line, in electronic forms that enable faster access, it is becoming more difficult to notice and process such information due to our inherent attention limits [112] and the overwhelmingly large amount of related as well as unrelated

information that is available to us [69]. People may not know the existence of potentially important information or they may not know how to find that information. Moreover, people are sometimes pressured to decide quickly what available information they should pay attention to and what they should not, potentially ignoring information that help maintain awareness and build social capital.

Finally, the continuous declination of social capital in virtually all aspects of the American society since the mid-1970's [93] also casts a shadow on community awareness. Being an important component of the civic society, academic research communities in this country have the responsibilities of investigating socio-technical answers to rejuvenate the sense of community in the greater society. In particular, as human-computer interaction researchers, we need to explore how technologies can help enhance or create social practices that fit people's current lifestyles [94].

Few technologies have been examined in helping people maintain community awareness. A large body of research has been focusing on supporting distinct awareness that is directly related to collaborative tasks in organized teams with common goals. These systems often shorten the physical distances between team members by providing virtual presences of people, enable them to contribute to the targeted products across time boundaries, and offer fine-grained control of their collaborative activities. However, these technologies demand full attention from the user and participation requires the user to spend a considerable amount of effort. Since community awareness is usually only loosely related to any specific collaborative tasks, if at all, and the benefits are rarely immediate, using these current technologies to maintain awareness at the community level becomes distracting, sometimes even overwhelming.

Unlike solution-seeking *groupware* systems, as those technologies are often referred to, the notion of *communityware* emphasizes on the discovery of knowledge and people in a diverse and amorphous community [52]. In particular, our research under the communityware umbrella focuses on how technologies can help people obtain as well as sustain a certain level of general awareness about the community and make better decisions in selectively paying attention to the vast amount of information that flows through the physical as well as the virtual work space.

More specifically, we explore an alternative design philosophy that aims to create appropriate opportunities for lightweight information exchange and allow people to easily take advantage of these opportunities. We use the term *opportunistic interfaces* to represent such designs. They present information in a timely manner without distracting the user's attention away from the tasks at hand. Furthermore, the user accesses the information by chance where the conditions for such chances to happen are well understood between the human and the machine, so that an instance of such opportunity is minimally intrusive to other tasks. Event though the application has a certain level of autonomy in deciding the information being shown, and therefore the exact piece of information presented to the user at the time of access is not always deterministic, the user can make sense and make use of the information with minimal amount of effort when such chance arises.

As such, it is often desirable to put community information at the periphery of attention first. In the physical world, for example, flyers posted on telephone poles and billboards set up along highways let people who pass by to process the information at their leisure, when they have the time and capacity. From the information consumers' point of view, they do not necessarily know what information they will see when they pass the fly-

ers or billboards. While information on those flyers or billboards can usually be safely ignored, this medium can heighten a person's awareness about a certain issue when the information is absorbed.

To summarize, opportunistic interfaces support opportunistic presentation of information and opportunistic access to information. This thesis seeks to show that computer supported opportunistic interfaces can provide lightweight alternatives in facilitating informal information gathering and exchange in academic research communities. This research contributes to the field of human-computer interaction by identifying opportunistic interfaces as a necessary design focus and by exploring such systems in promoting community awareness. Through the iterative process of developing two example applications, we will demonstrate the key challenges in designing opportunistic interfaces. From these exercises, we may begin to build a basic understanding of the role of opportunistic interfaces in communityware applications.

During the early stages of this research, we accessed the status of a real world academic research community, analyzed important factors that contributed to people's perception of the community, and characterized ways that could enhance awareness and improve communication. Extending from and integrating with existing technologies in supporting presence and conversation, we experimented with linking remote spaces together via constant "windows" of video and audio. Chapter 2 illustrates our assessments of the community and analyses of the issues and options. It also describes the implementations the video and audio windows projects as well as our observations in their use.

Evaluating existing research helps us better understand the problem domain and provides a theoretical background for alternative approaches. Chapter 3 surveys related

awareness and communication technologies, motivates the need for opportunistic interfaces, and develops a concrete set of objectives to guide our design of such systems.

An application at the periphery of the computer desktop is one of the possibilities to study opportunistic interfaces. We developed the "What's Happening" communication-bar that sits in a small area on the computer display, automatically presenting useful information without trying to grab the user's attention. Chapter 4 explains the operation and construction of this application, reports the design choices and trade-offs that we made, and shows how certain aspects of the application evolved in the design process.

While the communication-bar affords seamless switching between being in the focus of attention and being in the periphery, it inevitably competes for screen real-estate with applications that support more urgent or demanding tasks. The "What's Happening" screen-saver trades off this flexibility for the control of the entire screen, with the benefit of affording quicker information consumption through images. Chapter 5 provides a motivation and describes the implementation of a community awareness screen-saver, as well as the rationales behind our design choices.

Measuring the effects of opportunistic interfaces is an important, yet difficult step in this research. Chapter 6 outlines several challenges in evaluating community awareness applications as well as methods that we chose to use for gaining a good understanding of the usage and impacts. We also describe the details of our findings and what we can learn from the available data.

Finally, Chapter 7 summarizes of the lessons learned from the "What's Happening" applications and the research contributions of this thesis, and proposes several directions of future work in exploring opportunistic interfaces for promoting community awareness.

CHAPTER 2
Supporting Work

This chapter examines supporting work related to this thesis. We will start by assessing the state of in a real-world community and this will provide a backdrop for our research. Then we will describe two early projects in enhancing awareness and fostering communication, as well as our observations and results from these studies.

2.1 Sense of Community

The College of Computing at Georgia Tech is a fast growing scientific research community. It houses more than 60 faculty members and over 200 doctorate students, as well as a large number of Master's and undergraduate students, and staff members. As the community grew larger, the home building that was originally designated to the College (CCB) could no longer accommodate the entire community. To alleviate the space problem, the College had to move into two other locations also: the Centennial Research Building (CRB), and the Georgia Center for Advanced Telecommunication Technologies (GCATT).

This geographic separation seemed to contributed to people feeling an emerging sense of distance and unfamiliarity. Although CRB was only a short walk away from CCB and GCATT a short drive away, meeting with people in a different building was still viewed often as inconvenient. In addition, to better utilize available space, many people had to

relocate, some even multiple times. In an e-mail titled “Deteriorating Conditions at the College”, one graduate student said: “I feel that the College is no longer a real community. Instead, we're turning into fiefs and isolated research groups that barely encounter each other.” This e-mail stimulated discussions about what constitutes a community and why we needed a coherent community. Many people agreed that we need ways to improve the sense of community, even though a certain amount of separation was inevitable given the amount of growth that have occurred.

2.1.1 Survey Inquiries

To better understand people’s opinions about community issues, we deployed a written survey to nine graduate students, seven faculty members, and two staff members. We asked them to describe the sense of the community in the College, and whether they felt it needed to be improved. We also asked if in their work or study, they encountered any difficulty related to the physical separation of the College.

Few people, one doctorate student and two faculty members, responded that we had a good sense of community in the College and we did not need any improvement. Others agreed that there was room for improvement in the sense of community within the College. Most people said that the community was fragmented. They did not see people in other buildings as much as they would like to, and they did not know much about the research work being conducted by groups in other buildings.

Although the physical distances among the three buildings brought inconveniences and extra burdens, people consistently worried more about the reduced chances for casual interpersonal interactions. The trend to a growing number of isolated labs was worsening

the problem. People wanted more opportunities to talk with others outside of work related formal settings.

2.1.2 Interviews

We followed up with seven of the survey respondents by conducting informal interviews to gain a deeper understanding of workplace community issues. First, we asked them to define “the sense of community”. Consistent with definitions in the research literature [115], people described the sense of community as the feeling of social comfort, enjoyment, and companionship in the environment, the respect and trust toward other community members, and the sense of belonging to the community. They believed that characteristics of an organization exhibiting a strong sense of community would include:

- Members of the community generally know each other well. For example, one knows another person's research interests as well as hobbies and favorite restaurants.
- People are more likely to engage in spontaneous social interactions, such as casual chats in the hallways.
- People are more likely to share resources, such as computing facilities and time.
- People are more likely to attend community activities, such as seminars and parties.
- People are more likely to offer help in community activities and initiate efforts hoping to improve community life.

Generally being aware of artifacts and others in the environment was considered to have a number of benefits to people’s “quality of life” in the community. For example,

people felt more motivated in their work or study when they knew that someone else was doing the same at the same time, either in the same physical space such as a shared lab or the library, or in the same virtual space on-line. Chatting with other people during the break times in between work or study periods was relaxing and helped relieve stress. Knowing about other people also helped finding solutions to technical problems and setting up contacts for future collaborations. Knowing about other people's success motivated one's own work and built the impression that the overall community was successful. Above all, people felt that belonging to a special place with a strong sense of community would bring pride and would be quite rewarding.

However, people expressed some important concerns as well:

- Spontaneous interpersonal interactions were too often confined to individual buildings or individual research groups. A few people only interacted within small groups whose offices happened to be close by, and were not interested in other aspects in the College. In general, there were less inter-group interactions, less resource-sharing, and less variety in generating problem solving ideas.
- People were sometimes too busy for casual interaction opportunities. At the same time, those opportunities were becoming less frequent due to the separation of the labs and buildings. Although this situation might be less distracting and productivity might increase, people did not feel as fulfilled as they would like to.
- Some community events were poorly attended and it was difficult to solicit volunteers for College functions, especially graduate student volunteers.

For example, the graduate student council had long been looking for a volunteer to publish the graduate student newsletter. Even when a volunteer who helped create the one and only issue of the newsletter appeared, few people contributed stories to it. Another example was the grad-tea event: the College and industrial partners sponsored space and food for weekly tea-times, trying to encourage social interactions within the community. Each grad-tea needed volunteers to set up the tables and transport the food. Signing up volunteers was so difficult that the College had to cancel some of these events.

- Being unfamiliar with other people's work was considered a growing problem. More and more members of the community were becoming strangers.

People speculated that increased pressure in work or study at the College might also be contributing to these problems. Compared to the increasing size of the community, any individual's effort might be perceived as being too small to be useful. People might think that others would volunteer to do the work. They might not know what concrete ways that they could help because they were not aware of other people's volunteering effort. And lastly, they might not want to get involved in volunteering unless they could see it as beneficial to themselves. These feelings were hazardous to the sense of community in the College.

2.1.3 Questionnaire

To find out on a larger scale, how community members viewed the sense of community, we deployed an e-mail questionnaire to the faculty and graduate students (see Appendix A). We asked people to rate their familiarity with both research and social events in the

College, as well as their frequency of attending these events. In addition, we asked people to rate their familiarity with research work in groups other than their own, and frequency of interaction with those groups in academic and social settings. To obtain an estimate of people's satisfaction with the current state of the community, we also asked people to rate the sense of community in the College.

The ratings were on a 7-point scale with 7 representing the extreme in favor of a certain quality, for example, "very familiar" and "very good", and 1 representing the opposite extreme such as "very unfamiliar" and "very bad".

We received 59 responses to the questionnaire. Table 1 lists the averages of answers to the rating questions, and Figure 1 shows the frequencies of the answers.

Table 1: Summary of answers to rating questions in the initial community survey. Scores are given on a 7-point scale. A score of 1 represents "very unfamiliar", "rarely", or "very bad". A score of 7 represents "very familiar", "frequently", or "very good". The neutral score is 4.

Question	Average Rating	Standard Deviation
1) Familiarity with research events	4.8	1.65
2) Attendance of research events	3.5	1.42
3) Familiarity with social events	4.8	1.53
4) Attendance of social events	3.4	1.51
5) Familiarity with research in other groups	3.3	1.41
6) Academic interactions with other groups	3.1	1.68
7) Social interactions with other groups	3.7	1.83
8) Sense of community in the College	4.0	1.37

People had better than neutral satisfaction with their familiarity with research and social events, but reported less than neutral frequency of attendance. They were less than

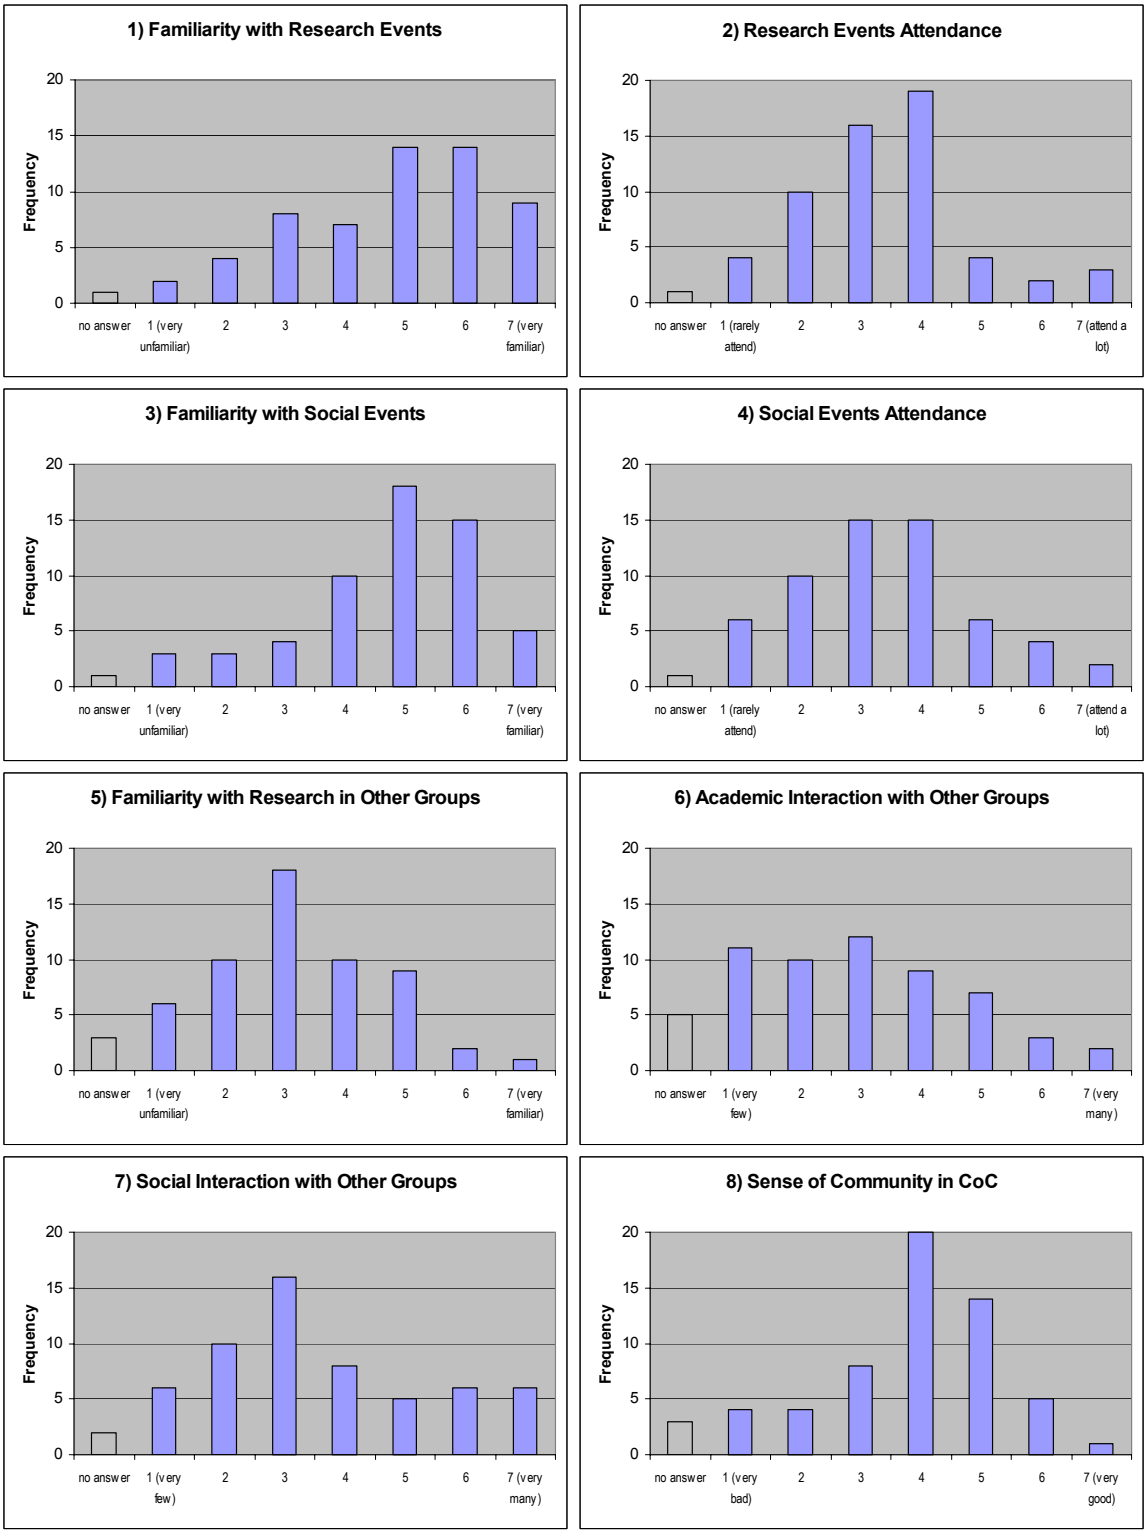

Figure 1. Histograms of answers to the rating questions in the initial community survey.

neutral familiar with research in other groups, and on average, they did not frequently interact with people in other groups, either. Overall, people rated the sense of community in the Collage a mediocre 4.

In the free-form comment section of the questionnaire, we asked people to identify the main contributing factors to the sense of community and provide ideas and suggestions on potential changes and improvements. One person commented that people should only do their jobs, accept what the community was and how it might change, and stop being concerned about the community. However, many other people recognized that a common sense of purpose and responsibility were important to the community. Another person wrote:

> One of the most important factors that contributes to a strong sense of community in the CoC is a common sense of responsibility for what we have now, and for the development of what we would like to have in the future. We must have an interest in participating in both social and academic activities with our peers, not because it directly benefits our careers/research/happiness, but because it can make the special experience of attending graduate school more enjoyable for everyone.

The majority of the answers from other respondents were consistent with our findings in previous surveys and interviews. The main contributing factors to the sense of community included “people knowing about other people's interests”, “opportunities to interact”, “common societal interests and values”, and “an interest of creating community”. People identified the physical separation and the lack of concern about the community as the most damaging factors to the sense of community. Many people expressed a desire of bringing the whole College back into one location. Others suggested “more realistic” options such as sponsoring informal social events and creating spaces where people would have the

opportunity to “bump into” other people and start a conversation. A few people proposed weekly or monthly community newsletters and on-line forums for open discussions.

In conclusion, our interviews and surveys agree with previous research on the relationship between physical proximity and informal communication. Loosely speaking, *informal communication* is the casual, interpersonal interaction of exchanging thoughts, messages, or information. It often starts by chance, such as “bumping into” another person in the hallway, and is usually shorter than formal meetings. In addition to helping resolve conflicts and coordinate activity under the conditions of uncertainty [63], informal communication helps people get to know others in the community, recognize other people's skills and achievements, build perceptions of others and relate to them, and learn the norms and culture in the work environment [114]. In essence, informal communication is an important vehicle for maintaining awareness and creating social capital.

Physical proximity is the predominant factor influencing the frequency, quality, and cost of informal communication. People located within a convenient physical distance have more opportunities to engage in low overhead, short, but frequent casual interactions while trying to accomplish other tasks. Consequently, they are more likely to collaborate with each other [62]. Without adequate physical proximity, the frequency of informal communication with ideal collaborators may decrease significantly, and as our data suggests, people may feel less connected to the rest of the community. The sharing of information and other resources may occur less often and many collaboration opportunities may be missed.

There was a consensus that we needed to care more about each other and the environment that we shared. We needed ways of getting to know about each other better, reaching

out to everybody in the community, and making informal communication more common, more convenient, as well as fun and useful. While we could not easily solve the problem of disappearing physical proximity, we could try to enhance awareness by exploring other means of communication.

However, people were cautious about community related communications — they did not want to be overwhelmed by the increased amount of information, even though they wanted better visibility of their own work within the community, as well as other people's work. Since people were reluctant to learn a technology unless they could benefit from it frequently, a tool solely for communication purposes might not find many users in this community because many other options already existed. Something that is lightweight, requires no training, and is integrated into what people are already using may have a better chance of becoming effective in enhancing awareness and improving communication.

2.2 Early Projects

Borrowing ideas from early media space systems, we conducted two projects to help people become more aware of others and to create opportunities for them to interact.

2.2.1 The Video Wall

Two previous experiments explored linking public spaces across distances with limited success. Xerox Palo Alto and Portland researchers used a continuously running video conferencing setup to connect conference rooms and offices between those two locations [86]. The VideoWindow system linked two separate rooms using large screen video and positional audio connections to make meeting with remote colleagues as easy as if the spaces were only separated by a sheet of glass [30]. To evaluate the possibility of using current,

more advanced teleconferencing technologies to support chance encounters and impromptu interactions, we set up real-time audio and video links between two common areas in CCB and CRB.

We installed an ATM audio/video transmitter and receiver in each building (see Figure 2). The transmitter digitized analog audio/video signals, and encoded them into ATM packets that were transmitted onto the network. The receiver decoded the packets that it received from the ATM network and converted the digital signals to analog. At each site, local audio and video data were captured by a microphone and a video camera, respectively, and sent to the network via the transmitter. Remote audio and video data were played back through a loudspeaker and displayed on a TV monitor at the same time.

In each of the two buildings, one set of video wall equipment was placed in a common area where lunches and casual meetings often took place. We posted signs near the TV monitors, stating that the setup was a feasibility test and would last about a month. In addition, we asked that the audio and video link be kept on whenever it did not interfere with

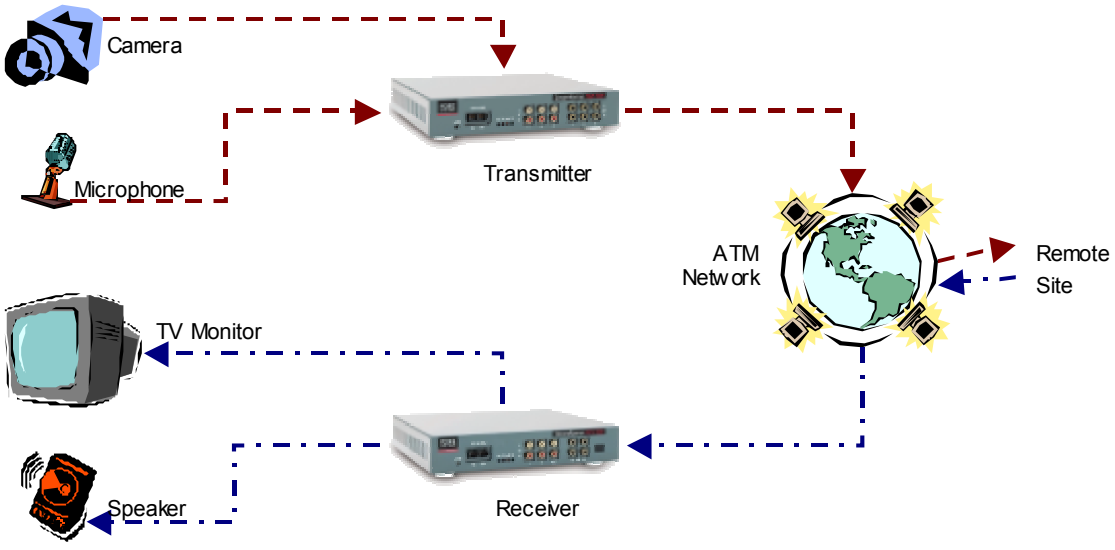

Figure 2. The video wall setup at each site. The arrow lines represent the directions of data flow.

other activities. We set the audio quality to stereo CD quality and the video quality to 30-frames-per-second NTSC. To help adjusting the camera at each site, we put loop-back video from the local camera in a small window on the local monitor, that same as the "picture-in-picture" mode seen on some commercial TV sets.

- **Observations**

Initially we tested the configuration by carrying on conversations across the video wall to see if the teleconferencing system could support remote meetings. We found that the loop-back "picture-in-picture" window obscured the view from the remote site. Ideally, the loop-back signal should be displayed on a separate, low cost, small monitor. Disabling the loop-back video made people standing outside of the camera view unaware that they were invisible to the remote site, however. This caused confusion and unease when people off-screen started to talk. Even when the loop-back video was available, auto-focus cameras with wide-angle lenses would have covered large spaces better.

Due to equipment shortage, we did not have good quality microphones for the video wall. People had to hold the microphones close to their mouths in order to allow people at the remote site hear them clearly. However, we did observe that people used the connection to chat, to ask other people to go to lunch together, or to find people at a remote location. A graduate student even Greek-danced with one of her friends across the audio/video link. Sometimes people would sit near the video wall and work on their own, and wave sparingly to other people at the other end of the virtual connection.

The video wall setup was generally welcomed in the CRB. People who worked there mostly felt more connected with the College and less "abandoned". However, at the CCB end, several people strongly objected to it. The microphone and camera in the CCB were

often turned off. Many e-mail messages were exchanged debating the video wall and people suggested that the following were possible reasons to the objection:

The audio/video link could break the assumption of privacy and indeed did make some people feel uncomfortable. The common area in CCB housing the video wall had a long history of prior use and people had formed the habit of having lunch there. Since it was an semi-enclosed room at the corner of the building, far away from classrooms and labs, not many people would pass by. Therefore, people did not expect other people overhearing their conversations. On the other hand, the CRB common area was new to those who moved to that building and the video wall was deployed in an open space in the middle of offices and open labs. People did not have high expectations about privacy in the CRB common area. Therefore, the video wall connection did not pose a threat to their casual conversations held in that area.

Another concern to some people was the extent of processing and storage of the audio and video signals. In reality, each ATM packet carrying audio or video signals lasted only a small fraction of a second and was not routed outside of the College's computer network. All reconstructed data were only played back and not recorded anywhere. However, nothing clearly communicated these facts. When the local TV monitor was turned off, a person might not know that the camera could be still on and he or she could still be watched remotely.

Due to the light traffic and the sentiments of many people against the CCB video wall, the virtual link between CCB and CRB was infrequently used for conversations after its novelty and people's curiosity wore off. After the equipment in CCB was moved to a newly renovated space without prior history of use and expectations, the video wall was

received better there. However, like other public-space sharing systems [58], its infrequent use did not seem to fully justify the cost of the equipment and maintenance.

- **Implications**

Retrofitting an existing space to support remote presence demands careful evaluation of its current place in the community culture to avoid conflicts with existing habits or expectations. The states of input devices and data processing systems should be clearly indicated. The spaces in which people are subject to being watched or heard by remote sites need to be clearly defined and understood. People should not have to pass the spaces when they do not want to participate in remote presence.

It is possible that the quality of current audio and video transmission technologies is simply not good enough to promote remote informal interaction. For example, the clarity of the audio or the resolution of the video may be unsatisfactory. In addition, traditional audio and video transmission mechanisms often suppress or distort subtle yet important interaction cues that people take for granted in face-to-face interactions, such as gaze directions, facial expressions, and speech tones [34]. In absence of these cues, people may feel uncomfortable and difficult in engaging in remote interactions without realizing what the problem is [100].

It is also likely that people do not feel any compelling reason to adopt this style of interaction because it does not help to inspire the topic or content of the interaction. For example, in a face-to-face meeting, people can easily share the sight of a model plane in the physical environment, and they may start a casual conversation about flight-control-interface design issues while examining the model. The model itself may not have the representations of the controls, but the human associative memory can stimulate related

thoughts, such as the controls and the design of them. In addition to supporting speech and vision in awareness systems, providing some level of topic information or content may help inspire people to start spontaneous interactions.

Furthermore, the use of the video wall is inevitably limited to those who pass by that space. There are a large number of people in the community who do not come near the video wall on a regular basis. For the overall awareness in the community to improve, we need to build technologies to reach those people as well.

2.2.2 The Electric Lounge

Desktop media spaces [12] are one alternative to video walls in trying to support informal awareness and interaction among community members. Instead of setting up expensive equipment at centralized locations, media spaces allow the community to transfer the cost to individual computer systems and can potentially connect more people. In addition, users have more control over media space applications running on their desktops and devices connected to their desktop machines.

We set up a simple media space called the "Electric Lounge" using public domain teleconferencing applications including the audio tool vat [57] and the video tool vic [73, 74]. Early systems such as the Portholes displayed a collection of periodically updated video images from remote sites, allowing a user to see quickly who was available in a group of people [28, 67]. An office sharing experiment used continuously running video links to connect pairs of offices together [26]. The Electric Lounge project aimed to explore such approaches in enhancing awareness of other people in a large community.

The advantage of using vat and vic to build a media space was that the source code of the programs was freely available, providing us the opportunity to modify the tools for our

needs. The disadvantage was that they mostly ran on Unix machines. They had limited support for Microsoft Windows systems and no support for Apple MacOS computers.

- **Maintaining Awareness**

The original audio-tool, vat, was used without modification. It listed the names of people who were connected to the Electric Lounge at the time (see Figure 3). By default, vat played back everything that was sent onto the network, but the user could mute all audio streams or specific ones from other users. The program did not broadcast audio data on start-up. The user needed to press the "talk" button before saying anything in order for other people to hear it.

Figure 3. The audio-tool, vat. This example shows that five people are listening in the Electric Lounge. The user can mute audio from an individual participant by checking the box before the person's name. The "listen" and "talk" check-buttons specify whether the program should play network audio and whether it should broadcast local audio signal to the network. The speaker and microphone icons indicate the devices that the program is using for audio output and audio input. The sliders adjust the volume or input gain of the audio devices.

We modified the video tool, vic, in a few areas. This program showed video streams from all video sources connected to the Electric Lounge in thumbnail pictures. The original vic program displayed several kinds of statistical information along with the names of the video streams next to the thumbnail images (see Figure 4). Most of these statistics were not related to our goal of supporting informal group awareness and were moved, leaving only the names (see Figure 5). An additional advantage of this modification was that we accommodated more thumbnail videos using the same or smaller screen space. To

Figure 4. The original vic program. It shows the thumbnail images of all video sources as well as a few statistics associated with each video stream. Clicking a thumbnail pops up a full-size view of the video. Checking the "mute" button stops a video stream. Unchecking the "color" button lets the program to decode that particular video in grey scale.

further ease navigation and locating specific video streams, we added a scroll bar in the main application window.

Users could click on a thumbnail picture to bring up a full-size window of the corresponding video source (usually a quarter of NTSC size, see Figure 6). It also displayed a one-line note posted by the sender of the video, and statistical information that used to be in the main application window.

The "Control Menu" item in the "Options" menu brought up the control panel (see Figure 7). The user could change the network bandwidth usage and frame rate limits,

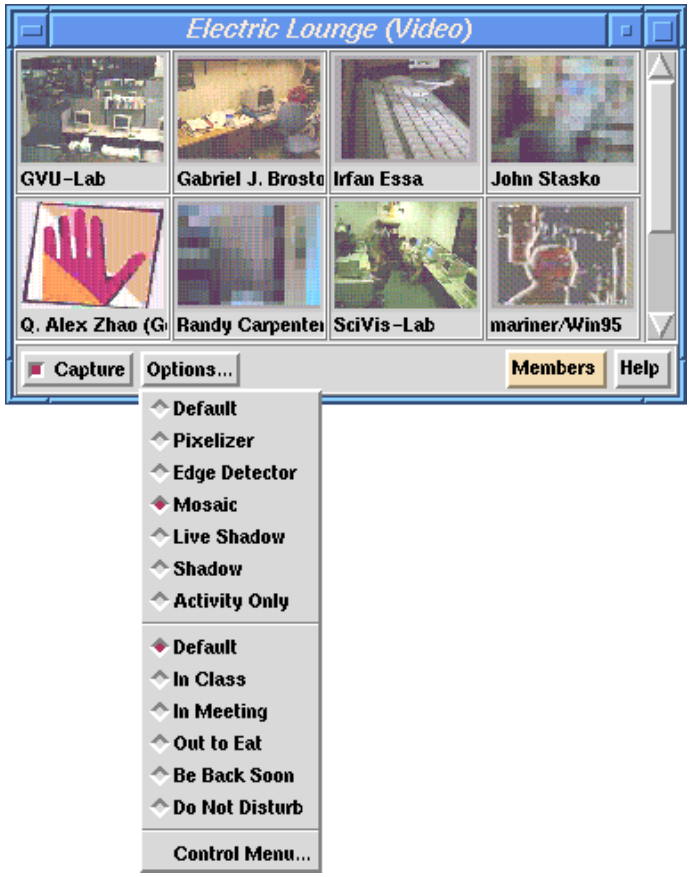

Figure 5. The modified vic program with the "Options" menu popped up. It shows the thumbnail images of all video sources labeled with the titles of each video stream. Clicking a thumbnail pops up a full-size view of the video. Toggling the "Capture" button stops or resumes the video capturing and transmission on the local machine. The "Members" button brings up a list of participants in the current video session.

which were by default set to ten kilo-bits per second and one frame per second, respectively. The user could also adjust other parameters such as the video capture hardware device, video format, and image size.

In an attempt to help locate people and provide casual awareness of work areas, we placed cameras in two labs: the GVU lab, and the SciVis lab. The video images were broadcast to the Electric Lounge and periodically saved to a Web page. There was no audio transmission associated with the lab-cams.

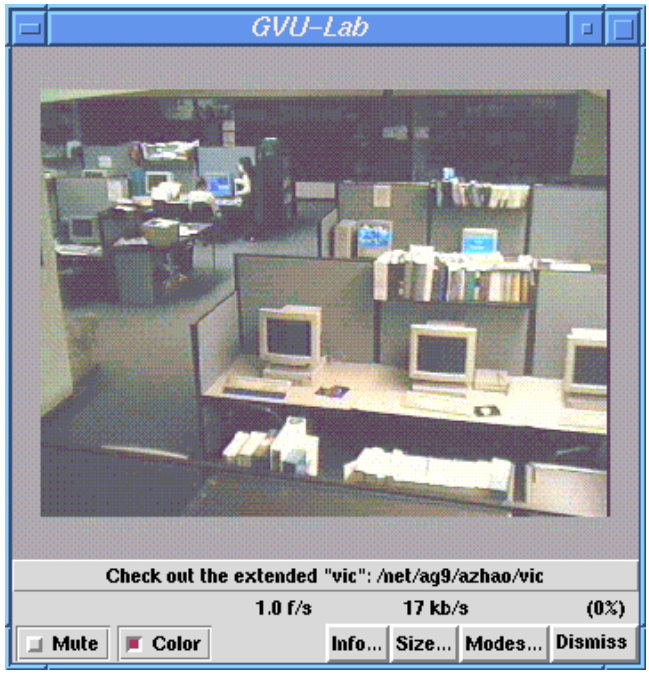

Figure 6. The full-size video display. This view shows a one-line note posted by the video creator below the video image. It also shows a few statistics about the video stream, such as frame rate, bit rate, and packet-loss rate. The “Info” button pops up a list of details about the video, such as source IP address and port number, the network protocol being used, and specific statistics about the particular network protocol and video decoder. The “Size” button allows the user to adjust the image size that the program uses to decode this video stream. The “Modes” button lets the user to control several parameters of the decoder, such as whether to use hardware decoding, if it is available.

- **Managing Privacy**

We made a number of changes to help the user control how much information about themselves was available to other participants in the Electronic Lounge. The first set of changes dealt with giving the user adequate feedback about as well as control of the state of the application. In the original vic, if one wanted to see whether video from the local machine was transmitted to the network or not, or wanted to start or stop the transmission, the user would have to bring up the control panel to accomplish this (see Figure 7). We added a "Capture" check-button to the main interface to show the state of the transmission and allow the user to easily toggle between transmitting and being silent. The "Members" button could bring up a list of all connected parties in the video session of the Electric Lounge, including the lurkers (those not transmitting videos). The "Members" button would also change color if the program detected lurkers. By clicking on the names in the members list, the user could choose to either receive or ignore future transmissions from the corresponding people. The "Options" menu helped with setting the clarity level of the video being broadcast from the local machine. In addition, the user could choose to post static signs. For example, people could post the "out to lunch" sign, telling potential visitors what they were doing and roughly when they were expected to be back.

To help people control the clarity of their virtual presence in the video streams, we added six image filters: the pixelizer, the edge detector, the mosaic filter, the live-shadow filter, the shadow-view filter, and the activity-only filter (see [51, 128, 129] for more detailed descriptions of these). The first five filters made video images less clear. The activity-only filter replaced the real video stream with a display of activity history: each line in an activity-only image represented a point in time, with the line at the bottom of the

image representing the most recently observed time, and the line at the top representing the time minutes earlier. The intensity of each line represented the observed video image intensity change level at the corresponding time: the brighter the line was, the more different the image captured at that time was from previous frames.

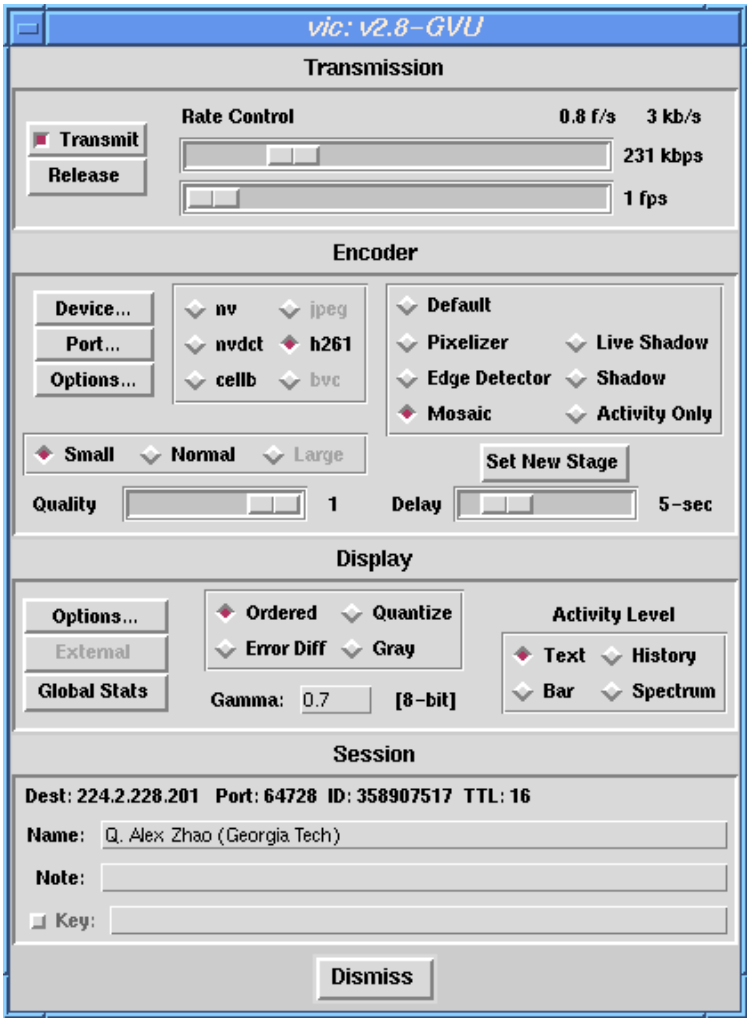

Figure 7. The vic local video control panel. The sliders at the top of the dialog allow the user to adjust the data and frame rate of the transmission of local video signal. The encoder section controls the parameters for encoding local video input, such as video device to use, the size, format, and quality of the encoded video image, as well as the privacy filter to use. The display section sets how data received from the network should be decoded and displayed. The session section allows the user to change the name that he or she wants to be represented, and post a quick note about the video.

The live-shadow and shadow-view filters required a static background scene to function properly. When the background scene was changed, for example when a piece of furniture was rearranged, the user should use the control panel to let the program take a snapshot of the scene again. The user would have enough time to leave the field of view of the camera before the new snapshot was taken.

A personal security issue was raised about the lab-cams: someone who worked in the lab during the night could be tracked by someone else and ambushed. In response to this potential risk, we modified the vic program that controlled the lab-cams to automatically pixelate the video images from 7pm to 7am and during weekends.

- **Observations**

Despite our efforts in promoting the Electric Lounge through e-mail announcements, newsgroup postings, and demonstrations, we had only about four to ten volunteers using the Electric Lounge at any one time. A number of factors contributed to the difficulty in the deployment. For example, a computer with audio and video capture capabilities was still a luxury in the College when we set up the Electric Lounge. Having the media space tools automatically start up after a user logged in involved tweaking Unix script files. In addition, the unavailability of corresponding MacOS tools and the initial instability of the video program turned away some users.

We asked people how they used the Electric Lounge in informal meetings and e-mail messages. We also solicited seven first-year Ph.D. students to experiment with the media space and report their opinions. We found that people did not use the audio-tool much, partly because broadcasting sound to everyone could be distracting and rude. When a user accidentally “saw” another person in the Electric Lounge, she most likely wanted to chat

with that particular person rather than everybody connected to the media space. The way to start a private conversation with someone in the audio-tool (clicking on the other person's name using the middle mouse button) was not visible in the interface. Furthermore, audio and video were handled by two separate applications and it was difficult to switch back and forth between them. Consequently, people preferred to use telephones when they needed to talk to someone.

Users appeared to become relatively accustomed to having the media space applications running on their desktops. They tend to ignore the videos when they were working on other tasks. However, people complained that those applications still took up a considerable portion of their screen real estate. In addition, the video tool competed for colormap entries with other applications. On an indexed-color display, colormap swapping occurred, causing screen flashing and distraction to the user. Moreover, the encoding, decoding, and displaying of video data required more computing resources than other typical desktop applications such as text editors or e-mail readers. Although the video tool executed at the lowest system priority, it inevitably affected the overall responsiveness of the computer systems.

Based on user feedback we found that, despite the occasional inconveniences, people used the videos to find others and check if the other person could be interrupted. As familiarity increased, people learned to better estimate the availability of others, even when the videos were blurred or pixelated. However, users of the video tool often settled on particular filtering modes and seldom changed to other filters. The pixelizer was the most popular filter, partly because it was the default filter and many people did not bother to set a different default. From informal interviews with the users, we found that this filter left enough

information that made the resulting video recognizable, while hiding much detail of the actions (see [13] for a follow-up study on the filters). People felt that the images filtered by the edge detector were not aesthetically pleasing. The shadow filters involved too much effort in setting up the backdrop images and making sure that the camera would not move. The activity-only mode gave too little information and required mental processing to understand that information.

Also from the interviews we found that after several months of continuous use, people generally valued the Electric Lounge as a tool for emotionally connecting to the community, especially for those in isolated offices or labs away from the main body of the College. From time to time, they would glance at the images to get a vague sense of status in the virtual proximity.

However, the Electric Lounge users noted that many people they would like to keep in contact with were not using the media space. The dilemma of the Electric Lounge was twofold. First, a critical mass of users was never achieved for it to be more useful. Second, if there was a larger user population, we would have had faced the information overload problem associated with increased number of video sources to be displayed.

People harbored differing views about the lab-cams. Some people cautioned that the introduction of the cameras violated the privacy of those who worked in the labs. Even worse, because there was no feedback given to people in the labs about the existence of the cameras, they most likely did not know that their privacy was being compromised. On the other hand, several people in satellite offices looked at the lab-cams videos from time to time because that was where some community activities took place and the lab-cams let them know what was happening in the labs — familiar places that they liked. They argued

that the privacy expectations in public lab environments were lower than those in private offices. In addition, since the lab-cams were placed high near the ceiling and far away from people working in the labs, it seemed to be difficult to identify individuals from the slow video images without high levels of familiarity. Therefore, the severity of possible privacy violations was likely to be low. However, in principle, the two groups all seemed to agree that the lab-cams did not bring direct benefits to those who were in the labs, possibly being watched.

- **Implications**

Informal awareness and interaction services ideally should be available to everyone in the community, either on the desktop or off the desktop, without dependency on computing platform, hardware resource, or network bandwidth. The information displays need to be as minimally distracting as possible and be economical with screen space usage. If accessing the information involves running an application program, that program needs to be easy to setup and use without external help. The program needs to be reliable, simple, but flexible, and customizable. It needs to convey understandable information even when the user population is large. When using a media space to support informal group awareness, it is best if the user has central control of all media parameters through a single interface.

An effective desktop portal into the virtually connected community has to promote frequent "sightings" in the virtual environment that result in conversations. In a community such as the College of Computing, meeting face-to-face is still possible, even though people are spread out in different buildings, making such meetings inconvenient. Therefore, the purpose of only *seeing* others may not justify the personal costs involved in setting up

the portal. For example, being able to see the copy room and the line waiting for the copy machine to become available, as one faculty member suggested, may bring other types of incentives to users and provide a virtual meeting place that community members are more likely to use. In addition, according to personal communication with several users, being able to hold conversations in-place, at least textually, will make the portal more useful and attractive.

Privacy is always going to be an issue that needs special attention, however. The concept has many different interpretations, and consequently, there are different methods to approach this issue [4, 11]. For example, there is the technical aspect of data protection that leads to the pursuit of better encryption algorithms. For the purposes of community awareness and this thesis, we will focus on the psychological aspect of the boundary of personal information. For example, people using the copy room may object to being watched by others whom they are not aware of.

2.2.3 Discussion

In conclusion, the video wall and Electric Lounge projects confirmed that the limitations on availability, accessibility, interaction styles, and shared context could prevent awareness and interaction systems from gaining wide use. Supporting the virtual presence of people and subsequently direct interpersonal interactions may be appropriate for improving awareness and communication in small groups, but it may not be appropriate in a large community. Relationships among a small group of people are often relatively stronger and people are more willing to involve relatively more effort in maintaining awareness and interacting with others. In a large community, however, such relationships are often weaker and it may be impractical or unnecessary to keep up with the presence of

many other community members. We need to find lightweight techniques to help communicate the state of the community and help people maintain those weak-ties with the community.

CHAPTER 3
Motivation

This chapter provides a theoretical motivation for opportunistic interfaces. We will start by examining different aspects of maintaining awareness with respect to communication and collaboration. We will briefly survey existing technologies for supporting distinct awareness that directly relates to interpersonal interaction, and analyze the characteristics of the general awareness pertinent to amorphous communities composed of mostly weak-ties. Finally, we will identify opportunistic interfaces as an alternative approach in maintaining one's awareness of the surroundings and other people, and devise a set of specific goals in designing community awareness applications.

3.1 Distinct Awareness and Related Technologies

The Merriam-Webster dictionary defines *awareness* as the state of "having or showing realization, perception, or knowledge" [77]. From the human-computer interaction perspective, awareness is "an adaptive, externally directed consciousness" [127]. Although maintaining awareness is not the eventual goal of any collaborative activities, it enables people to collaborate more effectively and smoothly [62].

A large body of research has focused on supporting the distinct awareness that tightly couples with interpersonal interactions, such as the physical appearances of people involved and the specific actions performed. We can loosely characterize these existing

awareness technologies using three primary domains: the function domain of people, actions, and artifacts; the time domain of past, present, and future; and the process domain of production, aggregation, and consumption.

3.1.1 Function Domain

Awareness and collaboration applications often have three primary functional areas to maintain awareness of to help people construct a basic understanding of the collaboration: the identities and states of the participants, the activities that the participants perform, and the artifacts that the participants manipulate. These functional areas are closely related to each other and often manifest together in individual systems.

A wide range of techniques have been developed to support virtual presences of people. MUD and MOO based systems such as the MediaMOO project [15] employ descriptive text to allow users imagine the appearances of other characters. In traditional instant messaging systems such as the MSN Messenger [79] and the AOL IM [5], the font styles and different indicators associated with a text label can tell whether the remote "buddy" is online or busy. Novel instant messaging systems such as Hubbub [55] may use audio cues to help identify people when their status change or when they initiate instant messages. Geometric shapes, such as the colored dots in Babble [29, 123], the concentric circles in Chat Circles [116], and the cartoonish icons in Peepholes [40], are sometimes used to convey presence. Graphical avatars in virtual worlds such as V-Chat [104] portray presences as well as emotions and personalities through virtual appearances and gestures. Finally, many systems shows images of remote users that ranges from the simple pictures in Piazza [56] and occasionally updated video snapshots in Portholes [28], to live video streams in our video wall project and mobile video units that support remote exploration [36].

In addition to simple status information about a remote user such as the online or busy states, where that person is located relative to the local user is also an important aspect of virtual presence. First, location information facilitates finding people. The Active Badge Location System uses the transmitters implanted in building access badges and a network of sensors in the building to pinpoint the locations of badge wearers, helping people reach each other easier [117]. Second, location information provides a social context and management mechanism for interpersonal interactions. For example, people usually feel more comfortable holding a conversation near each other in either the physical world or the virtual world than far away from each other. In many awareness applications such as Chat Circles, one shows interest in a conversation by moving closer to it and leaves a conversation by moving away from it.

Awareness of other people's activities in the context of collaboration helps a person assess overall progress and plan future actions in response. For example, the Conversation Board supports manipulation of graphical objects known as *conversational props* on a shared white-board, facilitating the exchange of ideas through free-form expressions [14]. Workspace awareness applications such as TeamRooms embed different collaborative activities on shared workspaces [96] and may use various focus + context techniques to help the user see other people's views, gestures, and actions [46]. Augmented reality applications such as Videoplace [64] and ClearBoard [53] may superimpose hand or face video streams from meeting participants with imageries of shared work surfaces and objects, giving the illusion of drawing on the same workspace or a transparent glass separating the users. Finally, virtual rooms in the DIVA virtual office environment show what people are working on and provide collaboration spaces for different tasks [106]. The

room metaphor helps arrange related work together, manage different styles of work, and reduce interference from unrelated topics [41].

People's actions transform objects, documents, and other artifacts. It is often useful to know the states of these artifacts, the changes that have been made, as well as those that are being made. For example, ShrEdit presents the most current version of a shared document while people are making changes to it [27]. The Fisheye Text Editor uses fisheye lenses to show the current state of the document and provide context for editing activities [39]. The virtual desks in DIVA support virtual grouping of shared documents, often represented as icons in the shared workspaces. The virtual briefcases in DIVA support convenient transportation and storage of the artifacts. In the TeleNotes system, shared documents can be passed along with annotations and dropped on personal workspaces not shared with other people [121]. Sometimes these documents can help initiate synchronous discussions through embedded "anchors" [19].

3.1.2 Time Domain

Many of the technologies that we have described so far support same-time awareness and allow people to collaborate synchronously. For example, DIVA shows what document people are working on at the moment. The Portholes video images show current activities in remote offices, such as talking on the phone or with a visitor.

Awareness of past and projected changes allow people more accurately analyze the past, assess the current state of the collaboration, and prepare for future activities. For example, Ginsberg and others described a virtual meeting room system that automatically record audio, video, mouse, and keyboard interactions, capture and index meta-events such as entering the meeting or switching from speech to drawing, and generate visualiza-

tions of the meeting history [38]. Many applications such as e-mail and Usenet news readers often mark folders or newsgroups differently to distinguish the ones that have been changed, i.e. have new messages added since the last review. TeamSCOPE tracks and organizes the progress of group projects on a web-based calendar [109]. A series of video snapshots from a remote office can be archived and played back later to discover patterns of activities in the office. Finally, visualizations such as Visual Who [25] and PeopleGarden [125] portray patterns of interaction and history of participation through the arrangements of color, orientation, size, and relative location of geometric shapes that represent people and tasks.

Awareness of pending tasks and actions helps people coordinate their activities effectively. Threaded Chats facilitates turn-taking in chat-room discussions by allowing people to see what one person is going to say while the message is being entered [103]. Software in-out boards and group calendars such as @Work [113] and TeamPortal [47] can tell people's whereabouts and help relay information. Using Bayesian models, Ambush predicts event attendance to facilitate intended casual interactions [83].

3.1.3 Process Domain

Awareness information often flows through a communication system via several stages: it is first captured on certain devices, then converted and aggregated to certain forms appropriate for the application, and finally transported and presented to the information consumer. In addition, feedback generated on this path in turn help choose what to capture or control how information is conveyed.

There are several different ways to initiate this process, and correspondingly different awareness technologies to support them. Kraut and others distinguished the initiation of

interpersonal interactions in varying degrees of spontaneity and formality into four categories [21]:

- *Scheduled*: a previously arranged interaction;
- *Intended*: in which the initiator looked specifically for the other parties;
- *Opportunistic*: in which the initiator intended to talk to the other persons and took the opportunity of a by-chance encounter to start the interaction;
- *Spontaneous*: in which none of the participants had any plans for the interaction and yet it started purely by-chance.

Scheduled and intended interpersonal interactions are often supported by proactive awareness technologies that provide tracking or glancing mechanisms to support formal and semi-formal interactions. For example, we often need to explicitly join a teleconferencing session to attend the online meeting. Montage [110] and RAVE [35] users can open a short video connection into another person's office, glancing into the remote space to see if there is any possibility for interaction.

Opportunistic and spontaneous interpersonal interactions account for the majority of the overall interpersonal interaction occurrences in the work environment [63]. Reactive awareness technologies create, detect, and present opportunities for interaction to the user with minimal involvement from the user [120]. For example, Cruiser [31] and RAVE let a user to initiate short audio/video connections to a series of remote offices, mimicking a walk down the hallway in an office environment, looking for whoever is available for interaction. Our video wall and Electronic Lounge projects let the user directly see who is around and what they are generally doing. Opportunity for interaction arises when people sense that other parties are in the virtual vicinity and become available for interruption.

Awareness system can let the user explicitly control or algorithmically determine the production of awareness information. In the DIVA virtual office, private rooms can be created so that only a particular set of people can see or enter these rooms. A person using an instant messenger may need to obtain permission from a remote user before putting that user on the list of people to monitor. Furthermore, instant messengers often allows users to directly set their online status such as "busy" or "away from desk". On the other hand, after detecting that two people are reading the same e-mail message or reading the same web page, Active Mail [45] and Piazza [56] display pictures or current video snapshots of the other party, and Contact Space [59] walks people's avatars toward each other.

Interestingly, social norms can also be utilized to control awareness information production. For example, a half-open office door in the physical world may signal that the occupant is in the office and available for important discussions, while a closed or mostly closed door may imply that the occupant is not in the office or does not want to be interrupted. This particular social norm can be used to control a person's video presence where video is transmitted only when the office door is open [16]. It can also be directly translated to the online world where the open and closed states of a door icon help convey different access permissions associated with the virtual room.

Awareness information originates in many different forms, such as the sight and sound in the remote space, the opening angle of a door, and the location of people and objects in the physical as well as online worlds. To help ease comprehension, reduce overload, and manage privacy, awareness information may need to be aggregated or abstracted to a smaller number of simpler forms [90]. For example, social activity indicators may transform and combine various data from several sensors to a simple line drawing [3]. Using

computer vision techniques, office activities such as “talking on the phone” can be extracted from video streams and described in textual form [18].

Depending on application requirements and restrictions, awareness information is presented using different approaches. Portholes shows video images from all awareness participants and is often displayed on secondary monitors. When screen real-estate is not available, only partial information can be shown and applications may use cyclic animations to provide adequate coverage of the whole information [75].

In summary, we have described three primary domains of supporting awareness in collaborative environments. We have touched on a number of awareness technologies that facilitates interpersonal interaction. The scope of awareness, however, has been mostly limited to the people that one already interact with or already knows.

3.2 General Awareness and Opportunistic Interfaces

In contrast to distinct awareness in tightly coupled groups that may immediately affect people’s activities, community awareness is a general form of awareness in an amorphous group of people, including those who we do not already know. This general awareness is relatively less detailed, is loosely related to people’s current activities, and often contain tacit knowledge about the environment. For example, we can learn a lot about the community and its people without actively acquiring information about them, perhaps even in an unrelated context. When we hear about a local research project on a news radio, we may unexpectedly gain some knowledge of what fellow researchers are working on. When we read an e-mail about a town-hall meeting, we become to know what the community is concerned about and what possible changes may come in the future.

Although this information is often processed without intentionally focusing our attention on it and without elaborate investigation of the meanings or implications, it keeps us aware of the states and shifts in the community as well as the states and changes of the people in-between interpersonal interactions. Not only does this background awareness benefit us practically in future interactions, preparing us better for casual encounters, but also it helps us connect better with the environment and other people, helps us build and maintain social ties with others, as well as communicates social norms and expectations.

E-mail is an effective method in delivering information and making contacts. It is used in formal, procedural settings, as well as informal, ephemeral ones [107]. However, in recent years, people are becoming overwhelmed by e-mail due to its popularity and the amplifying effect of misuse. The result is that often people process or filter e-mail too quickly to notice valuable information about their surroundings. A new idea mixed with routine status report messages, an introduction to a new group member, or a reminder of a community picnic may arrive at a time that the receiver of the messages decides to delay reading them. Sometimes those pending messages may never be read, or when they are read, it may be too late to make use of the information.

In the physical world, people post flyers on bulletin boards and telephone poles, set up billboards along highways, even use advertisements on television and radio broadcasts to convey information opportunistically. Information is initially "pushed" to potential audiences rather than "pulled" by them. Viewers or listeners are not obligated to pay attention to the pushed information. Information is delivered by chance and people can choose to pursue more details if they are interested.

Similarly, we can use computer-based technologies to achieve the same effect. Opportunistic interfaces, as the name implies, present information opportunistically. They seem to be a natural fit for supporting the general, background awareness that manifests in undemanding information. In particular, the opportunistic nature has two implications. First, information is presented by chance. Although one particular piece of information may or may not be interesting to the user, an effective opportunistic interface should maintain a high overall interest level in the user on a set of multiple instances of information, over an extended period of time. Second, the user notices or consumes the information by chance. The frequency of this happening will be related to people's subjective opinion about the particular communication mechanism. High frequency means that the user constantly pays attention to the information channel and may feel distracted. Low frequency implies that the user rarely uses the information channel and may consider it a waste of resources.

There are several places to build opportunistic interfaces, each have its own set of implications and limitations. On the computer desktop, media spaces, instant messengers, and chat media such as Zephyr [2] can be used to opportunistically monitor awareness. Systems such as Irwin [75] and the Awareness Monitor [17] aggregate dynamic information from multiple sources and may use small graphical indicators to signal the relevancy, magnitude, or urgency of the changes. However, these technologies compete screen real-estate against other applications and demand a considerable amount of effort in order to decode the information or find particular pieces of information. Spending a non-negligible amount of work on potentially peripheral knowledge may become a distraction rather than a help. Knowing that oneself can be potentially tracked by these tools, especially in media spaces where much more information about a person is available, may cause unease. Inter-

ruptions from other people using these tools may hinder the performance on primary tasks [22].

The once-popular PointCast [92] and other commercial "Internet Toolbars" show information in the edges or corners of computer desktops or in the screen-saver mode. The assumption is that they can provide the user with some information opportunistically while allowing the user to complete other tasks. However, the developers' business needs of selling advertisements and products are in conflict with the users' need of monitoring secondary information in the periphery of attention. The result is that these tools often use flashy animation to attract the user's attention away from their primary tasks. Furthermore, when co-existing with other applications on the user's desktop, these systems as well as Tickertape [32] often rely on continuously scrolling animation in order to fit long textual messages in one or two lines, also causing distraction to the user [72, 75].

Systems such as the Apple On-line Newspaper [66], the Aware Community Portals [98], and the Notification Collage [42] use large public displays to bring on-line content into the physical space where people work or pass by. The information being displayed keeps inhabitants of the shared spaces updated on local community news and events, technology news, and stories contributed by community members. It also gives people something to talk about when they are in the same place. These systems are best suited for organizations that have natural shared spaces, and they often require a certain amount of effort from the users to participate or use.

Calm technologies explore interfaces that stay in the user's periphery of attention while still providing some value to the user and shift smoothly into the focus when necessary [119]. The Digital Family Portraits use the picture frames to show rough estimates of

recent levels of activity in remote spaces [82]. The water lamps and pinwheels [23] and ambient displays in ambientROOM [54] convey information in the background, using unobtrusive physical objects, reflections, and sound effects. The Information Percolator uses computer-controlled water bubble generation in an array of water chambers to form text and imageries and communicate information through a decorative object [48]. Audio Aura provides auditory cues about tasks and colleagues while the user moves around in the physical work space [81]. These technologies do not distract users from other tasks or use valuable screen real-estate. However, they often support low information bandwidth awareness and require a certain amount of effort to decode and comprehend the information.

In summary, even though the general awareness of the on-going changes in the environment and attributes of people in between interpersonal interactions seems to contribute to the missing link for achieving the state of *readiness* [85] for spontaneous collaborations [8], it has not been properly supported in existing technologies. Current systems have limitations in three aspects of system design: low information utility, high involved effort, and obtrusiveness. Our goal in this research is to study opportunistic interfaces and develop a set of guidelines for the design and implementation of these interfaces. Specifically, we hypothesize that opportunistic interfaces should provide useful content, involve minimal amount of effort, and be non-distracting to the user. We believe that opportunistic interfaces can enhance awareness and improve communication in a distributed academic research community.

3.3 Design Goals and Trade-offs

The ultimate goal of our work is to enhance general awareness about the community and its members. In corollary, we need enough users to adequately study the effects of community awareness applications.

However, these goals are likely too high-level and abstract to guide the development of an opportunistic interface. Therefore, we have devised a set of more concrete objectives to drive the creation of community awareness tools throughout the research. These objectives are described below:

- **Provide useful and interesting content.**

In order to attract and sustain users, our opportunistic interfaces need to have content that people are interested in or find useful.

Content can come from user contributions or it can be collected automatically by computer programs. In the early stages of system deployment, people may not want or like to submit content because they may not know what is appropriate or what is the norm. Therefore, we cannot expect much user-contributed content in the beginning and we need some automatically collected content to help build a user base. However, people have wide variety of interests and it is difficult to identify topics that will appeal to all of our potential users. In addition, people's interests change and the changes may depend on factors that we may not be able to detect or control. Therefore, the practical solution is to assess what people's interests are and support the most popular.

In addition to the implied types of content that tie into the community life as described previously, we may need to include external information such as news or weather forecasts to increase the utility of the tools. This can also help us control the quantity of information

available through our systems. With too little information, awareness systems may not provide users the level of benefits to justify the cost of use and may become uninteresting as user's experience grows. On the other hand, with too much information, awareness systems can overwhelm the user and become distractions.

Along with automatically collected information, we need to encourage people to contribute their own content and allow them to informally express their interests and opinions. Potentially, the automatically generated content can lead to topics that people can talk about, while user-contributed stories and discussions tell much more about the individual users as real-life people. We hope that this participation and new content will help maintain existing user base and attract more users.

- **Encourage open, public, cross-group communication.**

Although people in the College of Computing community are connected organizationally, we are divided by walls that we can see, as well as walls that we cannot see. The walls that we can see are the physical office walls and partitions, the distances between locations where people work, and diversities in people's schedules and work practices that prevent us from chance encounters and impromptu interactions. The walls that we cannot see are the group boundaries, prejudices toward other groups, and general ignorance about other people that prevent us from effectively interacting and collaborating with each other even when such chance is present.

Rather than trying to lower the physical walls that we can see, we believe that by fostering an environment where people can openly discuss any topic they wish to discuss, with anybody in the community, we can lower those psychological walls that we cannot

see. This cross-group interaction can help expand one's social ties and connect with the environment [91].

To encourage cross-group communications is also the main reason why we have decided not to support any sub-group feature in our systems. Supposedly, a user can join a set of sub-groups as a filtering mechanism to control what groups of people will receive the user's postings, and in what groups the user will receive other people's postings. However, the sub-group idea can add much complexity to the system design in the aspects of initial group setup and policies on group creation and membership maintenance. In addition, selecting group preferences can be a burden to the user. And more importantly, sub-groups implicitly encourage division among people and discourage cross-group communication. Since denoting a message to a set of groups requires more processing on both the sender's and the receivers end, within-group communication may be better supported by more formal applications such as e-mail systems and newsgroups.

- **Complement existing communication methods.**

People use many different mechanisms and tools to receive, collect, and organize information. In addition, many people are very proficient and reliant on what they use. Similarly, a community has its established ways for communicating information about itself. Rather than trying to replace some of the existing communication methods which can cause disruptions and unease to community life, we need to focus on technologies that complement what already exists.

A concrete example may be to collect pointers to information sources and provide brief summaries about changes, rather than to build a self-contained information warehouse from the ground up. As long as we are flexible as where the information comes

from and where to point our users to, without restricting them to our ways of exchanging information, the users will find the best ways to integrate the opportunistic interfaces with what they already use.

The diversity in a large community predetermines that designing community awareness applications is a challenging task. In our example, the College of Computing is a particularly diverse environment. People have quite different backgrounds and life experiences. There are age and gender differences, family status differences, nationality, language, and culture differences, education differences, and many more. In addition, College of Computing people also have different goals and expectations on what they want to get out of this environment. Some care much about their peers and the welfare of the community. Some only want to do their jobs or just want to finish their degree and leave. And perhaps the majority of people fall between the these two extremes. Therefore, we are unlikely to find a single solution that appeals to everybody in the College.

The dynamics and complexity of people's work environment also predetermine that it is impractical to find a single perfect solution for enhancing awareness. Generally, a person may choose to process a piece of information, put it aside for later processing, or ignore it. Regardless of how one receives the information, how the person subsequently act on it depends on many different factors. From the information receiver's point-of-view, the contributing factors include their busyness, other tasks that they are handling, the perceived importance and urgency of the message, the receiver's interest level, their habits of handling such information, and even their mood at the moment. From the information delivery side, how the receiver acts on the information depends on the method and quality

of the presentation. Developing a set of complementary tools and choosing the right one for the situation may be more cost-effective than trying to find the single perfect tool.

- **Have a simple and calm interface.**

Opportunistic interfaces should not distract the user from their focus or primary tasks. Keeping up with the current affairs in the community and with perhaps a few of external information sources is unlikely to be a primary activity that a user frequently engages in. Therefore, awareness applications need to stay in the user's periphery of attention when the user is not directly interacting with them.

At times when the user is paying attention to the awareness application, the interface must allow the user to quickly finish what they want to do so they can return to their primary tasks. The awareness application is there to help enrich a user's work environment, not to replace their work activities. Because of their peripheral nature and the infrequency of direct interaction, opportunistic interfaces should be simple and intuitive so that people do not need to learn or remember much in order to use them. In addition, the applications should not need much configuration in order to fit in with people's work styles.

A number of promising ideas emerged from our brainstorming activities about potential interface designs. One possibility is a large flat panel display strategically placed somewhere in the College, showing various kinds of information in pleasant ways. Another possibility is to broadcast short segments of messages along with music on a local FM radio channel. The College of Computing community, however, does not have a shared physical space that lends itself well to displaying awareness information. In addition, the cost of these technologies and related licensing fees are prohibitive. After evaluat-

ing these constraints, we decided that the computer desktop would be our primary deployment platform.

Opportunistic interfaces on the desktop have potential because people already use such tools on their desktops to opportunistically monitor time, e-mail queues, system loads, and other data sources. Therefore, a desktop community awareness application would not be a completely alien concept. On the other hand, because of the array of information monitors that people already use, introducing another monitor becomes more difficult due to tighter screen real-estate constraints and higher user expectations on the quality of the new tool. Nevertheless, opportunistic interfaces do not have to stay on the desktop. We simply utilized this approach as a practical solution.

CHAPTER 4
The "What's Happening" Communication-bar

This chapter describes the first desktop opportunistic interface that we have created. We will start by showing how the "What's Happening" communication-bar is used and illustrating its user interface components. We will then describe its system architecture and implementation. Lastly, we will recount how the interface has evolved based on the feedback received from users.

4.1 System Overview

The "What's Happening" communication-bar is one of our community awareness systems. It is deployed on a user's desktop and potentially always visible (see Figure 8). The program shows short pieces of information about the local community such as announcements and calendar of events, as well as summaries of external content such as news reports and weather forecasts. These types of content are automatically collected by a server program and retrieved by the communication-bar clients. They typically last a day or until they are removed from the data sources. In addition, users can contribute content either by posting new stories or by following up on existing content in the built-in chat-rooms. A user posting can last from an hour to a week. Chat-rooms, on the other hand, are automatically cleaned up after being inactive for two hours.

While the communication-bar has a small size on the screen to minimize distraction when not in use, it supports a range of user activities to facilitate quick interaction when it is used by chance. Some examples of the activities include listing and flipping through available content, bringing up a standard Web browser to read more about any particular content, adjusting how frequently the program changes the display to show different information, selecting what external content should be shown, and marking a specific piece of information so that it would not be displayed again.

More specifically, different pieces of content are presented as short *blurbs* in the communication-bar and shown one at a time, in a cyclic manner. After displaying a blurb for a certain amount of time, and without any user input such as scrolling down the text, the

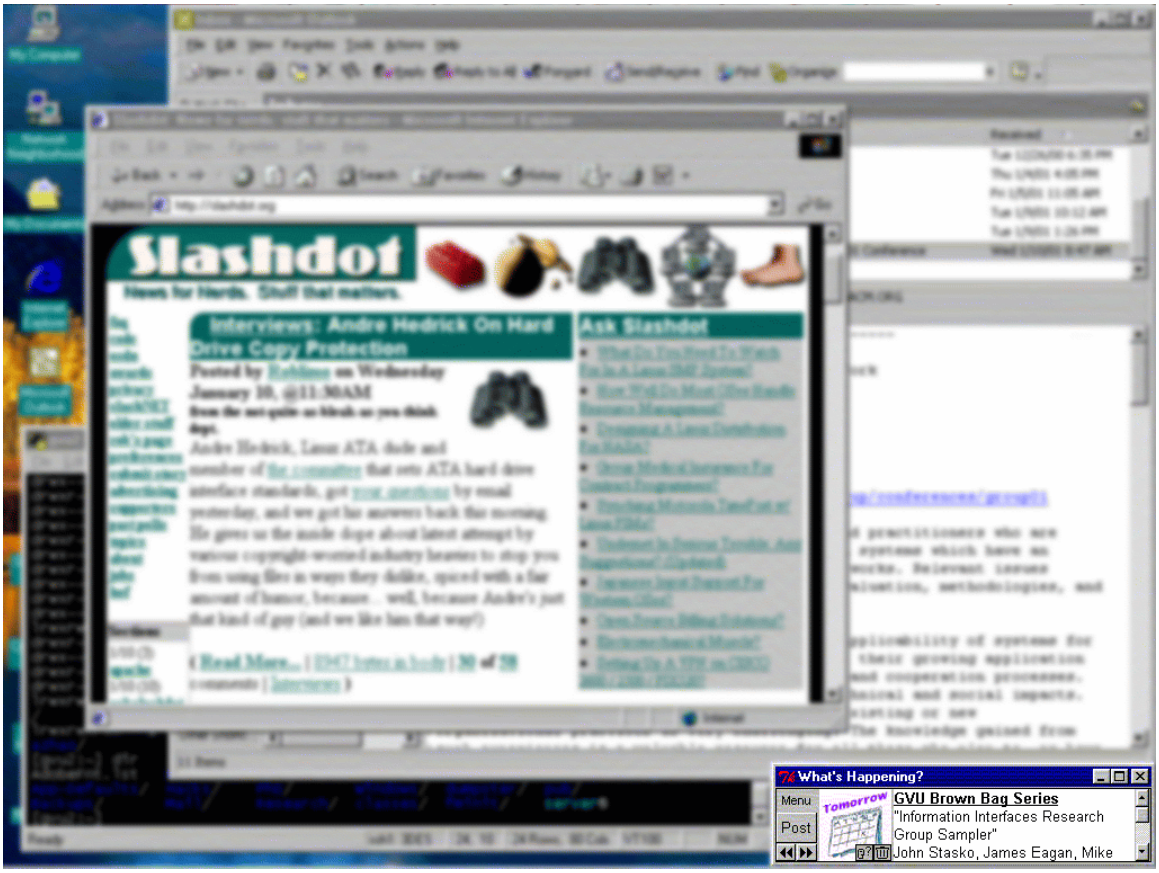

Figure 8. The communication-bar at bottom-right corner of the desktop.

program automatically replaces the old blurb with a new one using a smooth wiping animation. The delay is approximately one minute by default and it is adjustable.

We use multiple levels of detail to manage the display of the blurbs. The most prominent feature that the user sees is an image depicting the topic, source, or status of the current information, so that with a quick glance, a viewer can decide whether to pursue the information further or switch to other tasks (see Figure 9). A red flag at the top-right corner of the image signals that the blurb is being shown for the first time (see Figure 10 for several examples of different blurbs). The title of the blurb and a short summary provide the user with progressively more detail about the information. In addition, the title of a blurb signals the scope of the information: local content uses a bold font, external content

(a) Images conveying the topic of the blurb.

(b) Images conveying the source of the information.

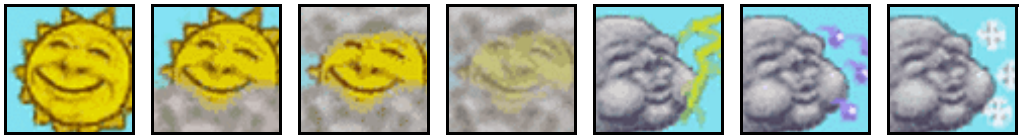

(c) Images conveying status such as the current weather condition.

Figure 9. Example blurb images.

uses the regular font. Finally, the user can bring up a web browser to see the original full-text article by simply clicking on the title.

Small iconic buttons at the bottom-right corner of the content image provide quick access to actions that can be performed on the current blurb. The trash-can button puts the blurb away so that the program will not automatically show it again in the future, even though the blurb is still stored in memory and can be accessed through a list of available blurbs. The head-and-question-mark button shows the activity level in the chat-room for the displayed blurb: the bigger and darker the head, the more recent that someone added to the chat-room discussion. Clicking on the head-and-question button brings up the corre-

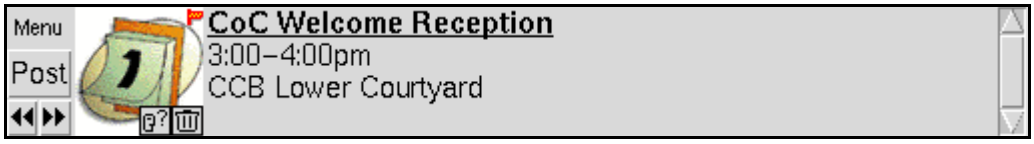

(a) A community event announcement.

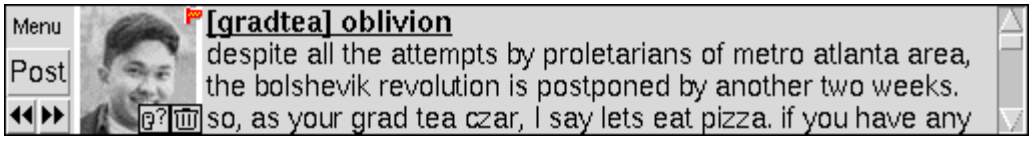

(b) A user-contributed grad-tea announcement.

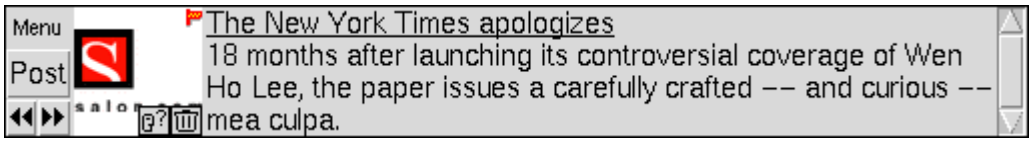

(c) External content: a Salon blurb in this case.

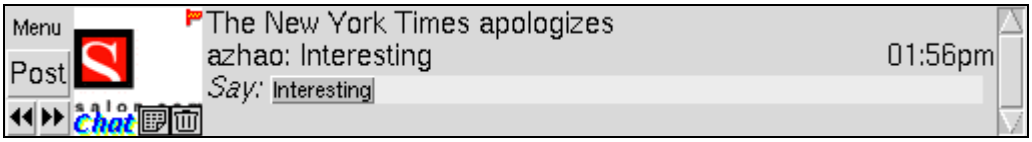

(d) The chat-room corresponding to the Salon blurb in (c).

Figure 10. Close-ups of the communication-bar.

sponding chat-room, which replaces the chat button with a document button that links the chat-room with the web page for the original blurb.

The buttons to the left of the content image control the general operation of the communication-bar. The “Menu” button brings up the application menu, which contains items to pause or resume automatic cycling, list available blurbs, and pop up the preferences dialog. The “Post” button brings up the post-a-blurb dialog. Finally, the two arrow buttons allow jumping to the previous or the next blurb in the cycle.

The preference dialog allows the user to customize the application to fit their needs (see Figure 11). A user can choose either horizontal or vertical layout (see Figure 12 for a

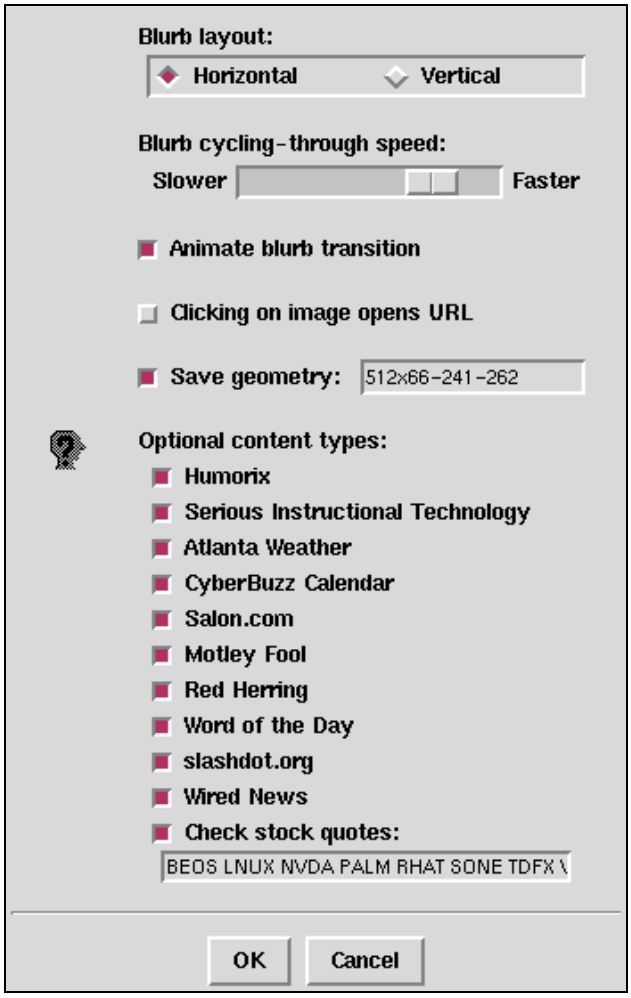

Figure 11. The preference dialog.

close-up of the vertical layout), adjust how long the program pauses when displaying individual blurbs, select whether the transition from one blurb to the next should be animated or not, and specify whether clicking on the content image should directly bring up the web page. In addition, users can save the geometry of the application through the preference dialog, thus enabling the main interface to appear at a fixed location on the desktop every time the program starts. Finally, the user can select the external information sources to monitor through the interface. Since community related blurbs should have higher priority than those from external sources, local content will always be automatically cycled.

The post-a-blurb dialog lets the user to submit a new blurb to "What's Happening" by specifying a title, a brief summary, and an optional web address for any details (see Figure 13). Since the precise moment that a blurb expires in "What's Happening" is usu-

Figure 12. Vertical communication-bar screen layout showing a "Happy Birthday" blurb. The height is adjustable when the application is laid out vertically.

ally not a critical matter, the communication-bar presents a few choices of expiration time instead of requiring the user to type in a time, hence reducing the effort involved in posting a blurb. The image shown to the left of the dialog will be displayed along with the submission. It can provide a space for self-expression in similar ways that the "zsigs" in Zephyr do for its users [2]. Clicking on the "Browse" button brings up a standard file browser that allows the user to select an image to use. The "Upload" button propagates the new image to other communication-bars.

4.2 Implementation Issues and Techniques

Structurally, "What's Happening" consists of a server running on a dedicated Unix system and client programs running on users' computers. The content server handles collecting data and extracting information from local and external sources, broadcasting information to clients, and relaying chat messages among the communication-bars. The client program is responsible for receiving content from the server, customizing the presentation of the blurbs according to user preferences, and submitting new blurbs and chat messages entered by the user to the server.

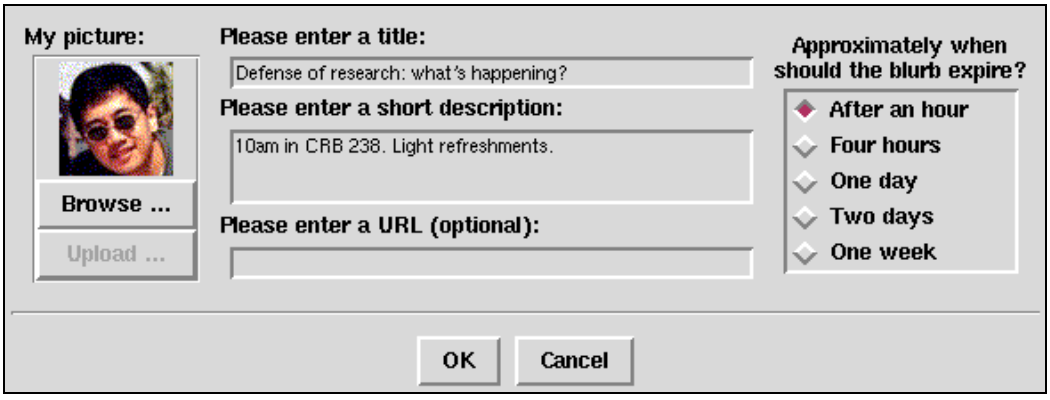

Figure 13. Dialog for posting a new blurb.

Much of "What's Happening" was implemented using the Tcl/Tk scripting language [88]. We chose Tcl/Tk for two main reasons. First, people in the College of Computing use an array of various flavors of Unix, Windows, and MacOS systems, and we needed to support as many different platforms as possible. Using a cross-platform development environment was more cost-efficient than developing in different environments on different platforms and making the different instances of the application not only conform to the same interface design but also collaborate with each other. Not only was Tcl/Tk available on all of the platforms that we wanted to support, but also the *same version* of Tcl/Tk was available on all of the platforms at the same time, providing superior interchangeability across platforms than alternatives such as Java [7]. Furthermore, Tcl/Tk requires much less computing power and memory usage than other environments like Java, allowing us to deploy the system on older computers with slower processors and smaller memory configurations, which were common in student offices.

The second reason for using Tcl/Tk was that we wanted to leverage the power of existing software instead of building a completely new environment. Tcl/Tk allowed us to easily *glue* small pieces of software together to build complex system logic [87]. This was especially true in a Unix environment in which the "What's Happening" server invokes other programs to process web pages and incorporates the results into usable content.

One of the challenges of building a client/server software system is to design a protocol that allows the clients and servers to easily communicate with each other. The "What's Happening" server and clients communicate simply by transmitting Tcl scripts in plaintext, bypassing the issue of extracting data from network packets. For example, a packet sent to a client may contain the following text:

```
append2ChatItem chat-key-123 {quartz: Lunch?}
```

Here, the server is telling the client to append the text to the end of the chat-room identified by the key. The first word is always the name of the Tcl procedure that will further process this packet. The rest of the packet is passed to the procedure as arguments. Curly braces prevent Tcl from splitting up the words in-between into separate arguments. In this example, the client procedure `append2ChatItem` will receive two arguments: "`chat-key-123`" and "`quartz: Lunch?`". See Appendix B for a detailed description of the client/server communication protocol.

Evaluating commands received from the network poses a security risk: malicious code can be embedded in the commands and cause interruption of services or damage to data when executed. In order to prevent the "What's Happening" program from executing malicious code, all commands are evaluated in a protected environment known as the Safe-Tcl interpreter.

Another challenge in building "What's Happening" involved extracting needed information from web pages. We wrote specialized parsers for several content sources using the Tcl built-in regular expression facility because those pages did not conform to any predefined format. This method incurs high programming costs. When a web site changes format, the parser for that web site must be re-implemented to accommodate the change.

On the other hand, several of the external content sources have XML backends that use the XML self-describing markup format [124]. Although different web sites may use different tags to mark up the same types of data, or they may structure the XML documents differently, building an XML parser is relatively easy because different types of data are

clearly separated. Specifically, we reused the TclXML extension [9] to implement a single parser for all XML content sources.

Putting the various features of the interface together and designing the presentation of multiple blurbs was an iterative process. We had many paper prototypes and brainstorming sessions before we built our first computer-based prototype. During the development, we had to weigh many different choices in the design space and sometimes made compromising decisions. Since at many times the advantages and disadvantages of different aspects of the design and implementation were not apparent, we had to iteratively revise and build upon previous generations of the system.

Incidentally, one piece of this evolution concerned the system's name. The first several iterations of the tool was called "News Flash". When we presented the system to colleagues and asked for input, some commented that the name had the same connotations of popular Internet news view-bars that used flashy animations to grab user's attention. As opposed to what that name suggested, the main purpose of this opportunistic interface was to enhance community awareness *without* distracting people from their regular activities. After discussing several alternatives, we chose "What's Happening" which captured our design goals more appropriately.

Since the "What's Happening" communication-bar would inevitably compete for user's attention with other applications running on the desktop at the same time, we needed to make it occupy the least amount of space possible, yet still be big enough to display awareness information. Instead of trying to find the *right* size which was an impractical exercise in trying to satisfy every taste, we decided to give the communication-bar a thin rectangular shape with a fairly small initial size so that it could be easily *docked* to a

corner of the desktop. The user could adjust the width of the application and make it bigger if they wanted.

We now describe the evolution of several aspects of the communication-bar. Reviewing some of the design choices considered and utilized at difference stages of the research helps to identify key issues in the creation of community-aware tools.

4.3 Interface Design Evolution

Figure 14 shows the first generation communication-bar layout. In parallel with the left-to-right English reading order, we placed the most abstract components of the interface to the left of the window and put progressively more concrete components to the right. Initially, we classified six categories of blurbs that might be available in our application and designed small iconic buttons that would allow the users quickly select the types of information that they might be interested in. The six categories were: community events, general news, user postings, traffic, financial news, and sports news. The corresponding buttons were laid out in a two by three array and were aligned to the left of the application window.

Next to the category buttons were three navigational buttons that would allow the user jump to the previous blurb, pause or resume automatic cycling of the blurbs, or jump to the next blurb. The 64-by-64-pixel content image, the title, and the summary were placed

Figure 14. First generation communication-bar screen layout showing a news blurb. The horizontal bar at the top was part of the window manager decoration, not part of the communication-bar interface.

next. We did not include a scrollbar in the design because we envisioned that mouse pointer entering the application window would signal that a user is reading the blurb and would cause the text to automatically begin smoothly scrolling up, revealing more of the blurb. Finally, three convenience buttons for quickly exiting the application, toggling the application to the front of the desktop or the background, and bringing up on-line manuals were embedded to the right of the application window.

In initial feedback, the image size was generally considered a good choice — small yet still comprehensible on 17" and 19" monitors that were popular in the College at the time. However, people pointed out that the interface was quite "busy". The presence of many buttons was distracting. Trying to display five lines of text made the font too small to be read comfortably. The automatic scrolling of the text could not accommodate the diversity in reading speeds and might make user feel not in control of the interface. Without help, people did not quite understand the meanings of the icons. Important features such as submitting a new story or accessing to the application settings were hidden in the interface. Finally, people wanted a way to list current stories available in the application and allow random access to specific blurbs, instead of having to sequentially go through available blurbs.

Based on this feedback, we simplified the design and created the second generation screen layout. We trimmed the categories down to two: news and chat. A category button in depressed state meant that the corresponding category would be included in the presentation. We changed the pause-resume button to a shape that mimicked the "repeat" buttons often found on home audio devices. The help button became a menu button which would include menu entries to bring up user preference dialogs and on-line manuals. Exit and

front-back toggle buttons that duplicated functionalities in window managers were removed. In addition, a scrollbar was added to the text area. As a result, the number of content selection and navigational buttons was greatly reduced (see Figure 15).

We used overview lists to support randomly accessing individual blurbs. Clicking on the minus sign in a news blurb would collapse that blurb and reveal a list of news titles. Then, clicking on a plus sign before a blurb title or double-clicking on the title would expand the corresponding blurb to fill the display area.

Since conversations in a chat-room sometimes move away from the original topic, the title of a chat-room does not always accurately reflect the current state of the discussion. The communication-bar would monitor what was being said in a chat-room and pick out keywords in the sentences, and the most frequently said keywords would represent that chat-room in the overview (see Figure 15-b). Based on either pre-specified user prefer-

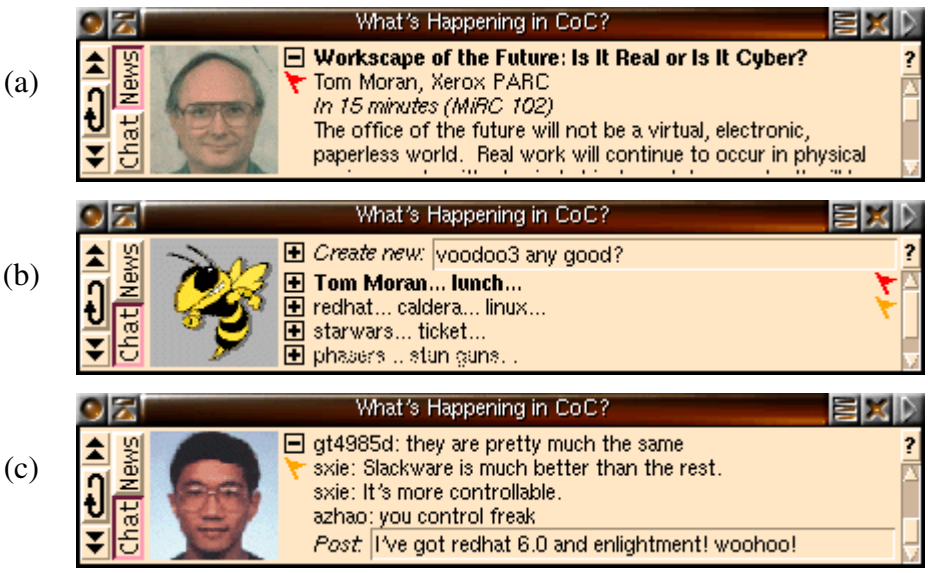

Figure 15. Second generation communication-bar screen layout. Three types of blurbs are shown: (a) a news blurb about an upcoming event, (b) a list of chat-rooms annotated by keywords with a facility to create new chat-rooms, and (c) the bottom of a chat-room with a text-entry widget for adding to the conversation.

ences or automatically tracking what the user had read or said, the program would recommend potentially interesting blurbs to the user by marking them using red flags. In addition, the program would highlight the chat-rooms where people had said something recently, and fade out those that seemed to be forgotten.

The communication-bar embedded text entry fields to support chat-room creation and adding to a conversation. The "create new" entry in the chat-room overview would create a new chat-room and make the entered text the first sentence in the chat-room. Similarly, the "post" entry at the bottom of the chat-room display would append the entered text to the chat-room discussion.

Reactions to the second generation screen layout were mixed. Although the appearance of the interface seemed simpler, the overall system design became more complex. Laying out jump-back and jump-forward buttons vertically did not map well to common controls on home audio devices where people would be most familiar with the horizontal orientation. Categories and filtering modes were confusing. The scrollbar did not align well with the text area and potentially gave the wrong impression of what could be scrolled. The advanced features, such as the recommendation system or hierarchical viewing support, would inevitably impose high development costs and might not provide much benefit to our users.

The consensus of user feedback was to simplify both the appearance of the interface and the interactions that the interface would afford. In our third generation design, navigational buttons were placed horizontally and only jump-to-previous and jump-to-next buttons were available. We allocated some of the saved space for a menu button and a post-a-blurb button. The menu included entries that were unlikely to be used often, such as stop-

ping or resuming automatic-cycling of blurbs, listing available news or chat blurbs, setting application preferences, etc. The post button was designed to provide an explicit, quick access to user content submission, in the hope to encourage such activities (see Figure 16).

We replaced the proposed comprehensive on-line manual by help tips that would be randomly selected and displayed in the common text area at application start-up. The reason was that the application should be simple enough that it should not need any help system to be usable.

Instead of supporting user selection of broad news or chat categories of blurbs to see, we opted to support user selection of specific data sources to include in the automatic cycling. One could, for example, select weather forecasts to be cycled in the application, but not Slashdot news [102]. Since we did not anticipate that users would adjust their selections very often, we put the choices in the user preferences dialog box (see Figure 11 on page 57) instead of on the main window as before.

If the user saw a particular blurb that was uninteresting, but the user did not want to filter out everything from the same data source, the person could click on the “forget this article” button to prevent the application from showing it again. Clicking on the “go to web page” button would bring up a full-text version of the article. Since these two buttons were associated with the specific blurb being displayed, they scrolled with the text and did not always appear on screen, further simplifying the display area.

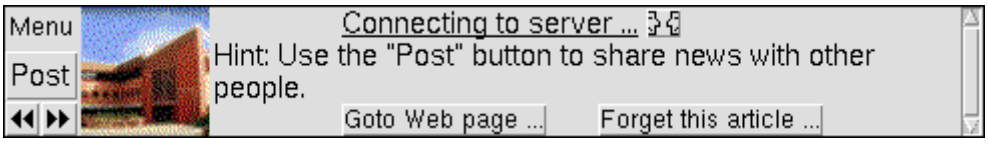

Figure 16. Third generation communication-bar screen layout. The application shows a usage hint when trying to establish a connection to the content server.

An article from a local source such as the College calendar of events would be more likely to be related to the local community than a blurb from an external source such as Slashdot. Visually differentiating local content from external content could raise the comprehension level after a quick glance at the display. Therefore, we chose bold font to display the titles of local content and regular font for other content types.

Creating chat rooms that did not relate to any content in the application was another activity that we did not anticipate the users doing. On the contrary, we expected that people would want to talk about information that they read on "What's Happening" and it was necessary to support quickly moving to the chat room if the user had something to say. In addition, we felt that there were incentives in providing an indicator [3] to show the activity level in the chat room corresponding to a specific blurb. Consequently, we designed a small iconic chat-activity button next to the blurb title (see Figure 16). The icon was composed of two faces looking at each other. If messages were recently added in the chat room related to the displayed blurb, the lines in the icon would become thicker. Clicking on the chat-activity button would bring the user to the corresponding chat-room.

The third generation screen layout was presented at an informal research seminar with an audience including people from other disciplines such as psychology, communication, and new media. After receiving positive feedback on the design, we decided to conduct a small scale deployment and observe the interface in real-life use. An e-mail was sent to the graduate students, faculty, and staff members in November 1999, announcing the availability of the "What's Happening" communication-bar.

Our users provided us with many valuable suggestions that led to further refinements to the interface. People commented that, in the third generation interface, they had to

scroll down to the bottom of the summary text in order to use the web-button or the remove-button, which was a tedious activity that needed simplification. In addition, those two text buttons occupied a fairly large amount of space. Many people also commented that they did not recognize the chat-activity icon and did not know what it would do. Based on these comments, we decided to change the web and remove text buttons to iconic buttons reminiscent of a document icon and a trash can icon, respectively. We thought that the icons would translate well to the actions that these buttons were associated with. We also moved those buttons to the title line, and changed the chat-activity icon to a design reminiscent of the user-group representations in MacOS. In addition, we implemented tool-tips that would pop up when the mouse pointer moved into the buttons and explain what action was associated with the button (see Figure 17).

Putting the small context buttons on the title line, however, did not effectively reduce unnecessary scrolling. Instead of requiring the user to scroll to the end to use the web-button, for example, the interface then required the user to scroll back to the beginning of the blurb if they had read the summary and scrolled past the title line. We revised the interface

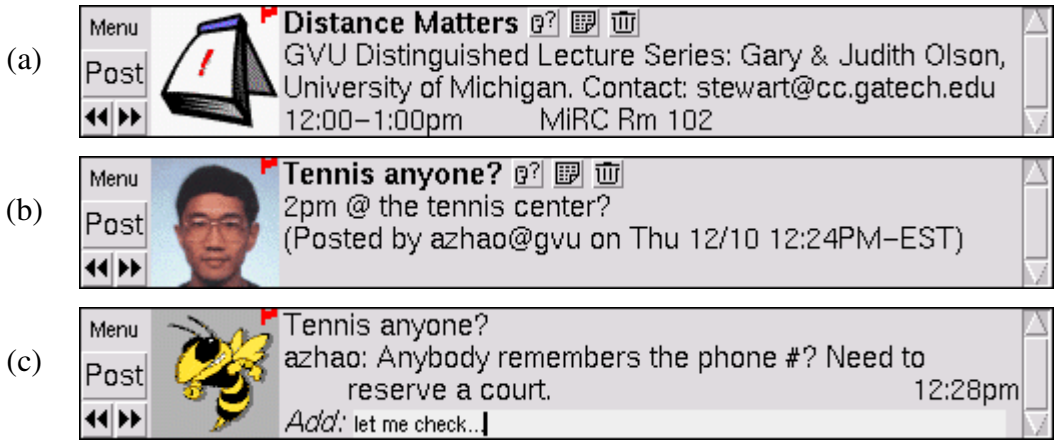

Figure 17. Fourth generation communication-bar screen layout. The application displays a calendar blurb in (a), shows a user submitted blurb in (b), and demonstrates participation in a chat blurb in (c).

further and used several techniques to help solve this problem. First, title lines were underlined and changed to behave like an anchor in a web page: if the mouse pointer moved into the title area, the title text would be highlighted; and if the user clicked on the title, the application would invoke the link action and bring up a web browser window to show the corresponding full-text article. This eliminated the need for a web-button on the title line. Second, we gave the user the choice of clicking on the image to bring up the full-text view. And finally, we moved the context buttons to the bottom-right corner of the content image. Since the image did not scroll with the text, the context buttons became always visible regardless where the user had scrolled to. Since the context buttons were relatively small, they did not occlude the content image much and did not affect the comprehensibility of the content image (see Figure 10 on page 56 in contrast to Figure 17).

Another set of comments suggested that a tall, narrow rectangular application would fit some desktop configurations better than a thin, wide one. Correspondingly, we added a choice that allowed the user to configure the interface vertically and adjust the height of the application window (see Figure 12 on page 58).

4.4 Design Issues

This section describes in detail some of the issues in designing the "What's Happening" communication-bar, the choices we made, and the rationales behind these choices.

4.4.1 Ordering the Presentation

A subtle, but important, issue in the communication-bar is the order that information blurbs appear in a cycle. The easiest way of deciding the order of the blurbs to appear is natural ordering where new blurbs are simply added to the end of the existing sequence.

Other methods may include alphabetical or random ordering. One of the drawbacks of these methods is that some of the blurbs from a single source may be scheduled next to each other in the sequence. Since these blurbs come from the same source, they are more likely to have similar topics. For example, several blurbs all related to Linux [68] from Slashdot may be displayed sequentially. In this case, watching the application may become repetitive and can lower the level of appeal of the application.

The "What's Happening" application should be pleasant to see when the user quickly glances it as well as when they spend a longer period of time looking at it. Therefore, blurbs on different topics need to be mixed well in the sequence of the cycling presentation. However, it is a difficult problem to calculate whether two texts are topically related. We simplify the problem by assuming that blurbs from different sources are unrelated so that we only need to avoid blurbs from the same source being scheduled together.

The algorithm for ordering blurbs works as follows: The program maintains the schedule in a list. Suppose that there are n blurbs, thus making $n + 1$ total *gaps* between two consecutive blurbs and at the beginning and end of the list. The gaps are for inserting additional blurbs. The program also maintains a *next-gap* pointer to the gap where the next insertion should take place. If we need to add a set of blurbs into the schedule, we add the first one to the gap at the next-gap pointer and move the pointer to the next gap that follows. Then we add the second blurb to the gap at the pointer and move the pointer again. We repeat this process until all blurbs have been added. If the next-gap pointer moves past the end of the list, the program creates a new set of gaps based on existing blurbs in the schedule, wraps the next-gap pointer around to the beginning of the list, then resumes. In order to mix all available blurbs, the program starts with an empty schedule, sorts the data

sources in ascending order according to the number of blurbs each source has, and progressively inserts blurbs in each data source into the schedule (see Figure 18 for an example run of this algorithm).

Although formal assessment of this algorithm is difficult, in our observation, it mixes blurbs from multiple sources fairly well. Even though the algorithm is more complex than natural ordering, it does not degrade the performance of the application much since the total number of blurbs is unlikely to be large.

4.4.2 Use of Animation

How the application transitions from displaying one blurb to the next can affect the level of distraction incurred by the user. Directly replacing the current display with the new one is the most straightforward way of changing from one instance of the display to the next. However, we suspected that this would appear similar to blinking and it would be distracting to users. In addition, we hypothesized that animation could help making transitions from one blurb to another less obtrusive. From informally discussing potential animations with several colleagues, we found that the tickering animation, which was popular in Internet news and stock monitoring applications, was distracting and therefore, inappropriate. Furthermore, we felt that the text should be static for a certain period of time after each transition animation had completed, to give the user a chance to read that blurb.

Empirical studies of animations in awareness applications did not exist in the literature when we designed the "What's Happening" communication-bar. To evaluate the suitability of the animations for blurb transitions, we prototyped three animations: fading, rolling, and wiping. Fading uses stippling to give the effect of one display of text dissolving while the next is crystallizing (see Figure 19). Instead of blending the images, which is a CPU-

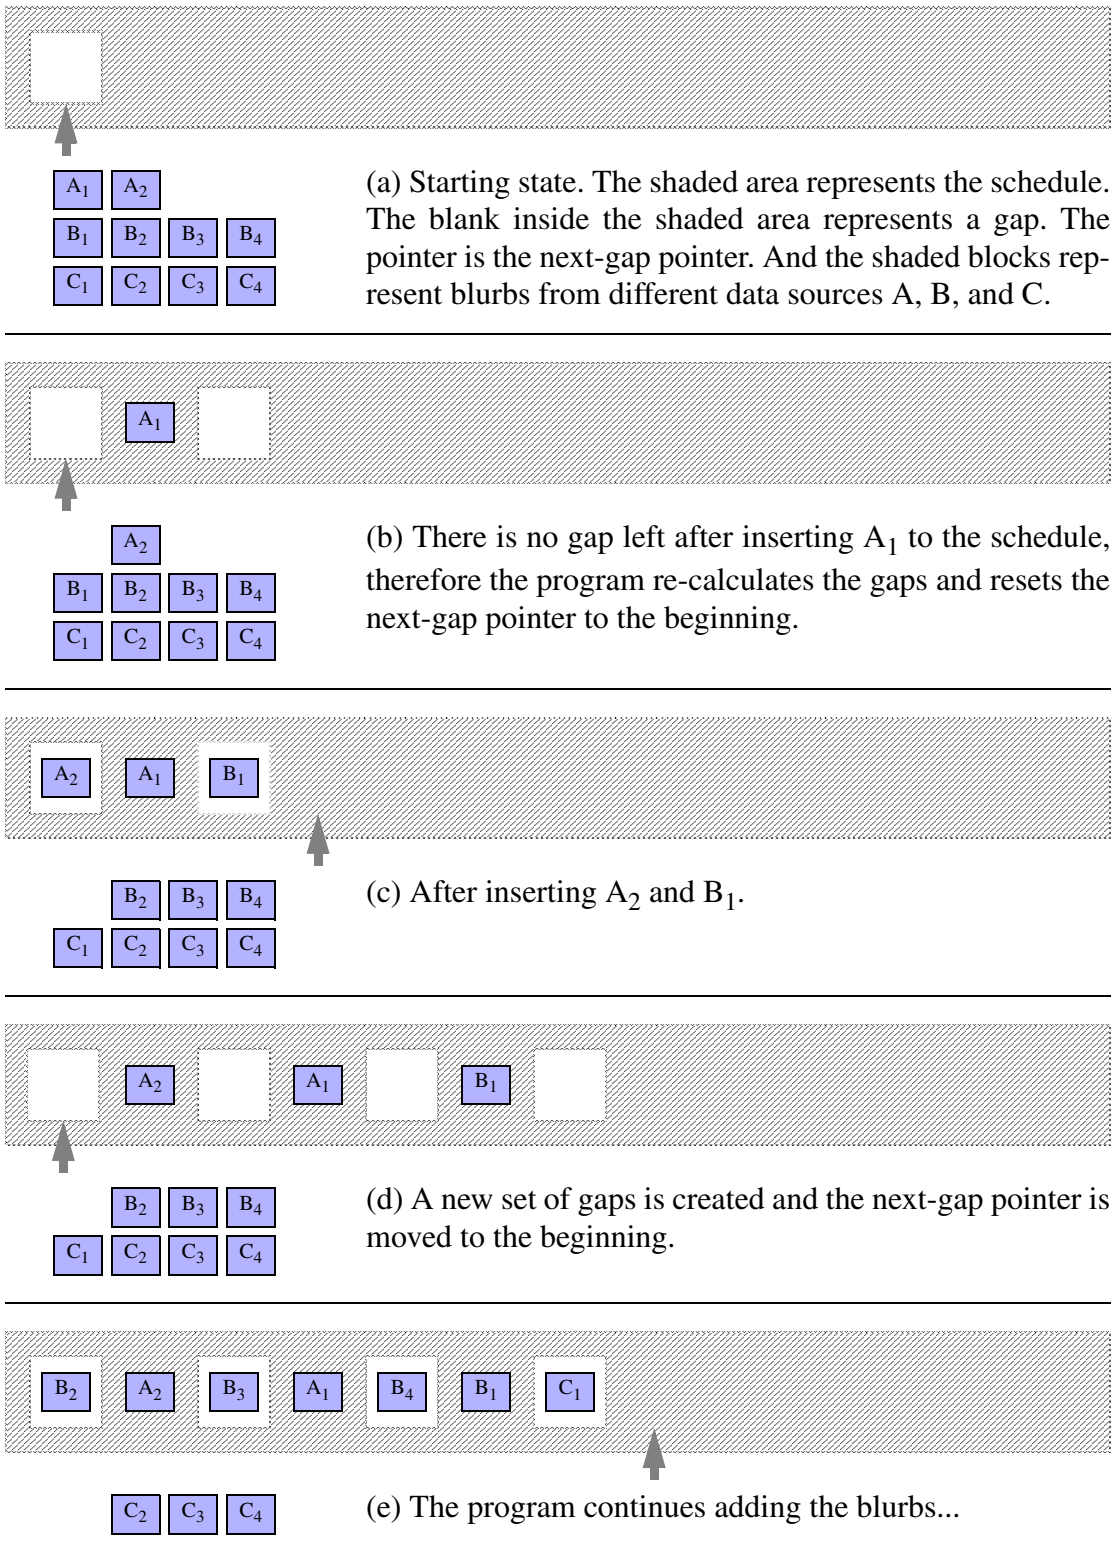

Figure 18. Initial steps to mix three sets of blurbs.

intensive operation, the program divides the image area into square regions and replaces one randomly selected region on the old image with the new one at each step[1]. Both the rolling and wiping animations gradually cover the display area with the new blurb. The difference is that rolling slides the new blurb in from the bottom of the display (see Figure 20), but wiping unrolls the new blurb down like dropping a curtain (see Figure 21).

Trial users who viewed the techniques found that the wiping animation to be the least distracting. In addition, the degree of distraction seemed to correlate to the sum of inten-

Figure 19. A time-lapse series of snapshots showing the fade animation. This shows five snapshots out of sixty-four steps that each round of animation takes.

1. Image fading was later implemented in the Agentk toolkit [76].

sity change at each pixel in the application window during each animation step. Employing this heuristic, we ruled out other animations in popular commercial presentation software such as PowerPoint [78], and implemented the wiping animation in the final version of the communication-bar.

Two separate research projects later confirmed our hypotheses [72, 75]. They found that directly replacing old text with new is intrusive. Instead, using animation can help the application to stay in the periphery with minimal distraction. The studies also found that moving text is more difficult to read and more distracting than static text. This helped to explain why the rolling animation was considered obtrusive. Finally, discrete animation,

Figure 20. A time-lapse series of snapshots showing the rolling animation. The new blurb slides up from the bottom of the screen.

which stops the motion for a moment when the information is in the view and resumes the motion at a later time, is less distracting than continuous motion. This supports our decision to pause between blurb-transition animations.

4.4.3 Content Collection

At the beginning of the development process, the communication-bar had a small set of automatically generated content items including the College calendar of events, local weather forecasts, and Slashdot news. After initial deployment of the system, we gradually added more information sources to the application based on user feedback. Some of

Figure 21. A time-lapse series of snapshots showing the wiping animation. The new blurb is revealed in-place of the old one, starting from the top.

the added external sources included Wired news [122] and user-customizable stock quotes from Yahoo! [126]. The complete list of external sources is shown in Figure 11 on page 57. The number of types of external content has increased to eleven.

When particular content is presented has changed as well. Initially, the blurb about a community event would appear on the day when the event took place. Users suggested that they wanted a chance to be reminded beforehand. Therefore, we changed blurbs to appear one day before the event with a different image labeled "tomorrow".

As we added more external information to the system, it became increasingly important to add local content, in order to avoid community-related information being overwhelmed by unrelated ones. However, some types of local information are relatively easy to generate automatically while other types are not. The *birthday file* is an example among the easy ones. The College of Computing maintains a public database of faculty, staff, and students' birthdays. Since the database has a fixed format, it is relatively easy for a program to search the database for records that match the current month and day, then compose a "happy birthday" message and add that to "What's Happening" (see Figure 12 on page 58). The privacy concern in this example is low because only partial birthday information is in the database: the month and day are available while the year is not. Nevertheless, we provided users the option of not participating in the "happy birthday" announcements.

Another example of easy-to-generate local content is the "happy hour" announcements. A group of graduate students in the College regularly hosts happy hours on Fridays and they use a special mailing list to send the announcements. When a message is received

on this mailing list, an e-mail filtering rule automatically forwards the message to "What's Happening" in the form of a happy hour blurb.

Although community-related messages often appear in e-mail and newsgroup postings, developing a computer program to automatically extract necessary information from an arbitrary text and decide its appropriateness for showing in "What's Happening" is quite challenging. For example, it is difficult to automatically compute the time when the message should expire. Some messages such as seminar announcements have clear ending times. Even if the type of message is known and an ending time is expected in the message, the program may still not be able to extract the time information without advanced natural language understanding algorithms due to the free-form nature of the messages. For those messages that are not clearly time-related, we can only judge their length of stay in "What's Happening" based on the situation. For example, a call-for-discussion message about future expansions of the College should probably stay longer than an announcement about surplus equipment available to be picked up.

Furthermore, different messages sent through the same channel may have different levels of appropriateness. For example, a wedding announcement may be welcomed by many users, but simple congratulation messages that follow may have much less appeal. Reliably judging the broad appeal of an arbitrary text is a difficult problem itself. More unappealing blurbs in the communication-bar may lower the signal-to-noise ratio and the appeal of the overall "What's Happening" system.

Since it is impractical to automatically evaluate the appropriateness and activation length of each e-mail message or newsgroup article, we have to manually choose and forward those that we see fit. Examples of manually forwarded messages include the call for

grad-tea hosts which solicits volunteers to host the weekly social events known as “grad-teas”, meeting announcements that are not listed on the official College of Computing calendar of events, and community related news and discussions not generally available through other communication methods. The coverage and quality of these manually forwarded messages will inevitably depend on how well we know the community as well as our personal judgements.

CHAPTER 5
The "What's Happening" Screen-saver

The "What's Happening" communication-bar provides one form of local information such as announcements and discussions, but there also are other sources of information about the community and its members. One in particular is the set of local web pages about people's research interests, political views, and even hobbies, travels, children, and many other topics. However, with the proliferation of personal web pages, it has become increasingly unlikely that people will browse the home pages of arbitrary strangers, even if they belong to the same community. Thus, showing some of the information opportunistically without requiring much effort from the user may provide more chances for people to learn about each other.

Information in a textual form may not be comprehensible to the user in a short glance. Conversely, images usually provide good representations of the content of web pages and they are easier to grasp than text, especially within a short time limit [49]. Therefore, it may be more effective to show the images on the web pages instead.

One fundamental problem, however, is that images generally do not fit easily in a small space such as that provided by the communication-bar. As an alternative, we decided to utilize a screen-saver as a community awareness tool. A screen-saver is an application that fills the entire screen with dynamically updated images, thus providing the space we need. It is activated after no user input has been sensed for a certain period of time, thus

providing opportunistic access to the imagery that the screen-saver paints. The original goal of screen-savers was to prevent the burning of a static image into the phosphor inside the cathode ray tube after hours of the same image being re-scanned. Even though current display technologies have made burn-in extremely unlikely, many people still use screen-savers so that there is something interesting on their screens when they come back to their computers.

The CollageMachine [61] and Mandala [49] use collages to facilitate browsing large sets of images and corresponding web sites. Based on this idea, we designed a server program to generate image collages. Each collage is a large JPEG image composed from smaller images that the server has collected on local web pages. The program repeatedly creates different collages and copies each one to a fixed location on the network.

We also developed a screen-saver client to periodically retrieve those collages from the network and display them on a screen. More specifically, the screen-saver checks whether there is a new collage every 45 seconds. In general, we cannot predict the frequency that the collages are updated due to the variability of system and network load, as well as the complexity of the collages. The more web images are added in one collage, the longer it takes the server to produce it. In practice, the delay between two updates ranges from 30 seconds to three minutes.

Figure 22 and Figure 23 show several examples of automatically generated collages. We have observed images about people's research, travel, kids, hobbies, pets, and other forms of self-expression.

We now describe the design of the collage server — how it collects images and how it combines the images together.

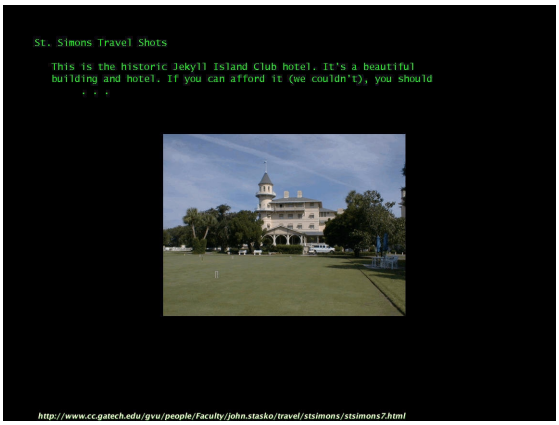

(a)

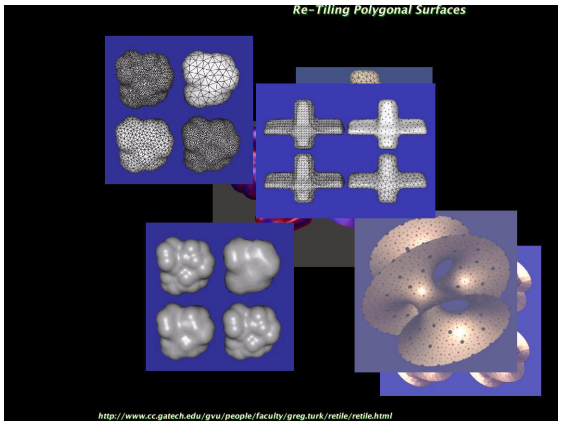

(b)

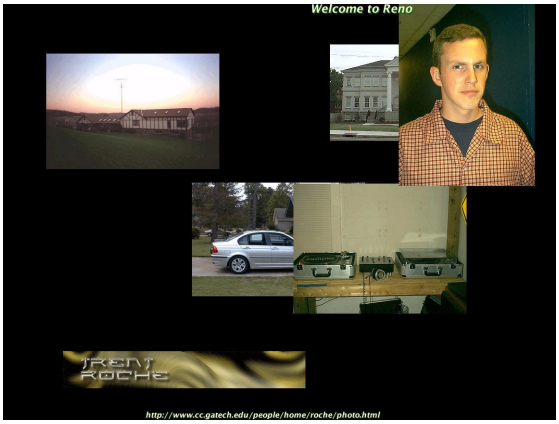

(c)

(d)

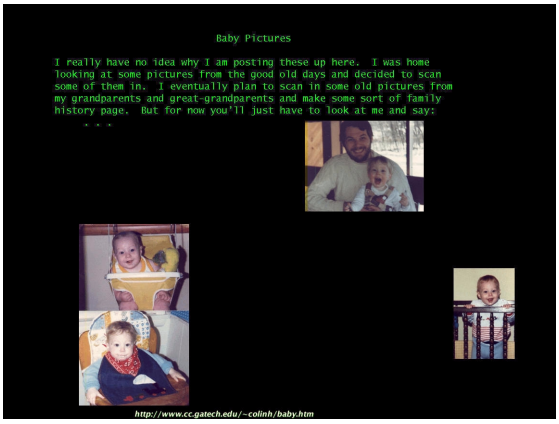

(e)

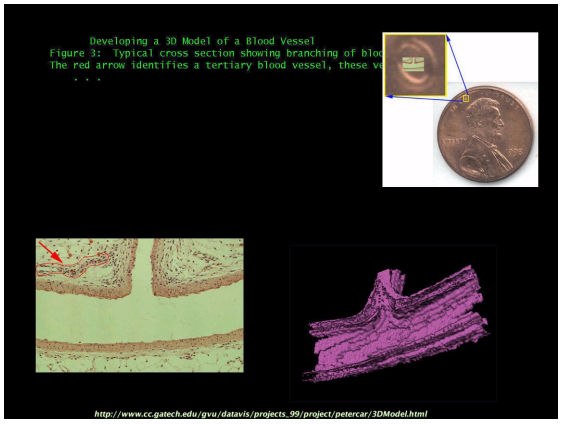

(f)

Figure 22. Example collages.

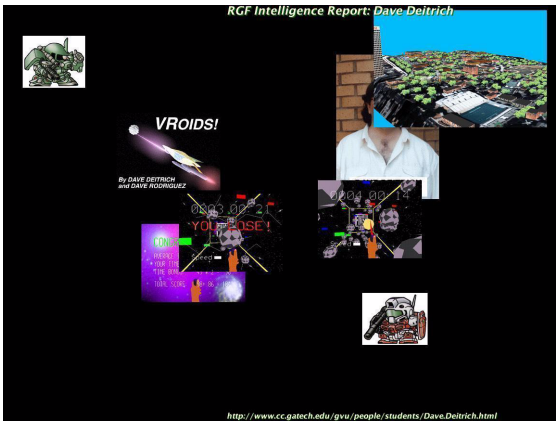

(a)

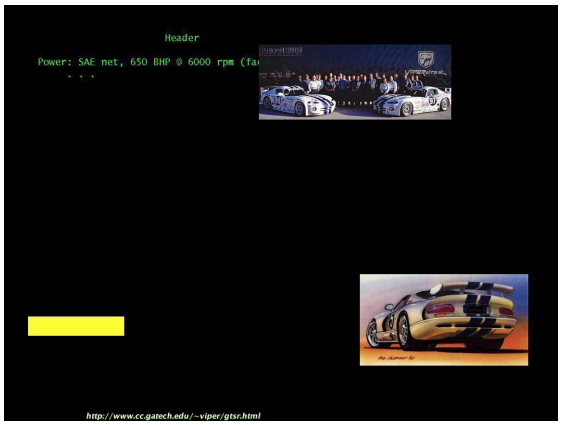

(b)

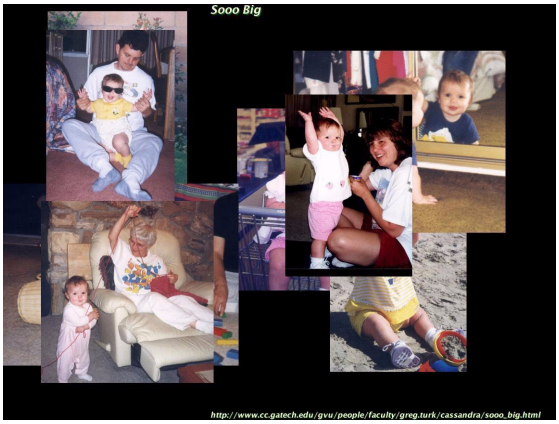

(c)

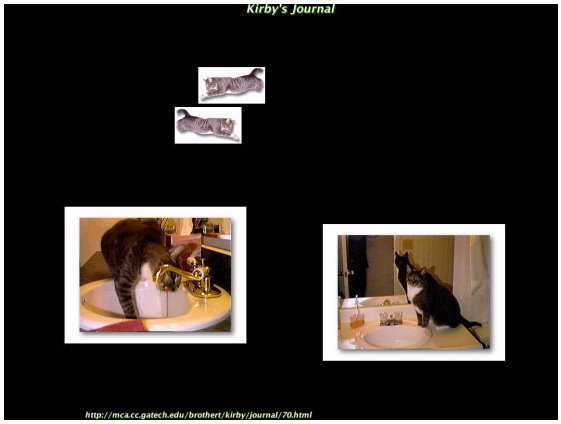

(d)

(e)

(f)

Figure 23. Example collages (continued).

5.1 Image Collection

Since the web tree on the College of Computing web server is accessible from local computers using network file sharing, the first generation collage server directly accessed images in the shared file hierarchy without parsing any referring web page. It defined a *leaf directory* as one that does not have images in its sub-directories. Using a depth-first-search algorithm, the collage server traversed the web site directory hierarchy and checked if each leaf directory contained any images. For each image-containing leaf directory, the collage server randomly selected images in all the enclosing directories along the path from the root directory to the leaf, and sequentially added the images at randomly determined positions to a blank canvas.

For example, suppose that the web tree was mounted as a Unix file system at `/net/www/`, and `/net/www/gvu/ii/community/` is a leaf directory that contains images. The collage server processed directories `/net/www/`, `/net/www/gvu/`, `/net/www/gvu/ii/`, and `/net/www/gvu/ii/community/`, in that order. For each directory, the collage server randomly selected a subset of the images, ignoring decorative elements such as thin separators, narrow borders, and small buttons, in order to avoid cluttering the collage. After adding images from sub-directories on that path, the collage server copied the resulting collage to a particular location for the screen-saver clients to fetch.

This algorithm was based on the hypothesis that directories near the root usually contain images more generic than those near the leaf. These more generic images were effectively placed behind those that have more specific meanings. We hoped that the collage would then convey those meanings better than one that only shows leaf level images.

This initial approach had three critical problems. First, the collage server selected images that were potentially unrelated due to the unpredictable manner in which people organize their web pages and images. For example, some people used a common image directory for a variety of their personal web pages. Randomly selecting images from this common directory often yielded a collage that did not tell a coherent story. Second, the collage server generated collages from images not referred by any page, potentially violating the intensions of their owners. Finally, some community members had separate web servers that could not be accessed by file sharing, preventing the collage server from searching for images on those sites.

Instead of presenting directory-oriented image collections, the current collage server presents actual page-oriented image collections. More specifically, all images added to a collage are embedded in or directly referred from a single web page. Since these images are more likely to be related, the collage based on these images is more likely to form a consistent story, allowing a casual viewer to get a rough understanding of what the collage is about with a short glance.

To gather the locations of web pages and images on all web sites in the College of Computing domain, we schedule an `ht://dig` web crawler [50] to index those sites every week and save the results to a text file. Each line of the index file records a web page and the URL to an image on that page. If the web page has multiple images, multiple lines are consecutively recorded for that page. For example, the index file may have the following segment:

```
http://chi.cc.gatech.edu/ -> http://chi.cc.gatech.edu/images/capture.jpg
http://chi.cc.gatech.edu/ -> http://chi.cc.gatech.edu/images/figure1.jpg
http://chi.cc.gatech.edu/ -> http://chi.cc.gatech.edu/images/figure2.jpg
http://fce.cc.gatech.edu/~bolot/ -> http://fce.cc.gatech.edu/~bolot/me.jpg
http://mca.cc.gatech.edu/kirby/ -> http://mca.cc.gatech.edu/kirby/fwd.jpg
```

The collage server then creates a list of unique web pages, arranges the list in a random order, and processes each page one by one. For each web page, the collage server randomly selects up to ten items from the image set associated with that page, and adds the selected images to a blank canvas to form a collage. Finally, the server draws the title of the page at the top of the collage and draws the location of the page at the bottom. It uses a shadow style to draw the annotations so that they are recognizable on both dark and light backgrounds.

If the web page has less than five images, the collage server inserts the first block of text on that page into the background of the collage. It first uses Lynx [70] to convert the HTML web page to plain-text. Then, it scans the result, removes separators composed in dashed lines ("–"), and skips the blank lines at the beginning. If the first line of text is followed by a blank line, that line of text is potentially the title of the web page and is saved in a text variable by the server program. Otherwise, the collage server searches through the lines of text and appends them to the text variable, until it reaches another blank line or the 16th line of text. Finally, the server converts the saved text to an image and adds it to the collage canvas.

Since the title and the first few sentences of a web page can convey the topic of the page quickly to a casual reader, a picture-with-text collage can alleviate the problem of having a smaller number of images and still tell the viewer what the page is about.

After making the new collage available on the network, the server sleeps for a short period of time to allow the screen-saver clients to update their displays and give the users a chance to see the collage before it is replace by the next one. The delay is the same amount of time as it takes the server to generate that particular collage, or 35 seconds, whichever is

shorter. Complex collages, those with a large number of images or large size images, tend to take longer time to generate. Since they generally take longer time to comprehend, we should give viewers more time on these collages, which is compensated by the longer delay in the server program.

5.2 Image Layout

The first generation collage server simply used the system built-in random number generator to calculate where to put an image into the collage canvas. Since the system random number generator only outputs deterministic pseudo-random numbers, the images tended to overlap and cluster on many of the collages. The resulting collages did not utilize available space efficiently and often were not aesthetically pleasing.

For the second generation server, we decided to spread out the images on the collages more. This algorithm divides the collage canvas into four quadrants and adds a fifth region

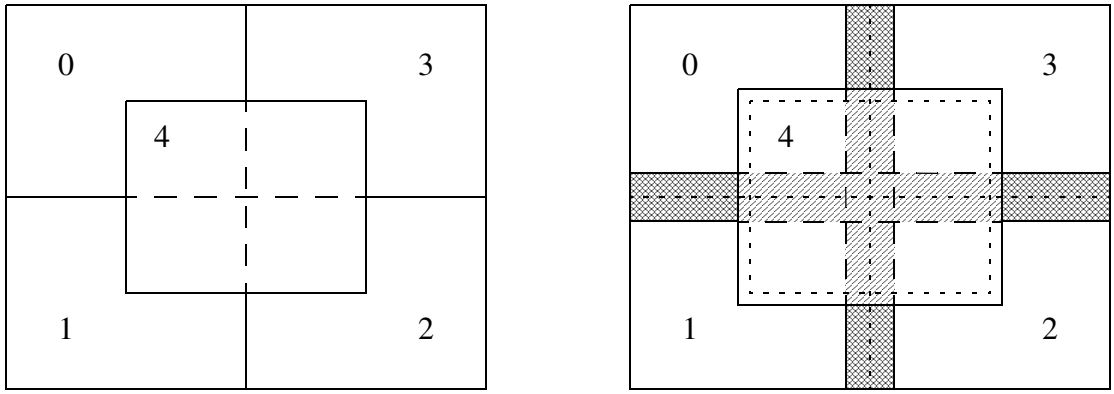

Figure 24. Dividing the collage canvas into five regions. Region 4 is drawn as if it is a semi-transparent sheet. The original regions are shown in (a). The expanded regions are shown in (b) where the dotted lines denote the original region boundaries and the shaded area represents overlapping among regions 0, 1, 2, and 3. Region 4 always overlaps with the other regions.

of the same size to the center of the canvas. The regions are labeled from 0 to 4 as shown in Figure 24-a.

The collage server then adds images to the five regions 0, 1, 2, 3, and 4 sequentially, iterating back to 0 after region 4. For each image to add, the collage server scales the image down to no more than 90% of the size of the region, if necessary, and places the image at a random location within the selected region so that it does not lay outside of the region.

If the number of images on one web page is less than five, at least one of the five regions will be "empty". Therefore, as said previously, we show a small amount of textual information to supplement the images. Conversely, adding ten images to a collage will put two images into each region. Adding additional images is likely to cause increased overlaps, more clutter in the collage, and the result to be less effective in conveying information.

This division method reduced the amount of overlapping and the images in the generated collages appeared more spread out. Trial users, however, commented that collages including four or five large images were not aesthetically pleasing. Some users suspected that when laying out large images with similar aspect ratios, the gaps between the images tended to become similar in size as well. While the images were not laid out exactly symmetrically, viewers might feel that the resulting collage suggested symmetry. Therefore, they might conclude that the collage was poorly laid out. Other users commented that they did not like those collages because of their monotonic appearance (see Figure 25-a).

The current collage server implemented an enhancement to the spreading algorithm that expanded each of the five regions by 10% to allow a small amount of overlap across

regions (see Figure 24-b). Reaction to this collage layout technique was favorable, and people thought that the new algorithm produced more attractive collages (see Figure 25-b).

5.3 Value-added Collages

In addition to web image collages, a screen-saver gives us a chance to provide value-added services to our users. For example, the collage server is a client of the "What's Happening" content server. It generates a collage based on a "What's Happening" blurb every seven minutes (see Figure 26).

Every 40 minutes or so, the collage server also generates a weather collage (see Figure 27). This collage contains images from The Weather Channel [118] that depicts current weather condition and temperature, weather radar scan, as well as the current air quality index [20].

(a) (b)

Figure 25. Comparing collage examples produced with different region set-ups. Image (a) was produced using non-overlapping quadrants. Image (b) was produced using overlapping quadrants.

For entertainment, the collage server generates a webcam collage every 50 minutes (see Figure 27). It contains images from AccessAtlanta [1] that shows scenes from around the city and the current broadcast image on a local TV station, WSB-TV.

Between 4:30pm and 7pm on workdays, the collage server builds a traffic collage every 3 minutes (see Figure 29). It contains a traffic map from Georgia Department of Transportation (DOT), highlighted by average highway speeds. It also includes images from highway cameras so that a viewer can access traffic conditions visually. To help users read the collage easier, the camera images are laid out in a way such that the location of an image on screen corresponds to the physical location of the camera. For example, the camera showing a segment of the highway at the northeast of the city is placed at the northeast corner of the collage.

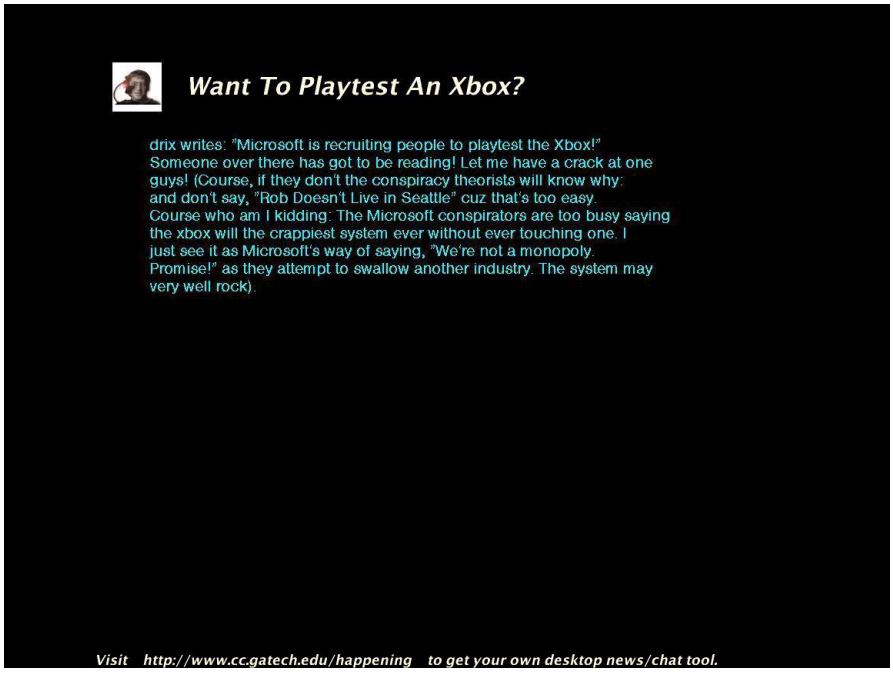

Figure 26. The "What's Happening" collage.

In practice, every other collage shown during rush hours is the traffic one. The higher frequency provides users increased possibility to closely monitor traffic conditions, a more demanding task during that time period.

To summarize, we have developed a screen-saver that opportunistically shows image collages that help people discover parts of the local web space, as well as news and external content that may help viewers in certain aspects of their daily lives.

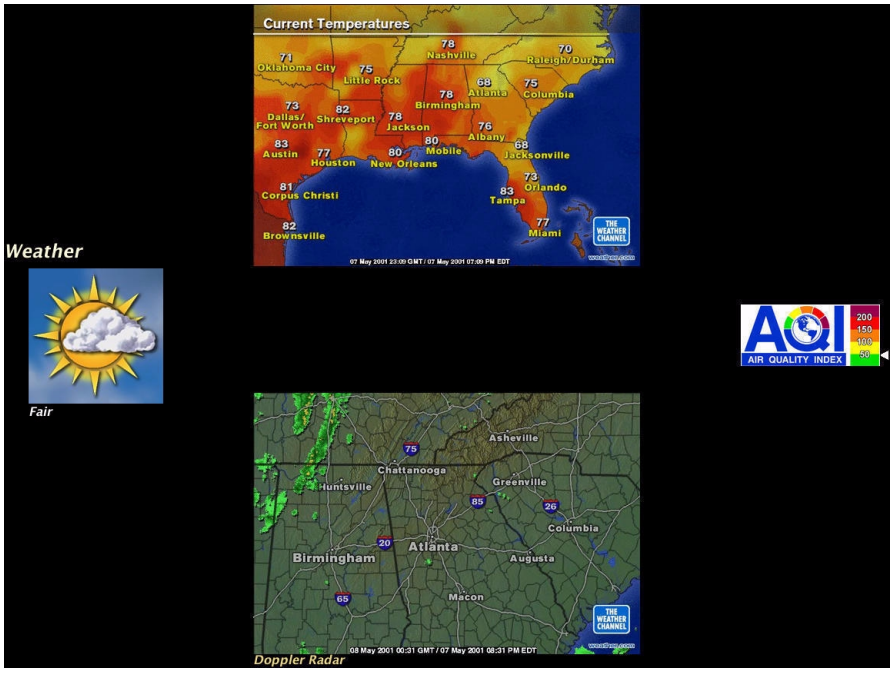

Figure 27. The weather collage.

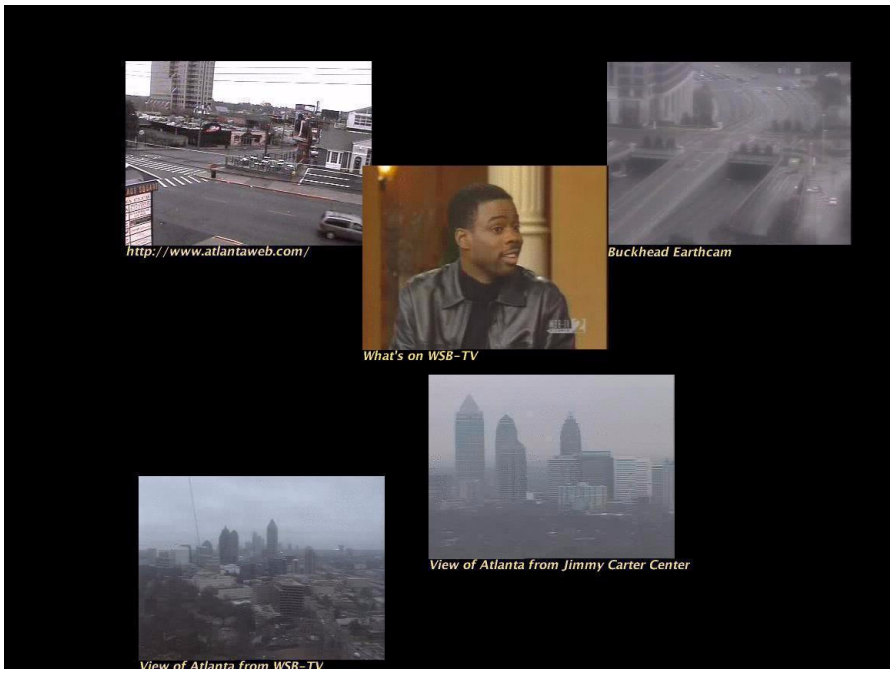

Figure 28. The webcam collage.

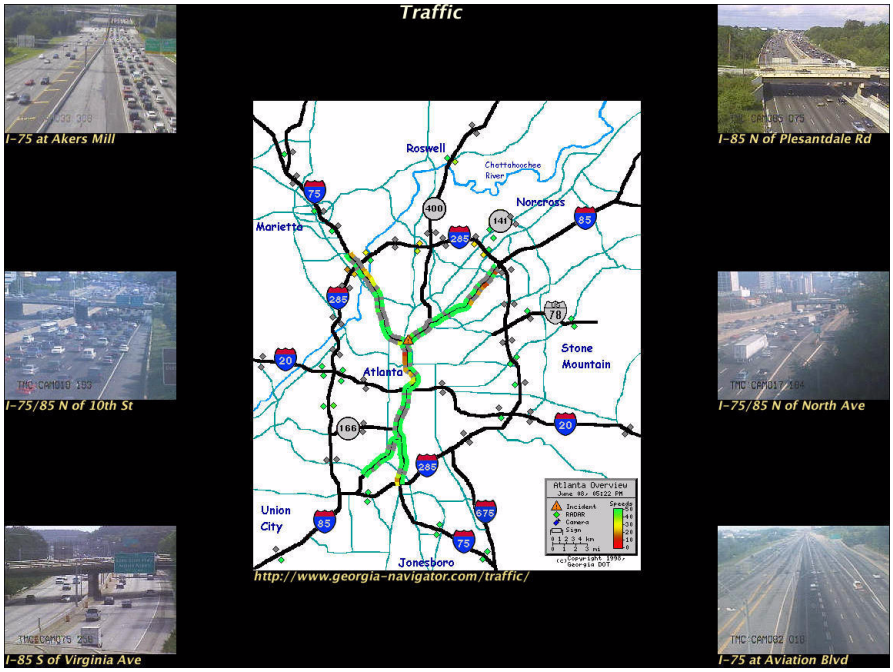

Figure 29. The traffic collage.

CHAPTER 6
Evaluation

Since the goals of our opportunistic interfaces are to enhance awareness and promote communication, we naturally want to assess how our designs impact awareness and communication. Specifically, we want to find out how widely our tools are used and whether they are useful. Furthermore, do people know more about the community and other community members from using the communication-bar and screen-saver? In this chapter, we will study the effects of those two opportunistic interfaces. We will start by reporting the process of deploying the "What's Happening" applications. Then, we will describe the challenges and outline our evaluation methods. Finally, we will present the results in evaluating each of these two opportunistic interfaces in terms of use and effectiveness, and what we can learn from them.

6.1 Deployment Status

We announced the availability of the third generation "What's Happening" communication-bar to the local graduate student, faculty, and staff through an e-mail message in November 1999 [130]. We described the awareness application as an on-going research project and explained that we would record basic usage statistics only for user interface evaluation purposes. We also presented the design in an informal inter-disciplinary research seminar. At the same time we discovered that, although the Tcl/Tk library was

available on Unix, Windows, and MacOS platforms, the identical communication-bar program did not work reliably on MacOS, often causing system faults from within the Tcl/Tk library.

Based on the suggestions that we received after the initial deployment, we revised the system and started a full-scale deployment in March, 2000, by sending a second notification e-mail message and posting newsgroup articles. In addition, we put flyers in people's mailboxes and posted them at various doorways throughout the College.

The Unix and Windows "What's Happening" screen-saver was announced in mid-September, 2000. We also put another set of flyers in the welcome packets that incoming graduate students received from the College, and showed them both the communication-bar and the screen-saver in a class. We put flyers about both applications in a newly renovated break room where many graduate students and faculty members had lunches and informal meetings, the same place where the CCB video wall used to be installed. At the same time, we started manually selecting interesting e-mail messages, newsgroup articles, and local web content, etc. that were difficult to automatically collect, and submitting them to "What's Happening".

6.2 Challenges

To evaluate our opportunistic interfaces in terms of usage, we need to know how many people are using the applications as well as how they use them. In particular, it might be beneficial if we knew when and for how long a user looked at the applications. Given this data, we could test how much reading frequency and reading time would influence awareness change. Additionally, we could use this data to pinpoint specific problem areas in the design.

However, it is difficult to accurately measure when a person looks at an opportunistic interface without tracking the user's focus of attention. For example, we cannot reliably know if and when a person has read a particular blurb on the communication-bar. We can infer that the user is reading the displayed information when the user interacts with the interface, such as when they click the web button, scroll down a blurb, or flip through a set of blurbs. Yet, during the time that the user is not touching the interface, they can still be reading the information shown in the application.

On the other hand, current technologies that support eye tracking are often intrusive and require careful calibration [101]. Therefore, they are not suitable for long term, free-form use and observation. Even if we could monitor when someone is looking at an application, we could not reliably infer that they have read what was shown or comprehended the information.

The usage of the screen-saver is even more difficult to quantitatively measure. At many times when the screen-saver is activated, the user is away from their desk. In addition, the deactivation of the screen-saver does not reliably indicate watching the latest collage, because the screen-saver can be deactivated accidentally or without the user looking at the screen.

It is also difficult to measure precisely the level of community awareness in people. Directly asking whether someone is more aware of the community may not yield valid results. Since this is an unfamiliar evaluative task, people may not be able to accurately report their awareness changes or correctly attribute the changes to using awareness applications. Furthermore, the lack of records of interactions among community members does not allow us to make any strong inferences on community awareness in the same way that

Putnam inferred social capital from such records [93]. Finally, introducing mechanisms for recording such interactions may affect community awareness itself or collide with people's privacy concerns.

Even if we observed enhanced awareness in the community, we may not be able to reliably attribute the effects to using the interfaces. A number of confounding factors may threaten the validity of such inferences [97]. Some of these factors are:

- *Diffusion*. One way of inferring causal relationships is to assign a control group and contrast awareness levels between users of the opportunistic interfaces and non-users. However, knowing that they are not using the tools may prompt people in the control group to pay more attention to community related information through other media, causing the measurements to untruthfully represent the real differences between the two conditions.
- *Selection, Experimenter Effect, and Mortality.* It is impractical to either force a person to use an opportunistic interface or forbid the person from using it. People choose to use the awareness applications for their own reasons rather than being randomly assigned. In addition, the familiarity between a person and the researcher may also influence whether the person uses the applications or not. Finally, existing users may stop using the tools and new users may start. Because of these individual differences, we may not be able to make equal comparisons between a user and a non-user, or between a user and another user.
- *History and Maturation*. Besides the introduction of any particular technology, other events such as a crisis [44] that happen at the same time may

prompt community members interact more and learn more about each other. On the other hand, any observed change of community awareness may be part of the nature process that takes place in a community and could have been achieved without the help of awareness applications.

Despite these confounding factors, we can still ask users to subjectively assess the use of the opportunistic interfaces and their impact. Perhaps the important result is not how the tools might have changed the community, but people's subjective satisfaction toward the opportunistic interfaces. Even though we may not be able to directly prove our hypotheses and can only make imperfect inferences from all evidence that we can find, we can still try to reduce the degree of uncertainty.

6.3 Methods

In order to support finding usage patterns, we built logging facilities in the "What's Happening" content server to record the Internet host addresses of the clients and the time when they connect or disconnect from the server. Since March 2000, we recorded several user activities such as jump to previous or next blurb, visit web pages, post new blurbs or adding to chat-room discussions.

Additionally, since we used the communication-bar and the screen-saver ourselves, we monitored the on-going state of the system and observed the quality and quantity of automatically collected content, and noted user participations.

In the summer of 2000, we conducted seven interviews with people on their use of the communication-bar (see Appendix C for the script that we roughly followed in the interviews). In these interviews, we asked people's opinions on the mechanics of the user interface such as widget placement and function, importance of the image, and obtrusiveness

of the transition animation. We also asked people to generally describe the application and followed their thoughts to find out what they were interested in, what they used the tool for, and what aspect of the community life, if any, that people might become more aware of. Furthermore, we asked people whether they posted blurbs or added to chat-room discussions, and how they used the tool as a communication medium. Lastly, we inquired about ideas of improvement and opinions of the weaknesses and strengths of the opportunistic interface.

Our final survey on the "What's Happening" applications included two parts: one part deployed in April, 2001, to people who were using at least one of the applications at that time, and another part deployed in June, 2001, to those who were using neither of the applications. We re-used the familiarity and sense of community questions in the initial community survey to estimate possible awareness changes (see Appendix D). In the user survey, we also re-used some of the questions in the communication-bar interview to gather quantitative assessments about the tools. In the non-user survey, we asked people why they were not using the applications.

We received 14 responses from our users to the April survey, including those from people who did not fill out the initial community survey. Among these responses, four were from the main College building, CCB, eight and two were from the satellite buildings CRB and GCATT, respectively; 13 of these users answered questions specific to the communication-bar, and 9 answered questions specific to the screen-saver. The surveyed users included one faculty member from CCB and one from CRB, and the rest were all students.

We received 55 responses from non-users to the June survey, gathered by hearsay and occasional intersection: 25 of whom were from CCB, 22 from CRB, and 8 from GCATT. These non-users included two faculty members and three support staff from CRB, and the rest were all students.

6.4 Comparing Familiarity and Sense of Community Estimates

Table 2 summarizes average ratings of people's familiarity with community events and research, as well as ratings of the sense of community in the College of Computing based on the responses to the final survey. Figure 30 shows the histograms of people's answers.

Table 2: Comparing answers to familiarity and sense of community questions. Ratings are given on a 7-point scale, where a rating of 1 represents "very unfamiliar" or "very bad", and a rating of 7 represents "very familiar" or "very good". The neutral rating is 4. The top number in each table cell represents the average rating, while the bottom number represents the standard deviation. The arrows highlight pairs of results that are statistically significantly different.

Question	Initial Survey	Final Survey	
		Users	Non-Users
Familiarity with research events	4.8 σ=1.65	5.6 σ=1.45	4.4 σ=1.59
Familiarity with social events	4.8 σ=1.53	4.7 σ=1.54	3.7 σ=1.60
Familiarity with research in other groups	3.3 σ=1.41	2.4 σ=1.09	3.4 σ=1.57
Sense of community in the College	4.0 σ=1.37	4.1 σ=1.41	4.1 σ=1.12

Comparing the ratings in the surveys, we find that the familiarity with research events among the "What's Happening" users is statistically significantly higher than that among the non-users ($p < 0.01$). Familiarity with social events among the non-users is statistically

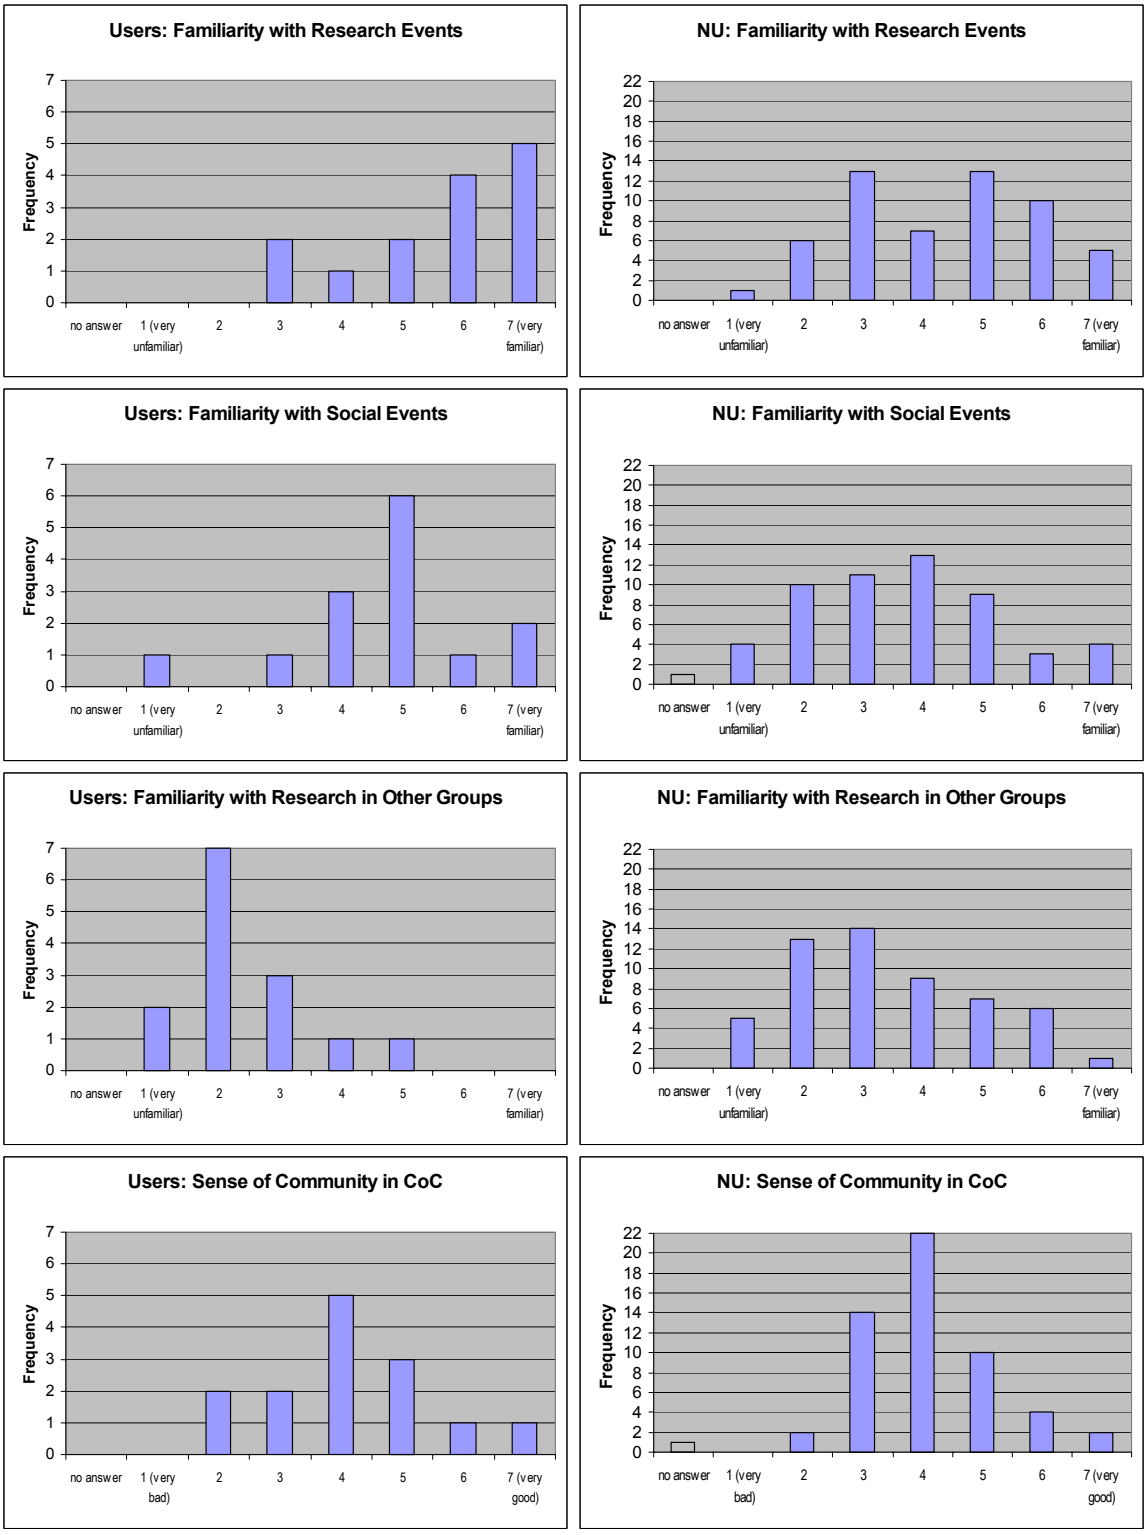

Figure 30. Histograms of answers to familiarity and sense of community questions. The graphs on the left represent results of the user survey, while the ones on the right represent the non-user survey. They should be viewed in comparison to those in Figure 1 on page 14.

significantly lower than the initial average rating and that among the users ($p < 0.01$ and $p < 0.05$, respectively). Familiarity with research in other groups among the users is statistically significantly lower than the initial average rating and that among the non-users ($p < 0.03$ and $p < 0.02$, respectively). Lastly, the sense of community level has remained virtually the same.[1]

Comparing survey responses from the non-users among different buildings, we find that the familiarity with research events in the main building, CCB, is statistically significantly higher than that in the CRB ($p < 0.01$, see Table 3). In addition, the familiarity with research in other groups among non-users in CCB is statistically significantly higher than that in the GCATT ($p < 0.01$).

Table 3: Comparing non-user survey responses among different buildings. Ratings are given on a 7-point scale, where a rating of 1 represents "very unfamiliar" or "very bad", and a rating of 7 represents "very familiar" or "very good". The neutral rating is 4. The top number in each table cell represents the average rating, while the bottom number represents the standard deviation. The arrows highlight pairs of results that are statistically significantly different.

Question	CCB	CRB	GCATT
Familiarity with research events	4.9 σ=1.54	3.7 σ=1.49	4.6 σ=1.51
Familiarity with social events	4.0 σ=1.68	3.5 σ=1.66	3.4 σ=1.30
Familiarity with research in other groups	3.8 σ=1.1.62	3.3 σ=1.49	2.4 σ=1.06
Sense of community in the College	4.2 σ=1.20	4.1 σ=1.17	3.9 σ=0.64

1. Statistical significance was tested using the two-tailed Student's *t*-test.

Even though we should not regard these differences in average ratings as conclusive evidence of any real differences in familiarity, primarily due to the small number of responses in the final survey for "What's Happening" users, we may speculate several reasons that help explain the differences.

In the time between the two surveys, the College of Computing community has become even more distributed. More people have moved into the two separate buildings distant from the main building as the College has continued to grow in size and composition. This physical separation may have been further lessening the familiarity in the community, especially in the two satellite buildings, CRB and GCATT. However, we see a higher level of familiarity with research events among "What's Happening" users than that among non-users. One possible explanation is that, information about research events is often conveyed through official announcements, which in turn are often expressed in similar forms. In addition, such information often contains only a few key points such as time, place, topic, and the one presenter. Therefore, this type of information is relatively easier to comprehend and remember, and the "What's Happening" applications seem to help enhance people's awareness about such events.

Conversely, information about social events is often conveyed through informal channels and may be expressed in many different styles. It often contains additional key points such as the reasons for the event, the theme of the event, directions to the location if it is an unfamiliar place, possibly more than one organizer and their humorous messages to make the event appear more appealing. Therefore, this type of information often takes longer to process and is relatively easier to be forgotten. Data from the surveys seem to suggest that

the "What's Happening" applications can help people maintain awareness about social events while non-users' familiarity with such events weakens.

One additional factor may contribute to the decline in familiarity with research in other groups among "What's Happening" users. While becoming better aware of research seminars and preliminary information about other research groups such as a glimpse of a design diagram, people might become aware of potential information that they did not already know. Since information about other people's research often carries more depth than one can fully appreciate in a glance, this awareness of the unknown might raise people's curiosity level and prompt them to give lower familiarity ratings.

In summary, although our surveys on familiarity and sense of community in the College should not be considered as a thorough measurement of community awareness, they provide us with useful estimates of the states of the community. More complete and precise studies in the future may help us reach more accurate conclusions. Next, we will evaluate each of the "What's Happening" applications in detail.

6.5 Evaluation of the Communication-bar

- **Frequency of Use**

Figure 31 shows the history of concurrent use of the communication-bar over time. From analyzing server logs, we estimate that at the beginning of the initial deployment of the tool, approximately 60 people tried the application in the first several weeks with about 25 becoming regular users who kept the application running on their desktops at all times.

The many suggestions about the communication-bar that we received seem to signal that people were interested in using this tool and seeing it improved. However, usage declined as we approached term project due dates, final exams, and the holiday season in

1999. As the Spring 2000 semester began and people came back to campus, we saw a gradual increase in usage.

Usage peaked at 38 concurrent users when we re-announced the communication-bar and started its full-scale deployment in February, 2000. It stabilized at around 20 concurrent users, perhaps after the novelty wore off. It declined again to around 10 concurrent users when the Spring semester ended and Summer semester started, since many people were off-campus during the summer.

We saw another increase in usage at the beginning of the Fall semester in August, 2000. The introduction of the "What's Happening" screen-saver in September, 2000, did not significantly affect the communication-bar usage.

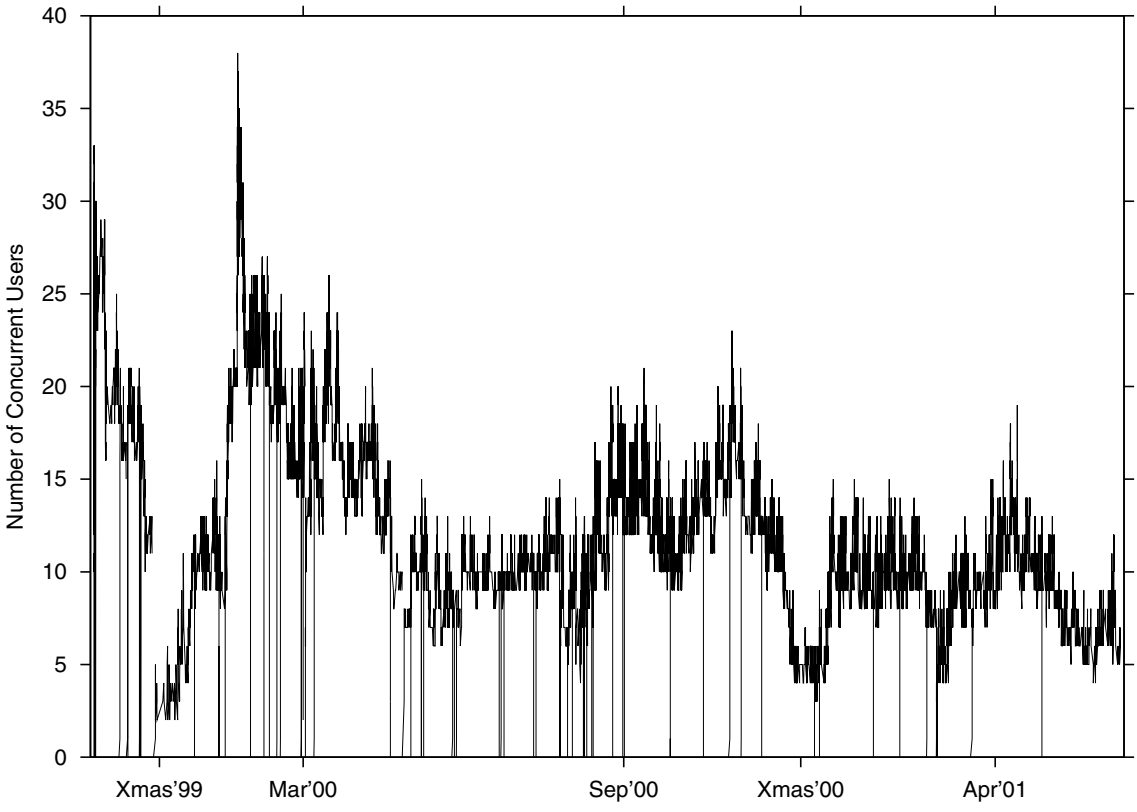

Figure 31. Number of concurrent users graphed over time.

The same pattern of decreasing usage around the holidays and increasing usage at the beginning of the Spring semester was observed at the end of 2000 as well.

- **Usage of Basic User Interface Functions**

Figure 32 illustrates the number of uses per day recorded for the navigational features that include the jump-to-previous button, the jump-to-next button, and selecting a blurb from a list. Event though the data is highly variable and does not show any clear trend, the use of the jump-to-next button seems to be more frequent than the other two methods.

Figure 33 illustrates the number of uses per day recorded for the go-to-web-page button and the hide-this-blurb button. Again, the go-to-web-page button usage data shows no clear trend. We did not record the specific web pages that users requested or the topics of the blurbs, although this may be beneficial in future application designs and evaluations. The frequency of using the hide-blurb button, however, seems to decrease over time.

This is consistent with our observation that there were different styles of interacting with the communication-bar. Most people used the tool in a passive manner: they looked at the information display when they were waiting for a lengthy computing task to finish, when they wanted to take a break, or when they were generally not very busy and their eyes came across an interesting blurb by chance. They described the communication-bar as similar to television in the sense that there was always information being displayed but they did not have to pay attention to it. A few people treated the “What’s Happening” communication-bar in a way similar to the morning newspaper: at the beginning of the day, they would click the forward button to go through all available blurbs and check if there was anything interesting to them. After reading the “newspaper”, their usage pattern usually became more passive, similar to other users’.

None of the users that we have interacted with permanently disabled the automatic cycling of the blurbs. Depending on many transient factors such as the level of involvement in other tasks, people's interest level of the current blurb, as well as their mood of the

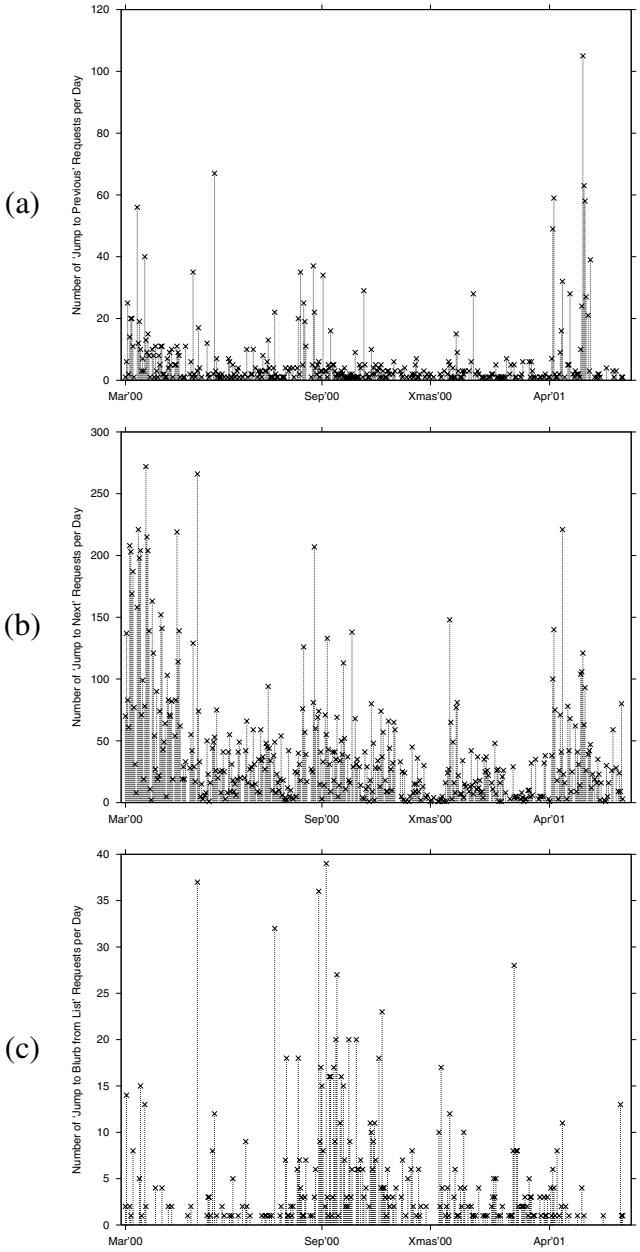

Figure 32. Usage of navigational features in the communication-bar graphed over time. Daily use of the jump-to-previous and jump-to-next buttons are shown in (a) and (b) respectively. The frequency of jumping to a specific blurb through the blurb list is shown in (c).

moment, they used the jump-to-next button when they wanted to see something else. They used the jump-to-previous button mostly if they happened to see something interesting and the display started to change before they could finish reading the blurb. Finally, they used the multiple levels of detail in the tool to gauge what was shown and used the web link feature to seek related information when they were interested. However, as people became accustomed to the application, they also became less inclined to remove and hide unwanted blurbs in the presentation cycle.

Two-thirds of the survey respondents used the communication-bar on Unix systems while the others used it on Microsoft Windows machines. Due to the instability of the Mac

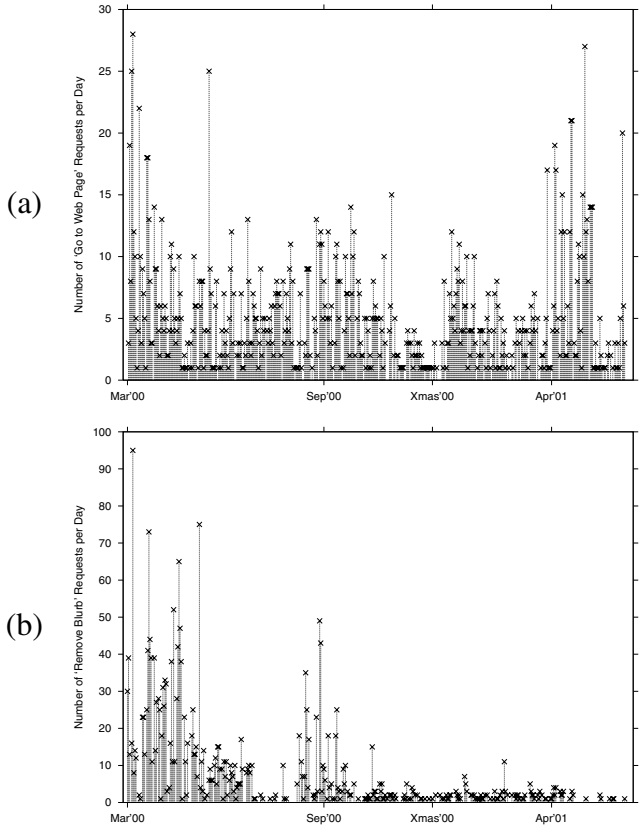

Figure 33. Usage of the content buttons on the communication-bar graphed over time. Daily use of the go-to-web-page and hide-blurb buttons are shown in (a) and (b) respectively.

Tcl/Tk library and hence the communication-bar program on this platform, people who tried the Mac version all stopped using it and were not included in the user survey. Many people used the horizontal layout. Only one person in addition to the developer used the vertical layout. Two-thirds of survey respondents put the tool at corners of their desktops while the others put it at edges. None used the center of the desktop to show "What's Happening". Some users even shrunk the application window to only display the content image and the first few words of the blurb title.

Many people kept the default selection of external content sources. Two survey respondents disabled local weather blurbs in the communication-bar with one of them commenting about having a window office. One survey respondent disabled stock quotes because of not owning any stock. A few selected only those web sites that they regularly read and used the tool as a filter — they used the web to pursue details when they felt interested by reading the summaries in "What's Happening". Adversely, one user disabled the content sources that he always read on the Web.

- **Subjective Rating of the User Interface**

Table 4 and Table 5 summarize people's ratings of the size of the application on their desktop displays and the obtrusiveness of the application (based on 13 responses to the user survey). Figure 34 and Figure 35 show the histograms of these ratings.

Table 4: Communication-bar user interface size rating. Ratings are on a 7-point scale where 1 represents "too small" and 7 represents "too big".

Question	Average Rating	Standard Deviation
On-the-screen size of the application	3.9	0.49

Table 5: Communication-bar user interface obtrusiveness ratings. Ratings are on a 7-point scale where 1 represents “not distracting” and 7 represents “very distracting”.

Question	Average Rating	Standard Deviation
Obtrusiveness of the presence of the application	2.2	1.14
Obtrusiveness of the transition animation	2.1	1.04

Although people thought that the on-screen-size of the application was appropriate, two users wanted the tool to automatically use more space when they read the current blurb. Overall, the ratings show that the application was relatively unobtrusive. However, one user commented that being curious about whether there was new information available in the communication-bar could be distracting at times.

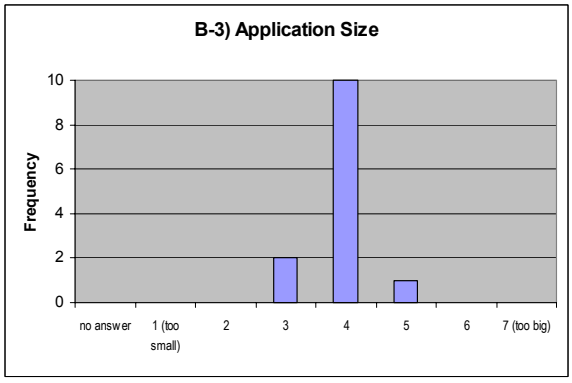

Figure 34. Histograms of communication-bar user interface size ratings.

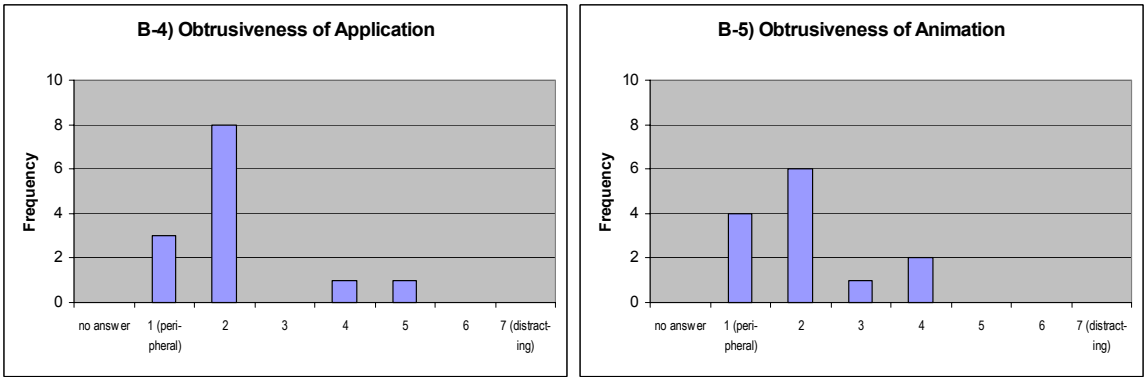

Figure 35. Histograms of communication-bar user interface obtrusiveness ratings.

In our interviews and informal interactions with the users, people expressed that the communication-bar was fairly unobtrusive as well. Once they became familiar with the tool, they tended to forget about its existence when they focused on other tasks. One user commented in an e-mail message that it was "the first community support tool that I've seen that is unobtrusive enough for me to actually keep it on my screen". Another user said that it was "a less intrusive way to keep up with information than, say, e-mail". Several others commented that the communication-bar reminded them about events in a better way at a better time. They did not feel pressured to check what was new or read what was displayed. When they had the leisure to read "What's Happening", they might read it more carefully and more completely, and perhaps were more likely to be interested to attend the events.

In general, people felt that the communication-bar was appropriate for showing information that did not require immediate response. In a few cases, the communication-bar even stimulated real-world interpersonal interactions. For example, people expressed birthday wishes when they bumped into a community member because they saw the corresponding information earlier in the communication-bar.

Table 6 summarizes the ratings in the final survey on the effectiveness of the communication-bar in keeping people updated on several domains of information. Figure 36 shows the corresponding histograms.

Table 6: Summary of communication-bar effectiveness ratings. Ratings are on a 7-point scale where 1 represents "not effective" and 7 represents "very effective".

Question	Average Rating	Standard Deviation
CoC events	4.5	1.33
CoC people's birthdays	5.5	1.85
Campus events	4.3	1.44
Local weather	5.2	1.17
Technology news	5.6	1.62
Financial news and stock quotes	5.4	1.38

The effectiveness of conveying each of these domains of information all received above average overall ratings. However, there were three main reasons that several survey respondents gave low effectiveness ratings. First, people wanted to be able to monitor certain types of information such as news about the Macintosh platform, that was not available in the communication-bar. In addition, a few users in the interviews expressed a desire to consolidate the many different types of reminder applications such as e-mail notifiers, calendars, and to-do lists, into one central place such as the "What's Happening" communication-bar.

Second, reading information with a high level of depth in a small window was difficult. In contrast to simple and short messages such as "town hall meeting at 5pm in MiRC 102", some blurbs such as one titled "3D without goggles" are inherently information-rich and require much longer text to tell the story. Since reading longer text in the communication-bar required more scrolling, several users suggested that they wanted to be able to quickly toggle between the regular small size display mode and a larger size mode to facilitate easier reading.

Finally, several people wanted urgent or time-sensitive information to be more noticeable than other types of information. For example, they were willing to allow reminders of events that they intended to attend to be shown slightly more distracting as the event was about to start. In addition, many people suggested that the application should show new information more frequently than old information so that they would have less chance of seeing something that they had already read.

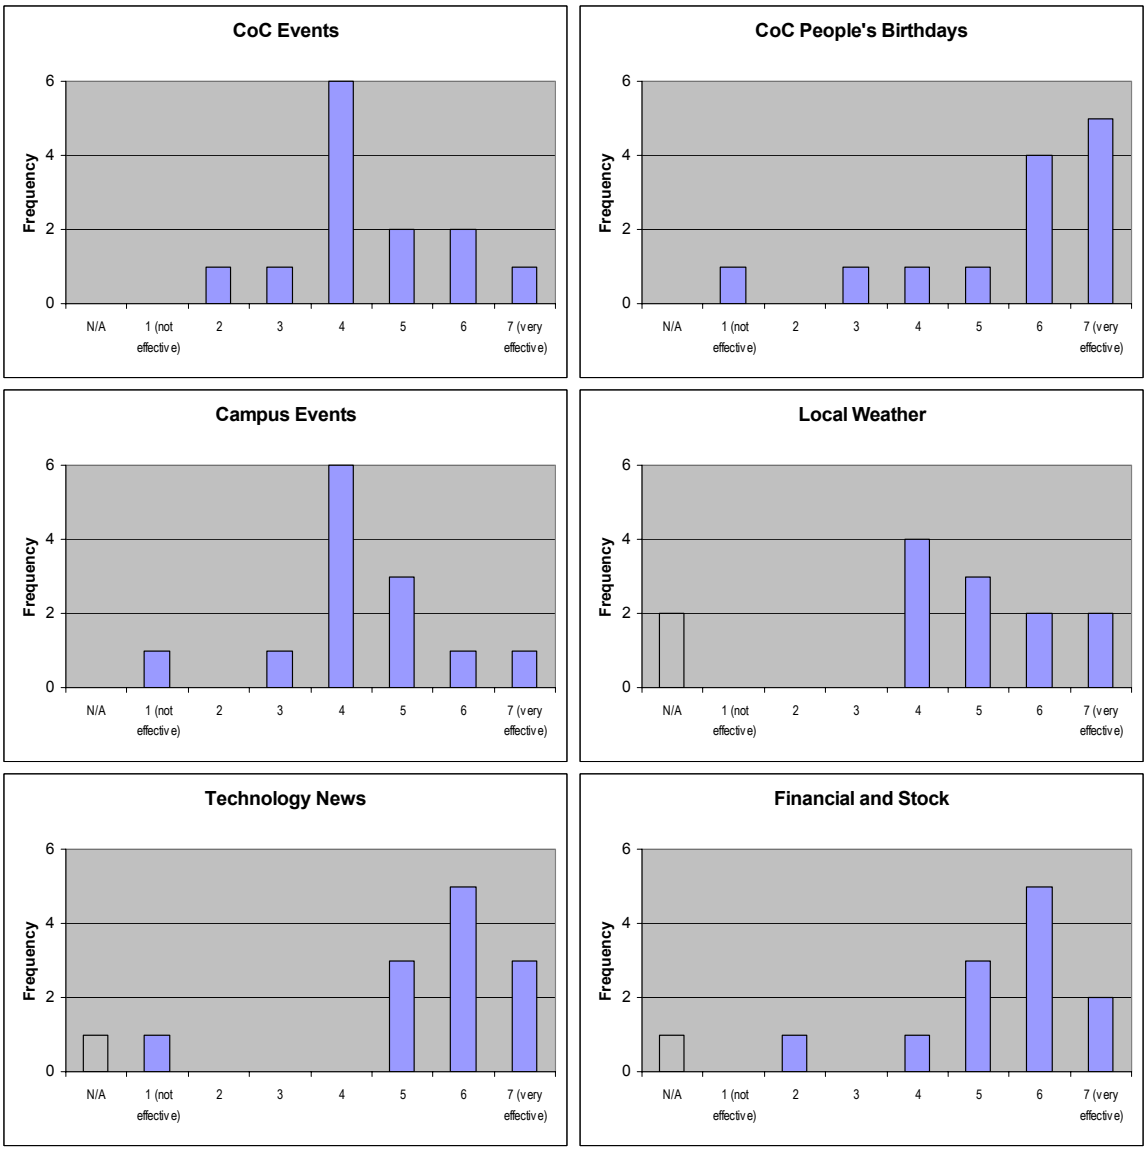

Figure 36. Histograms of communication-bar effectiveness ratings.

- **Participation**

Figure 37 shows the number of user posted blurbs and chat messages per day. Initially, we observed several types of user postings:

- Announcements such as congratulations for a student passing the qualifying exam, invitations to attend a demo event, and interesting or useful information found on the Web and passed along.
- Requests and offers such as those seeking tickets and selling used computer parts.
- Random opinions and calls for discussion such as guessing the outcomes of football games.

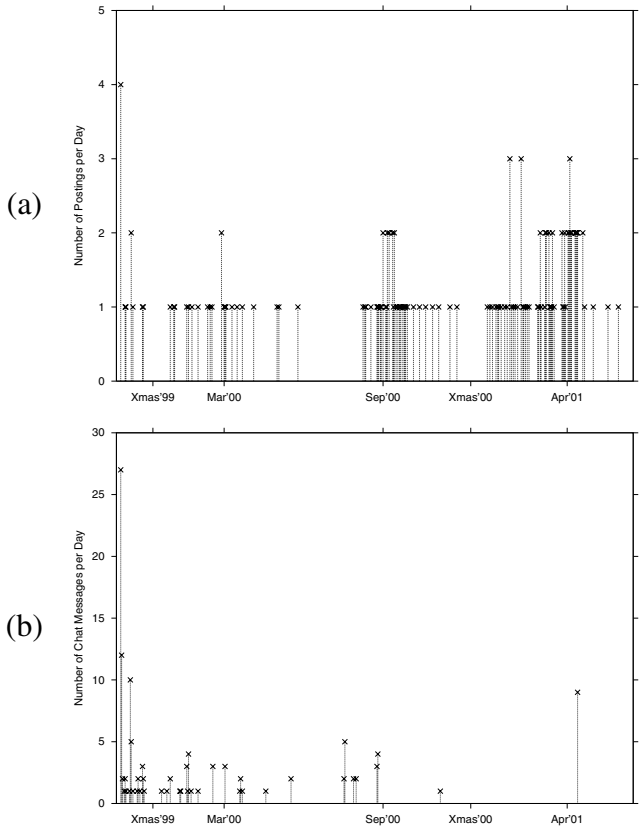

Figure 37. Daily blurb-posting and chat activities graphed over time.

Among the 16 people who posted blurbs, most only posted once, one person posted twice, one person posted three blurbs, and only the two people directly involved in this research each posted more than eight blurbs. Among the 45 people who posted chat messages, 37 of them each posted no more than three messages, six people each posted from five to eight messages, and only the two collaborators of this research each posted more than 15 messages.

Although a generic chat room was available through the WH interface, chat messages tended to relate to the content in a specific article, frequently involving comments or criticisms. For example, an automatically posted article about user identification in a software package spurred a debate on the controversies of privacy in ubiquitous computing.

Although several people had posted blurbs or engaged in chat-room discussion using the communication-bar at the beginning, participation declined to the level of about one posting per month, not counting blurbs forwarded by designated editors. People commented that even though everybody read approximately the same set of blurbs which provided similar frames of reference, they did not know who else was reading the one particular blurb that they were reading at the same time. In other words, certain amount of opportunity for interaction was lost due to the lack of presence awareness, as well as differences in choices of content and differences in scheduling of the blurbs in different instances of the client program.

People felt reluctant to post information because of several reasons: they did not know whether their posting would be interesting, they did not personally know the general audience of “What’s Happening”, they did not know what would be appropriate, and they thought other people might post the same information and make their own posting redun-

dant. In addition, a few users commented that they did not know whether their postings were followed up or not and did not get the feedback to encourage them to post other information.

Our evaluation of the communication-bar seems to confirm that this type of opportunistic interface is an appropriate tool to support awareness about less demanding information, in particular, information about the current affairs of the community and community members. However, it seems to also imply that the quality and the type of content affect usage and participation. People wanted to see more local content such as who was on leave or at a conference, recent awards or grants that other community members received, recent publications, academic and industrial visitors, etc. Since it is difficult to automatically decide what information flowing in the community is appropriate or relevant to post on "What's Happening", it is desirable that a small set of community members act as local content editors to filter information as well as collect information not already in electronic form. Other suggestions of increasing the utility and participation of the application included adding word games, showing titles of classes currently in session, and displaying random excerpts from people's personal web pages.

6.6 Evaluation of the Screen-saver

We set up four Windows NT machines in different labs to use the "What's Happening" screen-saver as the default logon screen-saver. The logon screen-saver is activated when the system is not in use by anybody, i.e. when no one is logged in on the console. This way, the screen-saver becomes a part of the physical environment, delivering information to whoever passes by. We also observed people leaving their desktops with the screen-saver running, making the tool available for others to see.

Table 7 and Table 8 list the average ratings of the 9 responses in the final survey on the effectiveness of the screen-saver in conveying different types of information and the interest levels on those types of information, respectively. Figure 38 and Figure 39 show the corresponding histograms.

Table 7: Summary of screen-saver presentation effectiveness ratings. Ratings are on a 7-point scale where 1 represents "not effective" and 7 represents "highly effective".

Question	Average Rating	Standard Deviation
CoC web images	5.1	1.83
"What's Happening" text messages	4.1	1.90
Local weather images	4.8	1.99
Highway traffic map and images	5.8	1.99
Local web-cam and TV images	4.3	2.12

Table 8: Summary of screen-saver personal interest ratings. Ratings are on a 7-point scale where 1 represents "not interested, ignored" and 7 represents "interested, always watch".

Question	Average Rating	Standard Deviation
CoC web images	6.1	1.05
"What's Happening" text messages	4.8	1.92
Local weather images	4.4	2.13
Highway traffic map and images	4.6	2.07
Local web-cam and TV images	2.3	1.50

The data seem to show strong individual differences in evaluating the effectiveness of the screen-saver. Some users liked the way that the images were placed. Others sometimes found the collages unpleasing due to the randomness of the layout algorithm. Some users were satisfied with the depth of information displayed. Others wanted a convenient way to

see details about the last collage such as the web page where the images were collected, but this was not implemented in the screen-saver. Finally, one user commented that the time that each collage stays on the screen should be made shorter, “perhaps 15-20 seconds but not a minute or two”.

There was much variability in people’s interest levels on the different collages as well. Some users liked to see other people’s vacation pictures while one survey respondent

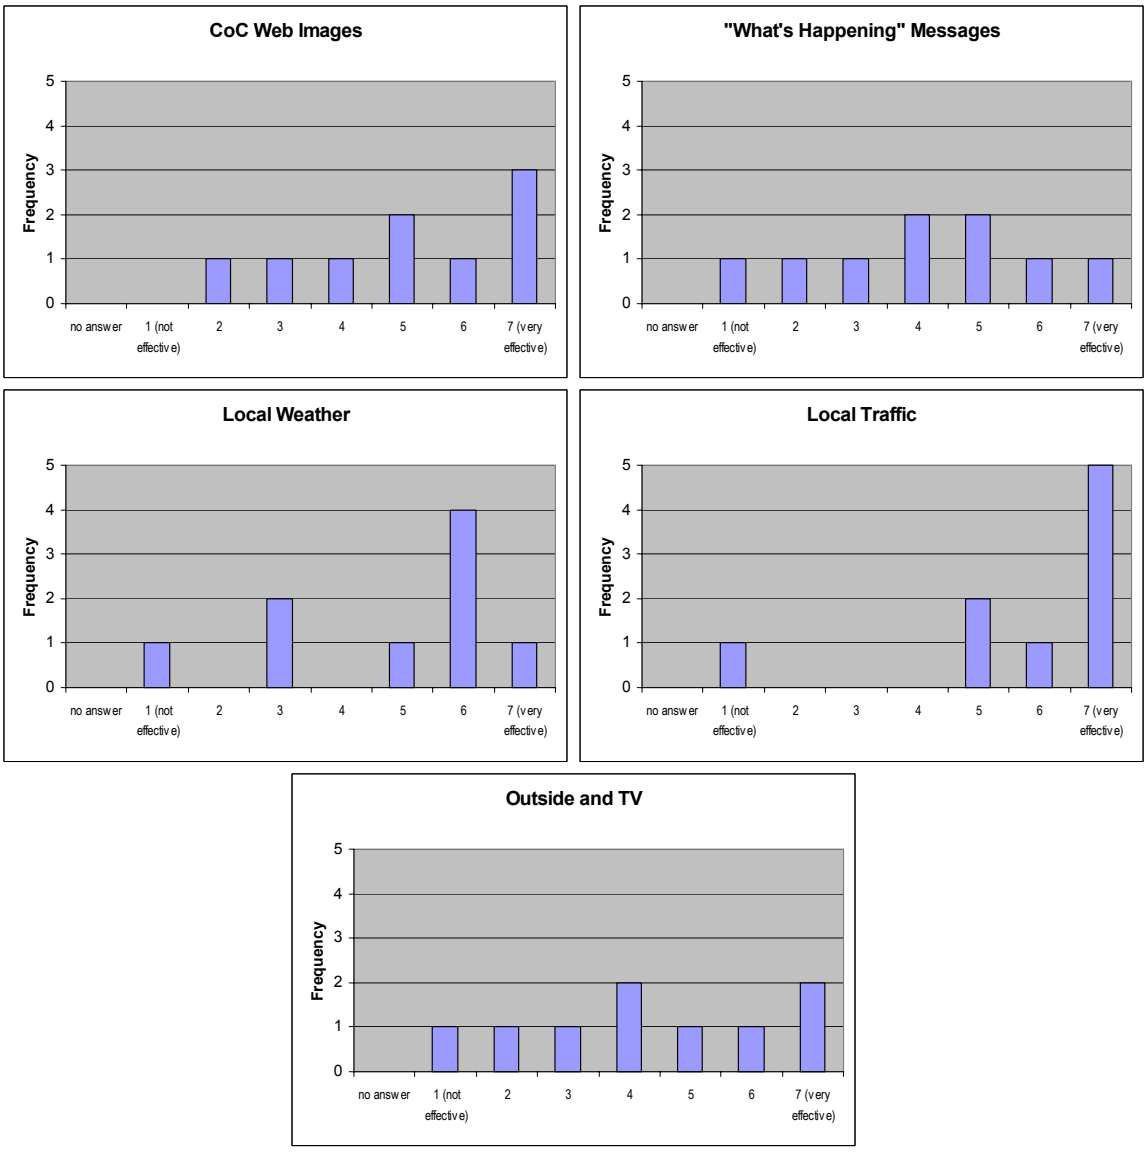

Figure 38. Histograms of screen-saver presentation effectiveness ratings.

wanted a way to filter out these images. Having a window office decreased at least one user's interest level on the weather collage. In addition, a few people were not interested in the traffic collages because they did not need to ride on the freeways to get home. Overall, people were interested in collages of local web images but not web-cam collages.

A few non-users dismissed the screen-saver because they would always be interacting with their desktop systems when they were in their offices. When the screen-saver was

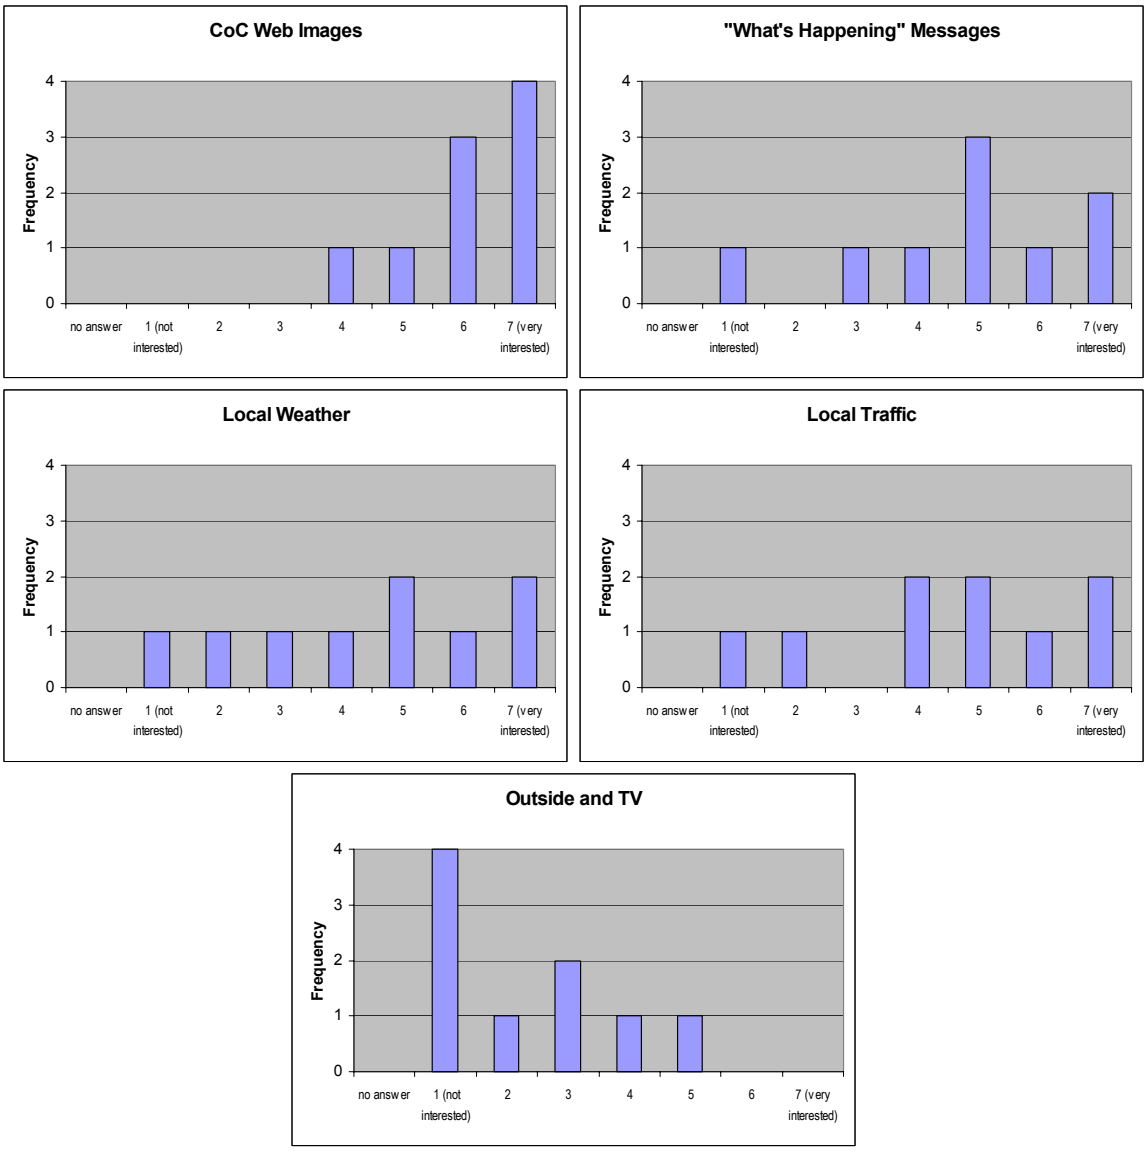

Figure 39. Histograms of screen-saver interest ratings.

activated, they were virtually never around. People who used the screen-saver had a slightly different work style — they often stopped to read, think, or deal with other activities next to the computers, which gave the screen-saver more chance to activate, and consequently, they saw the screen-saver display more often. One user commented:

> By the nature of a screen-saver, it's usually on when I'm not around or not paying attention to the screen. However, the collages seem to be the most useful for me... I like seeing the pictures that people put up on their web sites. Unfortunately, it seems that a lot of times you get image collages of navigation images rather than of actual interesting images. When large chunks of text appear on the screen, I find myself more inclined to skip it.

People sometimes received information that they initially did not know of. For example, one user was not aware that a certain member in the community got married until he saw the wedding pictures on the screen-saver. People also commented that the screen-saver had a certain entertainment value. Sometimes they talked with other people about what they saw on the screen-saver. Occasionally, external visitors would find a particular collage interesting or useful, and be intrigued by the screen-saver. In addition, several people said that the screen-saver provided them a convenient, but extremely opportunistic way of checking certain useful information, such as weather and traffic maps.

People did not, however, report that having the screen-saver activated was a distraction to the conversation with their guests. A few people commented that the screen-saver even gave them something to talk about and facilitated social information exchange. Even though talking about what was being shown on the screen-saver did not seem to directly help with their tasks at hand, people argued that they felt more content afterwards and they might indirectly benefit from the information in the future.

6.7 Reflections

While having 10 to 20 communication-bar users plus probably not many more screen-saver users in a several hundred person community is not a success story, finding out why people choose not to use the applications may help us better understand opportunistic interfaces and communityware. When shown pictures of the "What's Happening" applications and asked why they were not using the tools, 26 of the 55 non-users whom we surveyed, including the only three Mac users in this study, said that they were unaware of the communication-bar, and 21 of them indicated that they were unaware of the screen-saver (see Table 9 and Table 10).

Table 9: Reasons for not using the communication-bar. Items are sorted in descending order of frequency.

Reason	Frequency
Unaware of the communication-bar.	26
Haven't got a chance to try it.	8
Haven't got a chance to put it in start up script.	4
Don't use push applications; always pull for information.	3
Not sure if it's useful.	3
Not enough interesting information, not enough users.	3
No fixed computer to use; moving from one to other.	2
Would be tempted to look at it and be distracted.	1
Only use full-screen applications.	1
Program was unstable.	1
Sounds too complicated.	1
Could not make the application to work.	1
It slowed down the computer.	1

Table 10: Reasons for not using the screen-saver. Items are sorted in descending order of frequency.

Reason	Frequency
Unaware of the screen-saver.	21
Haven't got a chance to try it.	8
Not sure if it's useful.	8
Don't use any screen-saver.	6
Already see it on other people's computers.	3
No fixed computer to use; moving from one to other.	2
Sounds too complicated.	2
Don't use push applications; always pull for information.	1
Couldn't remember where the instructions were.	1
Haven't got a chance to put it in start up script.	1
Could not make the application to work.	1
It slowed down the computer.	1

One explanation of the lack of wide-spread adoption is perhaps that the perceived benefits of using the "What's Happening" applications did not overcome the associated costs. Despite our efforts in promoting these applications, many people did not seem to remember hearing about them or having seen them being used. Even among those who were aware of the applications, many did not feel that the software could be useful, at least not useful enough to try them or make the tools automatically start up when they logged into their computer systems. In addition, installing, running, and setting up the applications in login scripts were still non-trivial tasks despite of our efforts in making these tasks easy and writing comprehensive web pages to help people with these tasks. The flexibility of using many different, and incompatible, desktop environments in the College further complicated this problem. Lastly, however we made the communication-bar interface small

and lightweight, it still competed with other applications for screen space and color usage. Since very few people in the College had the luxury of large, high color-depth screens or second monitors to comfortably and persistently display the interface, the costs of running that particular application were still relatively high.

Although the administrative staff in the College were important sources of formal, official information as well as informal, unofficial information, they largely were unable to use the "What's Happening" applications because they use the Mac operating system exclusively while the communication-bar was unstable on MacOS and the screen-saver was unavailable on this platform. Therefore, it is important to improve the stability of the program and provide Mac users, including the administrative staff and a few faculty members and students, equal opportunity in accessing the "What's Happening" applications.

Finally, the ambiguity of accountability was another deficit of the "What's Happening" deployment. Non-users were sometimes reluctant to try the applications because they perceive them as short-term, transient research efforts. They occasionally asked the author where to ask for help if there were any technical problems and whether it would be maintained after the author had left. Users also wondered what would happen to the awareness applications after the author's forthcoming graduation. Therefore, it is advisable to create an entity responsible for the development and maintenance of community awareness applications and communicate this information to community members.

6.8 Summary

We have observed many different ways of using the "What's Happening" applications. Many people use them in a passive manner and they consider the tools as simple and unobtrusive. The two awareness tools provided users opportunities to preview, review, and

filter information that they might receive from other channels at the users' convenience. In addition, these applications provided users opportunities to discover not only information immediately useful to tasks at hand, but also information about the community and other community members that could benefit the users in the future.

While opportunistic interfaces are appropriate for delivering awareness information, they do not automatically provide or generate such information. In addition to user interface issues uncovered in the evaluation, content collection and creation need to be improved in future awareness applications, perhaps with the help from creative editors and leadership in the organization. Nevertheless, we believe that, coupled with useful and interesting content, opportunistic interfaces can enhance community awareness and encourage communication.

CHAPTER 7
Conclusions and Future Work

Community awareness is the general knowledge about community members and the community, such as people's backgrounds and interests, their on-going activities, and historical and current state of the community as well as gradual changes in the community. Even though community awareness rarely has immediate impact on people's daily work or study, enhancing community awareness improves communication, resource sharing, and collaboration in an academic research community. This thesis has examined technological supports for maintaining community awareness in a large and distributed real-world organization. The scope of this thesis spans from the early projects of supporting casual awareness across physical distances to the development of a basic understanding of opportunistic interfaces in maintaining community awareness.

The focus of this thesis was on the use of opportunistic interfaces that exist in the periphery on users' computer desktops for extended periods of time and convey useful information opportunistically when such chances arise. The development of two example opportunistic interfaces explored the benefits and limitations of various design choices as well as challenges that application designers may face. The evaluation of these two opportunistic interfaces investigated the usage of our particular designs and their effects on community awareness.

Initial interviews and questionnaires assessed the state of a local academic research community and identified several primary components of community awareness. People valued common knowledge of each other as well as the inclinations of engaging in casual conversations, attending community events, and offering help in community functions. People wanted better visibility of community members' achievements and more opportunities for spontaneous interpersonal interactions, as well as better management of information in ever growing amounts.

The early projects of the video wall and Electric Lounge experimented with providing virtual presence of people to create chance encounters and facilitate informal communication across physical distances. They exposed various limitations and difficulties in maintaining awareness in a large, amorphous group of people and furnished basic motivations for lightweight, unobtrusive alternatives in fostering community awareness.

The development of the "What's Happening" tools demonstrated two example designs of community awareness applications, emphasizing on providing useful and interesting content through simple and calm interfaces. The evolutions of these examples contrasted the implications of two design choices — a communication-bar that co-existed with other applications on the user's computer desktop versus a screen-saver that had exclusive control of the entire display area. In order to avoid distracting the user from other tasks and give the user a rough estimate of the information in a short glance, the communication-bar utilized a small on-screen size, a few simple controls, a layout that supported information presentation in progressive details, and a smooth, gradual animation that minimized sudden changes in appearance when switching from one instance of information to the next. The screen-saver, on the other hand, gave up the capabilities of customization and screen-

sharing in exchange for the advantage of being able to show large images to convey information that was otherwise difficult to describe. It focused on creating visually pleasing image collages and mixing image presentations from different sources to let the viewer discover interesting aspects of the community.

Customization is another one of the design dimensions in opportunistic interfaces. On one extreme, all sources of information need to be specified by the user. For example, on a My Netscape start-page, the user may choose any combination of information sources from a set of many [84]. On the other extreme, the information and presentation are universal among all users so that everybody receives the same. For example, information in a traditional newspaper is fixed across all subscribers and does not allow personalization. In the design of the "What's Happening" communication-bar, we chose to display community related information to all users in order to promote community awareness, while allowing people to choose from the set of external sources in order to provide flexible incentives for using this application.

The evaluation of the "What's Happening" tools cautioned designers about several obstacles in assessing community awareness and the effectiveness of community awareness technologies. While exposing several areas of improvement for future designs, the evaluation showed that opportunistic interfaces could convey information in the periphery and unobtrusively enrich one's information space without much effort. In addition, it showed that the effectiveness of community awareness applications depended on not only lightweight interfaces, but also the information content communicated through these interfaces. To foster community awareness, designers should couple awareness technologies

with practices in the community culture, such as contributing editors that help create interesting content, set examples for other users, and encourage participation.

Overall, this research has developed an understanding of the roles that opportunistic interfaces play in promoting community awareness. We may conclude that opportunistic interfaces help balance the tension between avoiding distraction and maintaining awareness, and can support community awareness in a large, distributed academic research organization. Even though this thesis only focused on opportunistic interfaces on computer desktops, the design elements can be translated to other platforms as well. For example, the communication-bar can be embedded in other applications, in door displays, or as part of the physical work space, while the screen-saver seem to be readily applicable to public information displays.

To increase the effectiveness of opportunistic interfaces in enhancing community awareness, developers should focus on three main areas of design: content generation, content presentation, and interaction between the user and the application. Since it is difficult to automatically collect, let alone create interesting content, it is important to supplement information that can be automatically collected with contributed stories from perhaps volunteers who hold central places in the community casual information exchange network. The presentation of the content should be unobtrusive to prevent the application from distracting the user, while providing high information utility to justify the costs involved in using the application. On average, the benefits of all the chances of accessing the information presented should outweigh the costs in order to sustain use over time. Therefore, the interaction between the user and the application should be simple and light-

weight, without requiring much effort from the user, to minimize the costs associated with each instance of use.

7.1 Improvements and Extensions

There are a number of natural improvements and extensions to the design and evaluation of the "What's Happening" applications that will expand our understanding of opportunistic interfaces for promoting community awareness.

Approaches to increase the benefits of using community awareness technologies should start with building better content. As previously described, volunteers may help gather and create interesting and entertaining content that identifies with the local community, that attracts newcomers and sustain experienced users. In addition, the community needs to build an interest in community awareness and establish a role of developing and maintaining community awareness technologies. Although our data showed only a few differences in the opinions toward the state of the overall community among people in different buildings, they most likely have different interests and expectations related to their work environments or job functions. Studying people's general interests and issues of concern may help designers discover the different types of content that appeal to different types of users.

Several users suggested integrating the "What's Happening" applications with other types of information that they wanted to monitor. Future extensions may allow people to write simple "plug-ins" to the communication-bar or screen-saver and mix information from user-defined sources with what is available now. Another extension may shift collage generation to the screen-saver client side and use the same algorithm to monitor external web sites listed in the user's bookmark file. Finally, techniques in attentive information

systems and user modeling research may help dynamically infer and monitor users' specific interests, tailoring the information to better fit people's work styles and increase the benefits to individual users [71, 99].

Making the programs easier to install and optimizing their usage of computing resources may help reduce the costs in adopting these new technologies. Putting the applications in default user profiles and new system installations so that these programs automatically appear when a person logs into a new account or a new type of computer may help promote usage and reduce complexity[1].

In our evaluation of the "What's Happening" applications, one person complained that having to constantly move around without staying at any particular one desktop computer system hindered the use of those awareness applications. The utility of the opportunistic interfaces depends on the frequency of chances that the user "passes by" the application. This seems to suggest that a public display is still needed to allow opportunistic access to community related information. Utilizing public displays in shared physical spaces may reduce the impact on individual user's attention and the need for running the awareness tools individually.

When developing the "What's Happening" applications, we had to make design decisions based on the availability of underlying technologies and balance the simplicity of the designs with the power and flexibility of the applications. As such, several new ideas in system architecture may fundamentally improve our existing designs. For example, the current communication-bar uses heavyweight TCP connections to the content server to

1. It should be easy to deactivate these programs if the user wants to do so, in order to avoid making the user feel "forced" to use the applications.

receive blurbs. An alternative design based on multicast UDP may conserve network resources and provide better scalability [6].

Even though the appearance of the communication-bar was considered unobtrusive and the size was appropriate, a few users had other applications partially covering their instances of the communication-bar, obscuring the menu and navigational buttons. Several other users wanted to hide those buttons when not using the application. This seems to suggest that one area of improvement is to make the interface use less screen real-estate and make the buttons less distracting. One possible change is to move the menu, post, and flip backward and forward buttons to a transparent layer on top of the image display area and arrange them along one side or at the corners. When the mouse pointer is outside the application window, the buttons are highly transparent, giving faint representations of themselves. When the mouse pointer moves into the application window, the buttons fade in to a semi-transparent state, allowing the user to see the actions that they are representing. Previous experiments suggested that overlaying semi-transparent user interface objects with displayed information could be an effective way of utilizing small screen spaces [60].

Although at many times the collage server generated intriguing and visually pleasing image collages, at other times it generated either sparse-looking ones with only a few small images or crowded ones with one set of images nearly completely obscuring another set. In addition, the current image selection method focuses on web pages with embedded or linked images. Consequently, people who have many such web pages will have a higher chance of getting an image collage generated based on the images that they publish. It might be beneficial to use a different method that gives each person equal chance of being

selected. An alternative layout algorithm may be extended from [10] that limits the amount of overlapping and detects when the collage canvas is almost full. Other possibilities include the spiral layout in Mandala [49] and artistic compositions using aesthetic templates in Kandinsky [33]. Empirical studies need to be conducted to compare the results of the different algorithms in visual appearance and information comprehension.

The evaluation of the "What's Happening" tools required the recording of users' interactions with the awareness applications. To reduce the effort involved in running the applications, we did not use any authentication mechanism such as asking for the login name and password to identify a user. Instead, we compromised on a less accurate, host-address-only journaling facility. Alternatively, we may ask the user to log in when the applications are used for the first time. A software identification number, also known as a "cookie", is then generated for the user and stored securely in the user's private file space. This cookie will help identify the user in the future and allow more precise evaluation of the use of the applications.

7.2 Future Directions

This thesis also opens new research opportunities in understanding the task of maintaining awareness and the benefits of different technologies. For example, awareness applications fit naturally on a second monitor on the computer desk [43, 80]. What assumptions can be made in this situation? What design changes should be made? Are they more effective on the second monitor? These are some of the research questions that need to be addressed.

Integrating community awareness technologies into applications that people already use as well as existing or new social practices may help people better understand the bene-

fits and further lower the costs in adopting these technologies [66, 89]. How should designers integrate opportunistic community awareness techniques with other awareness and communication technologies? For example, the communication-bar may support certain rudimentary awareness of other people using the tool, such as how many they are and how many are seeing the same blurb. We may also replace the chat-room images with the Babble social proxy to support social transparency and basic awareness on people's activities [29, 123]. What level of awareness about other people is worthwhile to support in opportunistic interfaces? How should this awareness be represented? In addition, other integration candidates may include the Squeaker's Collaborative Radio project where "What's Happening" content can be inserted sparingly in between music playings [108] and the Xerox PARC Cover Up project where information can be added to print job burst pages [65]. What effects on community awareness can we expect from these technologies? What existing social practices are appropriate to enhance through community awareness technologies and what new ones can we create?

How can community related information be presented opportunistically yet still relating to the context of people's work activities? For example, when a person is writing a research paper on software agents, an opportunistic interface may show a reminder about a seminar on a related topic and this person may take advantage of this opportunity to learn new information. What role does context play in opportunistic interfaces and community awareness [24]?

A fun experience in using opportunistic interfaces is more likely to attract users than one that feels like work. How can we design entertaining practices that supplements com-

munity awareness? What lightweight interfaces afford such practices? And eventually, how do we evaluate the play aspect in community awareness?

Finally, our local College of Computing community is a unique environment in many ways. Although it is geographically distributed, regularly meeting people in different buildings face-to-face is still possible. This condition may have influenced people's interest in awareness technologies and their tolerance level on related costs. In addition, academic research communities often have knowledge-oriented emphases that may support the development of communityware applications, while other communities may have product- or service-oriented emphases that perhaps consider community awareness at different priorities. It is important to study opportunistic interfaces and communityware applications in the contexts of other academic research communities and non academic research ones, as well as smaller and larger communities.

7.3 Summary of Contributions

This research contributes to the field of human-computer interaction by exploring opportunistic interfaces for promoting community awareness. Community awareness is a relatively unexplored area and is becoming increasingly more important in large, distributed academic research communities. Previous research has primarily focused on supporting the distinct awareness that directly relates to people's tasks at hand or has concrete ramifications on people's daily work or study, rather than the background awareness about the general attributes of people and the community as well as their on-going activities. This thesis thus provides a framework for investigating opportunistic interfaces as an alternative design philosophy and their use in maintaining community awareness.

In summary, the research work described in this thesis made the following contributions:

- **Identified characteristics of opportunistic interfaces and requirements for maintaining community awareness.** Existing groupware technologies mostly emphasize on awareness of immediate surroundings and actions in close-knit collaborative situations. Community awareness, on the other hand, lacks details of activities and often include knowledge about weakly connected community members. Opportunistic interfaces deliver information by chance, in the periphery, and may be more appropriate for maintaining community awareness with less effort and less distraction.
- **Iteratively designed example opportunistic interfaces for promoting community awareness in a real-world academic research community.** The communication-bar calmly cycles through blurbs of information in a small window on the user's computer desktop. The screen-saver shows collages of images collected on local community web servers in full-screen. The evolutions of these applications demonstrated various trade-offs and design choices in balancing awareness and distraction.
- **Evaluated the use and effectiveness of example opportunistic interfaces in promoting community awareness.** Carefully designed, unobtrusive and lightweight opportunistic interfaces can help people enhance their awareness on simple topics, maintain their awareness on complex ones, and discover new information about the community.

- **Analyzed the drawbacks of example community awareness applications and developed recommendations for future designs**. In addition to improve the benefits of using the applications and reduce the costs, we would still need leadership to help people understand the benefits of a strong sense of community, to set examples, tones and conventions, and to encourage participation. The role of the opportunistic interfaces will then be to make these guiding activities easy, to facilitate the direction that these examples set forth, and to amply the impacts these examples make.

These contributions extend our understanding of community awareness, the relationship between technology and community awareness, and the utility of opportunistic interfaces for promoting community awareness.

APPENDIX A
Initial Community Questionnaire

We deployed the following questionnaire in the local community before introducing the "What's Happening" tools. See Section 2.1.3 for details.

```
Some members of our community have expressed concerns about the sense of
community in CoC during this period of growth. To help us better understand
this issue, I propose that we start by asking "how much do you know about
others in the College?"

Please fill out this short survey and email it back to me (or put a
hardcopy in my mailbox if you prefer). Information about you, if it's
available to me, will be kept confidential and separate from the survey
results.

Please answer each of the multiple-choice questions by placing an "X" mark
within the pair of [ ] that corresponds best to your opinion on a 7-point
scale.

1) How would you rate your familiarity with the CoC research-oriented
events that are being scheduled (e.g. faculty recruiting talks, CoC
distinguished lecture series lectures, etc.)?

       1        2        3        4        5        6        7
      [ ]      [ ]      [ ]      [ ]      [ ]      [ ]      [ ]
      very                                                  very
      unfamiliar                                            familiar

2) How often do you attend CoC research-oriented events?

       1        2        3        4        5        6        7
      [ ]      [ ]      [ ]      [ ]      [ ]      [ ]      [ ]
      rarely                                                attend
      attend                                                a lot

3) How would you rate your familiarity with the CoC social events that are
being scheduled?

       1        2        3        4        5        6        7
      [ ]      [ ]      [ ]      [ ]      [ ]      [ ]      [ ]
      very                                                  very
      unfamiliar                                            familiar

4) How often do you attend CoC social events?

       1        2        3        4        5        6        7
      [ ]      [ ]      [ ]      [ ]      [ ]      [ ]      [ ]
      rarely                                                attend
      attend                                                a lot
```

5) How would you rate your familiarity with research work in groups other than your own? Here group refers to the smallest recognized people entity (for example, the animation group fits our definition of "group" better than, say, the GVU Center).

1	2	3	4	5	6	7
[ ]	[ ]	[ ]	[ ]	[ ]	[ ]	[ ]
very unfamiliar						very familiar

6) How many CoC people outside of your group do you interact with in research/academic settings?

1	2	3	4	5	6	7
[ ]	[ ]	[ ]	[ ]	[ ]	[ ]	[ ]
very few						very many

7) How many CoC people outside of your group do you interact with in social settings?

1	2	3	4	5	6	7
[ ]	[ ]	[ ]	[ ]	[ ]	[ ]	[ ]
very few						very many

8) How would you rate the sense of community in CoC?

1	2	3	4	5	6	7
[ ]	[ ]	[ ]	[ ]	[ ]	[ ]	[ ]
very bad						very good

9) In your opinion, what are the most important contributing factors to the sense of community in CoC?

10) If you were to change a few things about the CoC community, what would those things be?

Thank you very much for your cooperation!

= Q. Alex Zhao ~{0"UT~}
http://www.cc.gatech.edu/~qiang.a.zhao/
mailto:azhao@cc.gatech.edu voiceto:404-894-9390 faxto:404-385-1253
Graphics, Visualization & Usability Center, Georgia Inst. of Tech.

APPENDIX B
Client-Server Communication Protocol

This section describes the "What's Happening" client-server communication protocol. A plain-text copy of this document is available on College of Computing Unix systems as `/usr/local/happening/API.txt` to allow other developers to communicate with the "What's Happening" server in their applications.

The table below lists several parameters of the protocol.

Table 11: Connection parameters for client-server communication.

Protocol version	20000316
Server host name	`happening.cc.gatech.edu`
Server TCP port numbers	53841 for text commands 53842 for image transmissions

In the following discussion, text enclosed in angle brackets (<>) represents parameters which will be substituted by real data in client-server communications.

- **Connecting to the Server**

  The client opens a TCP connection to the server on the text command port and sends

  ```
  refreshAll <protocol version #>
  ```

  with the protocol version number replaced by that the client can understand. Then server responds with an acknowledgment with the same keyword:

  ```
  refreshAll <protocol version #>
  ```

except that the protocol version is the one which the server understands. If the client protocol version does not match the server protocol version, it should disconnect and either automatically upgrade itself to the latest version or prompt the user to upgrade.

If the protocol versions match, the client should remove all the old articles, if any, from the previous successful connection. The server will send update messages for all current content afterwards. If the connection breaks up for some reason, the client should use exponential backup to try to reconnect.

- **Updating Content**

The server broadcasts the following message for adding new content or updating an existing one:

```
updateSubItem <article category> <article key> <data>
```

where the category may be "calendar", "slashdot", etc.; the key is a unique identifier within the category; data is the real content of the blurb in the form of a Tcl list — the first element is the title, the second is the image file name within that category, and the rest is another list which is the body of the news. Elements of the body list should be displayed as separate lines. The article key is generally in the form of "*preamble*&<URL>" -- the embedded URL tells what web page to go to for more details on the article.

- **Posting a New Blurb**

To post a new blurb, the client sends the following to the server:

```
startNewsItem <key> <life> <user> <title> <image> <body>
```

Here the implicit category is "User". The life parameter specifies how soon, in seconds, the blurb should expire and be removed. The user id is generally the user's login

name, the MacOS Chooser name, or in the form of "unknown@<hostname>". The client is responsible for generating the unique key in the form of:

```
<user>%<process-id>/<sequence-number>#<random>&<URL>
```

- **Acknowledging a Posted Blurb**

  When the user posts a new blurb or creates a new chat-room, the client should wait for the following message from the server before showing the result of the posting to the user:

  ```
  gotoItem <category> <key>
  ```

- **User-Initiated Blurb Cancellation**

  A user may cancel their own blurbs. In this case, the client sends the following message to the server:

  ```
  cancelNewsItem <key>
  ```

- **Server-Initiated Blurb Removal**

  When a blurb expires, the server sends the following message to delete it:

  ```
  removeSubItem <category> <key>
  ```

- **Creating a New Chat-room**

  A user may create a new chat-room to comment on an existing blurb.

  ```
  startChatItem <original category> <key> <user> <opening>
  ```

- **Adding to a Chat-room Discussion**

  When a user types something into the chat-room text entry, the client should send

  ```
  append2ChatItem <key> <user> <line>
  ```

  to the server. The server then broadcasts the following message to all clients:

  ```
  append2ChatItem <key> <user: line>
  ```

Here the implicit category is “Chat”.

- **Uploading and Downloading Images**

The client should use the image port to upload or download images. When uploading an image, the first line that the client sends to the server is the name of the file. For security reasons, the server will remove all directory components from that name and prefix “`User/`” to the result. Finally, both parties should switch the connection to *binary* mode and the server saves everything it receives thereafter to the file.

Downloading is similar to uploading, except that the first character on the first line to the server is an “`@`” sign. The rest of the line is a Tcl list of the form:

```
{<category> <filename>}
```

Again, the client sends the image data in binary mode.

APPENDIX C
Communication-bar Interview Script

The following is a script that we roughly followed in conducting user interviews to evaluate the "What's Happening" communication-bar.

```
0) General info about the user

  Name  _____________________    Date    __________

  [ ] PhD               area: _____________________
  [ ] Master's          area: _____________________
  [ ] Undergrad         area: _____________________
  [ ] Faculty           area: _____________________
  [ ] CNS
  [ ] Administrative staff

  Years with CoC:       _____           Current Building:       _____

  Primary desktop system(s):
        [ ] Unix        [ ] Windows     [ ] MacOS

1) Mechanics

  - Where on the desktop do you put the application?
    What's its orientation (horizontal vs. vertical)?

  - Does the application come up automatically when you login to a computer?
    - [ ] If yes, do you
          [ ] logout when you leave for the day or
          [ ] do you keep your computer logged in?
    - [ ] If no, how often do you start up the application?
          [ ] twice a day or more
          [ ] once a day
          [ ] a few times a week
          [ ] once a week or less

  - How would you describe the user interface?
    - How does it work and what do those buttons do?
    - Do you use the skip backward/forward buttons? How often?
    - Do you use the "go to web page" button? How often?
    - Do you use the "trash this" button? How often?
    - Do you use the cycling toggle? How often?
    - Do you use the list news/chat menu options? How often?
    - Do you change preferences? How often?

  - Do you change the preferences? (what are the current prefs)

  - On a scale of 5 (1 being "too small" and 5 being "too big"), how would
    you rate the on-the-screen size of the application?
```

size: 1 2 3 4 5

- On a scale of 5 (1 being unnecessary and 5 being very important), how would you rate the importance of each of the following features in an article:
  image: 1 2 3 4 5
  headline: 1 2 3 4 5
  summary: 1 2 3 4 5

- How would you describe the animation?
  On a scale of 5 (1 being "not at all" and 5 being "too much"), how would you rate the distractiveness of the animation:
  distractive: 1 2 3 4 5

2) Awareness

- How would you describe what this application is?
  - Why do you use it?
    If you stopped using it, why?
  - What prompted you to put it in your startup script?
    Or, why didn't you put it in your startup script?

- Do you find yourself:
  [ ] purposefully stopping and looking at the application, or
  [ ] do you look at it purely by chance, or
  [ ] some change in the display catches your attention?
  How would you characterize the frequency of that happening?

- On a scale of 5 (1 being "not interested at all" and 5 being "highly interested"), how would you rate your interest level in the following areas (independent of the app):
  CoC events: 1 2 3 4 5
  CoC b-days: 1 2 3 4 5
  campus events: 1 2 3 4 5
  weather: 1 2 3 4 5
  tech-news: 1 2 3 4 5
  financial: 1 2 3 4 5
  sports news: 1 2 3 4 5

  Among the ones you are interested in, does the application keep you informed on what's going on? (too much info, too little, ...)

- Is there other type of content you'd like to see but doesn't exist in this application?

- What has the application made you more aware of?
  How so (could you give some examples)?
  Have you noticed anything related to the CoC community?

3) Participation and communication

- Do you post articles?
  - If not posting, why not?
    Could you think of any situation when you do want to post?
  - Would you want to post more often or less often?
    Why?

- Do you use the chat feature?
  - Do you
    [ ] create new chat rooms, or
    [ ] follow up other people's messages, or both?
    How often?
  - If not using it, why not?
    Could you think of any situation when you do want to use it?
  - Would you want to use the chat feature more often or less often?

Why?

4) General comments

- How could the application be improved?
  What needs to be changed/added/deleted/etc.?

- What does the application do best? What's its biggest strength?

APPENDIX D
Final Survey on "What's Happening"

The following is the final survey to the "What's Happening" communication-bar and screen-saver users. It was deployed in April, 2001, and was delivered in person with a few exceptions where an e-mail was sent.

Please fill out this short survey and email it back to me (or put a hardcopy in my mailbox if you prefer). Information about you, if it's available to me, will be kept confidential and separate from the survey results.

Please answer each of the multiple-choice questions by placing an "X" mark within the pair of [ ] that corresponds best to your opinion on a 7-point scale.

A) The following four questions ask about your general impressions of the College of Computing community.

1) How would you rate your familiarity with the CoC research-oriented events that are being scheduled (e.g. faculty recruiting talks, CoC distinguished lecture series lectures, etc.)?

[ ] 1 [ ] 2 [ ] 3 [ ] 4 [ ] 5 [ ] 6 [ ] 7
very very
unfamiliar familiar

2) How would you rate your familiarity with the CoC social events that are being scheduled?

[ ] 1 [ ] 2 [ ] 3 [ ] 4 [ ] 5 [ ] 6 [ ] 7
very very
unfamiliar familiar

3) How would you rate your familiarity with research work in groups other than your own? Here group refers to the smallest recognized people entity (for example, the animation group fits our definition of "group" better than, say, the GVU Center).

[ ] 1 [ ] 2 [ ] 3 [ ] 4 [ ] 5 [ ] 6 [ ] 7
very very
unfamiliar familiar

4) How would you rate the sense of community in CoC?

[ ] 1 [ ] 2 [ ] 3 [ ] 4 [ ] 5 [ ] 6 [ ] 7
very very
bad good

B) This section contains questions specifically about the "What's Happening" info-bar. If you do not use that application, please skip to section C.

1) What desktop system(s) do you use the info-bar on?

[ ] Unix [ ] Windows [ ] MacOS
[ ] other people's desktops

2) Where on the desktop do you put the application?

[ ] northwest [ ] north [ ] northeast
[ ] west [ ] center [ ] east
[ ] southwest [ ] south [ ] southeast

3) How would you rate the on-the-screen size of the application?

[ ] 1 [ ] 2 [ ] 3 [ ] 4 [ ] 5 [ ] 6 [ ] 7
too small too big

4) How would you characterize the info-bar's presence on your computer desktop?

[ ] 1 [ ] 2 [ ] 3 [ ] 4 [ ] 5 [ ] 6 [ ] 7
peripheral attention-grabbing
non-distracting distracting

5) How would you characterize the animation that transitions between articles?

[ ] 1 [ ] 2 [ ] 3 [ ] 4 [ ] 5 [ ] 6 [ ] 7
pleasing jarring
non-distracting distracting

6) How would you rate the effectiveness of the info-bar in keeping you updated about the following topics:

CoC events:
[ ] 1 [ ] 2 [ ] 3 [ ] 4 [ ] 5 [ ] 6 [ ] 7
not very
effective effective

"Happy Birthdays" to CoC people:
[ ] 1 [ ] 2 [ ] 3 [ ] 4 [ ] 5 [ ] 6 [ ] 7
not very
effective effective

campus events:
[ ] 1 [ ] 2 [ ] 3 [ ] 4 [ ] 5 [ ] 6 [ ] 7 [ ] N/A
not very disabled
effective effective

Atlanta weather:
[ ] 1 [ ] 2 [ ] 3 [ ] 4 [ ] 5 [ ] 6 [ ] 7 [ ] N/A
not very disabled
effective effective

technology news:
[ ] 1 [ ] 2 [ ] 3 [ ] 4 [ ] 5 [ ] 6 [ ] 7 [ ] N/A
not very disabled
effective effective

financial news and stock quotes:

[ ] 1 [ ] 2 [ ] 3 [ ] 4 [ ] 5 [ ] 6 [ ] 7 [ ] N/A
not effective very effective disabled

7) Do you have any general comments about the info-bar? Is the interface easy to use? Why do you use this application? Or, if you stopped using it, why?

C) This section contains questions about the "What's Happening" screen-saver. If this is not applicable to you, please skip to section D.

1) What desktop system(s) do you use the screen-saver on?

[ ] Unix [ ] Windows [ ] MacOS
[ ] other people's desktops

2) How effective are the following types of collages in communicating information about their respective topics?

CoC web images:
[ ] 1 [ ] 2 [ ] 3 [ ] 4 [ ] 5 [ ] 6 [ ] 7
not at all highly effective

The "What's Happening" text messages:
[ ] 1 [ ] 2 [ ] 3 [ ] 4 [ ] 5 [ ] 6 [ ] 7
not at all highly effective

Weather:
[ ] 1 [ ] 2 [ ] 3 [ ] 4 [ ] 5 [ ] 6 [ ] 7
not at all highly effective

Traffic:
[ ] 1 [ ] 2 [ ] 3 [ ] 4 [ ] 5 [ ] 6 [ ] 7
not at all highly effective

Atlanta Web-cams:
[ ] 1 [ ] 2 [ ] 3 [ ] 4 [ ] 5 [ ] 6 [ ] 7
not at all highly effective

3) How much interest do you have personally in the following types of collages:

CoC web images:
[ ] 1 [ ] 2 [ ] 3 [ ] 4 [ ] 5 [ ] 6 [ ] 7
not interested ignored interested always watch

The "What's Happening" text messages:
[ ] 1 [ ] 2 [ ] 3 [ ] 4 [ ] 5 [ ] 6 [ ] 7
not interested ignored interested always watch

Weather:
[ ] 1 [ ] 2 [ ] 3 [ ] 4 [ ] 5 [ ] 6 [ ] 7
not interested ignored interested always watch

Traffic:
[ ] 1 [ ] 2 [ ] 3 [ ] 4 [ ] 5 [ ] 6 [ ] 7

not interested interested
ignored always watch

Atlanta Web-cams:
[ ] 1 [ ] 2 [ ] 3 [ ] 4 [ ] 5 [ ] 6 [ ] 7
not interested interested
ignored always watch

4) Do you have any general comment about the screen-saver? Why do you use it? Or, if you stopped using it, why?

D) That's it -- No more questions. Thank you very much for your help!

Bibliography

[1] AccessAtlanta. At `http://www.accessatlanta.com/`.

[2] Ackerman, M. S. and Palen, L. The Zephyr Help Instance: Promoting Ongoing Activity in a CSCW System. In *CHI'96 Conference Proceedings*, 268-275. ACM, 1996.

[3] Ackerman, M. S. and Starr, B. Social Activity Indicators: Interface Components for CSCW Systems. In *UIST'95 Conference Proceedings*, 159-168. ACM, 1995.

[4] Agre, P. E. and Rotenberg, M. *Technology and Privacy: The New Landscape*. MIT Press, 1998.

[5] America Online. AOL Instant Messenger. See `http://www.aol.com/aim/`.

[6] Ammar, M. H., Almeroth, K., Clark, R. J., and Fei, Z., Multicast Delivery of Web Pages OR How to make Web servers pushy. In *Proceedings of the Workshop on Internet Server Performance*, Madison, Wisconsin, June 1998.

[7] Arnold, K. and Gosling, J. *The Java Programming Language*. Addison-Wesley, 1996.

[8] Axelrod, R. *The Evolution of Cooperation*. Basic Books, 1984.

[9] Ball, S. TclXML. Available at `http://www.zveno.com/zm.cgi/in-tclxml/`.

[10] Bell, B. A. and Feiner, S. K. Dynamic Space Management for User Interfaces. In *UIST 2000 Conference Proceeding*, 239-248. ACM, 2000.

[11] Bellotti, V. What You Don't Know Can Hurt You: Privacy in Collaborative Computing. In *HCI'96 Conference Proceeding*, British Computer Society. Springer-Verlag, 1996.

[12] Bly, S. A., Harrison, S. R., and Irwin, S. Media Spaces: Bringing People Together in a Video, Audio, and Computing Environment. *Communications of the ACM*, vol. 36, no. 1, 28-47. ACM, January 1993.

[13] Boyle, M., Edwards, C. and Greenberg, S. The Effects of Filtered Video on Awareness and Privacy. In *CSCW 2000 Conference Proceeding*. ACM, 2000.

[14] Brinck, T. and Gomez, L. M. A Collaborative Medium for the Support of Conversational Props. In *CSCW'92 Conference Proceeding*, 171-178. ACM, 1992.

[15] Bruckman, A. and Resnick, M. The MediaMOO Project: Constructionism and Professional Community. In *Convergence*, 1:1. Spring 1995.

[16] Buxton, W. A. S. Living in Augmented Reality: Ubiquitous Media and Reactive Environments. In Finn, K., Sellen A., and Wilber, S. (Eds.), *Video Mediated Communication*, 363-384. Erlbaum, 1997.

[17] Cadiz, J., Fussell, S., Kraut, R., Lerch, F. J., Scherlis, W. The Awareness Monitor: A Coordination Tool for Asynchronous, Distributed Work Teams. Unpublished manuscript. Demonstrated at CSCW'98.

[18] Chen, D. and Gellersen, H. W. Recognition and Reasoning in an Awareness Support System for Generation of Storyboard-like Views of Recent Activity. In *GROUP'99 Conference Proceeding*, 356-364. ACM, 1999.

[19] Churchill, E. F., Trevor, J., Bly, S., Nelson, L., and Cubranic, D. Anchored Conversations: Chatting in the Context of a Document. In *CHI 2000 Conference Proceeding*, 454 - 461. ACM, 2000.

[20] Clean Air Campaign at `http://www.cleanaircampaign.com/`.

[21] Coleman, J. S. Social Capital in the Creation of Human Capital. American Journal of Sociology, 94:S95-S120. University of Chicago Press, 1988.

[22] Cutrell, E., Czerwinski, M., and Horvitz, E. Notification, Disruption and Memory: Effects of Messaging Interruptions on Memory and Performance. To appear in *Proceedings of Interact 2001*, Tokyo. 2001.

[23] Dahley, A., Wisneski, C., and Ishii, H. Water Lamp and Pinwheels: Ambient Projection of Digital Information into Architecture Space. In *CHI'98 Conference Summary*, 269-270. ACM, 1998.

[24] Day, A. K. Understanding and Using Context. To appear in *Personal and Ubiquitous Computing*, Special issue on Situated Interaction and Ubiquitous Computing, 5(1), 2001.

[25] Donath, J. S. Visual Who: Animating the Affinities and Activities of an Electronic Community. In *Multimedia'95 Conference Proceeding*, 99-107. ACM, 1995.

[26] Dourish, P., Adler, A., Bellotti, V., and Henderson, A. Your Place or Mine? Learning from Long-term Use of Video Communication. *Computer-Supported Cooperative Work*, 5(1), 33-62. Kluwer, 1996.

[27] Dourish, P. and Bellotti, V. Awareness and Coordination in Shared Workspaces. In *CSCW'92 Conference Proceeding*, 107-114. ACM, 1992.

[28] Dourish, P. and Bly, S. Portholes: Supporting Awareness in a Distributed Work Group. In *CHI'92 Conference Proceeding*, 541-547. ACM, 1992.

[29] Erickson, T., Smith, D. N., Kellogg, W. A., Laff, M. R., Richards, J. T., and Bradner, E. Socially Translucent Systems: Social Proxies, Persistent Conversation, and the Design of "Babble". In *CHI'99 Conference Proceeding*, 72-79. ACM, 1999.

[30] Fish, R. S., Kraut, R. E., and Chalfonte, B. L. The VideoWindow System in Informal Communications Shared Video Spaces. In *CSCW'90 Conference Proceeding*, 1-11. ACM, 1990.

[31] Fish, R. S., Kraut, R. E., Root, R. W., and Rice, R. E. Video as a Technology for Informal Communication. *Communications of the ACM*, Vol. 36, No. 1, 48-61. ACM, January 1993.

[32] Fitzpatrick, G., Parsowith, S., Segall, B., and Kaplan, S. Tickertape: Awareness in a Single Line. In *CHI'98 Summary*, 281-282. ACM, 1998.

[33] Fogarty, J., Forlizzi, J., and Hudson, S. E. Aesthetic Information Collages: Generating Decorative Displays that Contain Information. In submission to *UIST 2001*.

[34] Gaver, W. The Affordances of Media Spaces for Collaboration. In *CSCW'92 Conference Proceeding*, 17-24. ACM, 1992.

[35] Gaver, W., Moran, T., MacLean, A., Lövstrand, L., Dourish, P., Carter, K., and Buxton, W. Realizing a Video Environment: EuroPARC's RAVE System. In *CHI'92 Conference Proceeding*, 27-35. ACM, 1992.

[36] Gaver, W., Sellen, A., Heath, C., and Luff, P. One is not Enough: Multiple Views in a Media Space. In *InterCHI'93 Conference Proceeding*, 335-341. ACM, 1993.

[37] Georgia Navigator. At `http://www.georgia-navigator.com/traffic/`.

[38] Ginsberg, A. and Ahuja, S. Automating Envisionment of Virtual Meeting Room Histories. In *Multimedia'95 Conference Proceeding*, 65-75. ACM, 1995.

[39] Greenberg, S. A Fisheye Text Editor for Relaxed-WYSIWIS Groupware. In *CHI'96 Conference Companion*, 212-213. ACM, 1996.

[40] Greenberg, S. Peepholes: Low Cost Awareness of One's Community. In *CHI'96 Conference Companion*, 206-207. ACM, 1996.

[41] Greenberg S. and Roseman, M. Using a Room Metaphor to Ease Transitions in Groupware. To appear in Ackerman, M., Pipek, V., Wulf, V. (eds.) *Beyond Knowledge Management: Sharing Expertise*. MIT Press.

[42] Greenberg, S. and Rounding, M. The Notification Collage: Posting Information to Public and Personal Displays. In *CHI 2001 Conference Proceeding*, 514-521. ACM, 2001.

[43] Grudin, J. Partitioning Digital Worlds: Focal and Peripheral Awareness in Multiple Monitor Use. In *CHI 2001 Conference Proceeding*, 458-465. ACM, 2001.

[44] Godwin, M. Nine Principles for Making Virtual Communities Work. *Wired*, 2.06: 72-73. 1994.

[45] Goldberg, Y., Safran, M., and Shapiro, E. Active Mail – A Framework for Implementing Groupware. In *CSCW'92 Conference Proceeding*, 75-83. ACM, 1992.

[46] Gutwin, C., Greenberg, S., and Roseman, M. Workspace Awareness in Real-Time Distributed Groupware: Framework, Widgets, and Evaluation. In *Proceedings of the HCI'96 Conference on People and Computers*, 281-298. Springer-Verlag, 1996.

[47] Handel, M. and Wills, G. TeamPortal: Providing Team Awareness On the Web. In *Personal and Ubiquitous Technologies*, vol. 5(1). Springer-Verlag, 2001.

[48] Heiner, J. M., Hudson, S. E., and Tanaka, K. The Information Percolator: Ambient Information Display in a Decorative Object. In *UIST'99 Conference Proceeding*, 141-148. ACM, 1999.

[49] Helfman, J. I. *Image Representations for Access and Similarity-Based Organization of Web Information*. Ph.D. thesis, Department of Computer Science, University of New Mexico, July 1999.

[50] ht://Dig Group, the, at `http://www.htdig.org/`.

[51] Hudson, S. E. and Smith, I. Techniques for Addressing Fundamental Privacy and Disruption Tradeoffs in Awareness Support Systems. In *CSCW'96 Conference Proceeding*, 248-257. ACM, 1996.

[52] Ishida, T. Towards Communityware. In *Proceedings of the International Conference and Exhibition on the Practical Application of Intelligent Agents and Multi-Agent Technology* (PAAM-97), 7-21. 1997.

[53] Ishii, H. and Kobayashi, M. ClearBoard: A Seamless Medium for Shared Drawing and Conversation with Eye Contact. In *CHI'92 Conference Proceeding*, 525-532. ACM, 1992.

[54] Ishii, H., Wisneski, C., Brave, S., Dahley, A., Gorbet, M., Ullmer, B., and Yarin, P. ambientROOM: Integrating Ambient Media with Architectural Space. In *CHI'98 Conference Summary*, 173-174. ACM, 1998.

[55] Issacs, E., Walendowski, A., and Ranganathan, D. Hubbut: A Wireless Instant Messenger That Uses Earcons for Awareness and for "Sound Instant Messages". In *CHI 2001 Extended Abstracts*, 3-4. ACM, 2001.

[56] Isaacs, E. A., Tang, J. C., and Morris, T. Piazza: A Desktop Environment Supporting Impromptu and Planned Interactions. In *CSCW'96 Conference Proceeding*, 315-333. ACM, 1996.

[57] Jacobson, V. and McCanne, S. vat - LBNL Audio Conferencing Tool. At `http://www-nrg.ee.lbl.gov/vat/`.

[58] Jancke, G., Venolia, G. D., Grudin, J., Cadiz, J. J., and Gupta, A. Linking Public Spaces: Technical and Social Issues. In *CHI 2001 Conference Proceeding*, 530-537. ACM, 2001.

[59] Jeffrey, P. and McGrath, A. Sharing Serendipity in the Workplace. In *CVE 2000 Conference Proceeding*, 173-179. ACM, 2000.

[60] Kamba, T., Elson, S. A., Harpold, T., Stamper, T., and Sukaviriya, P. Using Small Screen Space More Efficiently. In *CHI'96 Conference Proceeding*, 383-390. ACM, 1996.

[61] Kerne, A. CollageMachine: Temporality and Inderterminacy in Media Browsing via Interface Ecology. In *CHI'97 Conference Extended Abstracts*, 238-239. ACM, 1997.

[62] Kraut, R., Egido, C., and Galegher, J. Patterns of Contact and Communication in Scientific Research Collaboration. In *CSCW'88 Conference Proceeding*, 1-12. ACM, 1988.

[63] Kraut, R. E., Fish, R. S., Root, R. W., and Chalfonte, B. L. Informal Communication in Organizations: Form, Function, and Technology. In Oskamp, S., and Spacapan, S. (Eds.) *People's Reactions to Technology in Factories, Offices, and Aerospace*, The Claremont Symposium on Applied Social Psychology, 145-199. Sage Publications, 1990.

[64] Krueger, M. W. Videoplace. *Artificial Reality II*, 34-41. Addison-Wesley, 1991.

[65] LaMarca, A. At http://www.lamarca.org/anthony/.

[66] Leahy, L., Houde, S., and Bellamy, R. An On-line Newspaper for Improving Communications within Communities. *The Apple Education Learning Technology Review*, Issue 3, 14-23. Apple Computer, Spring 1998.

[67] Lee, A., Girgensohn, A., and Schlueter, K. NYNEX Portholes: Initial User Reactions and Redesign Implications. In *GROUP'97 Conference Proceeding*, 385-394. ACM, 1997.

[68] Linux Documentation Project. Linux Frequently Asked Questions with Answers. Available at http://www.linuxdoc.org/.

[69] Lyman, P. and Varian, H. R. How Much Information? Available at http://www.sims.berkeley.edu/how-much-info/.

[70] Lynx. Available at http://lynx.browser.org/.

[71] Maglio, P. P., Barrett, R., Campbell, C. S., and Selker, T. SUITOR: An Attentive Information System. In *IUI 2000 Conference Proceedings*, 169-175. ACM, 2000.

[72] Maglio, P. P. and Campbell, C. S. Tradeoffs in Displaying Peripheral Information. In *CHI 2000 Conference Proceedings*, 241-248. ACM, 2000

[73] McCanne, S. and Jacobson, V. vic: A Flexible Framework for Packet Video. In *Multimedia'95 Conference Proceeding*, 511-522. ACM, 1995.

[74] McCanne, S. and Jacobson, V. vic - Video Conferencing Tool. At http://www-nrg.ee.lbl.gov/vic/.

[75] McCrickard, D. S. *Maintaing Information Awareness in a Dynamic Environment: Assessing Animation as a Communication Mechanism*. Ph.D. thesis, College of Computing, Georgia Institute of Technology, June 2000.

[76] McCrickard, D. S. and Zhao, Q. A. Supporting Information Awareness Using Animated Widgets. In *the 7th USENIX Tcl/Tk Conference Proceedings*, 117-127. USENIX, 2000.

[77] Merriam-Webster's Collegiate Dictionary. At http://www.m-w.com/.

[78] Microsoft. Microsoft PowerPoint. See http://www.microsoft.com/.

[79] Microsoft Network. MSN Messenger. See http://www.msn.com/.

[80] Miller, T., and Stasko, J. The InfoCanvas: Information Conveyance through Personalized, Expressive Art. In *CHI 2001 Extended Abstracts*, 305-306. ACM, 2001.

[81] Mynatt, E., Back, M., Want, R., Baer, M., and Ellis, J. Designing Audio Aura. In *CHI'98 Conference Proceeding*, 566-573. ACM, 1998.

[82] Mynatt, E. D., Rowan, J., Craighill, S., and Jacobs, A. Digital Family Portraits: Supporting Peace of Mind for Extended Family Members. In *CHI 2001 Conference Proceeding*, 333-340. ACM, 2001.

[83] Mynatt, E. and Tullio, J. Inferring Calendar Event Attendance. In *IUI 2001 Conference Proceeding*, 121-128. ACM, 2001.

[84] Netscape. My Netscape, at `http://my.netscape.com`.

[85] Olson, G. M. and Olson, J. S. Distance Matters. To appear in Carroll, J. (ed.) *HCI in the New Millennium*. Addison-Wesley, 2001.

[86] Olson, M. H. and Bly, S. A. The Portland Experience: A Report on a Distributed Research Group. In Greenberg, S. (ed.) Computer Supported Cooperative Work and Groupware, 81-98. Academic Press, 1991.

[87] Ousterhout, J. K. Scripting: Higher Level Programming for the 21st Century. *IEEE Computer*, March 1998.

[88] Ousterhout, J. K. *Tcl and the Tk Toolkit*. Addison-Wesley, 1994.

[89] Palen, L. and Grudin, J. Discretionary Adoption of Group Support Software. To appear in Munkvold, B. E. (ed.), *Organizational Implementation of Collaboration Technology.* 2001.

[90] Pedersen, E. R. and Sokoler, T. AROMA: Abstract Representation of Presence Supporting Mutual Awareness. In CHI'97 Conference Proceeding, 51-58. ACM, 1997.

[91] Pickering, J. M. and King, J. L. Hardwiring Weak Ties: Individual and Institutional Issues in Computer Mediated Communication. In *CSCW'92 Conference on Proceeding*, 356-361. ACM, 1992.

[92] PointCast. At `http://www.pointcast.com/`.

[93] Putnam, R. D. *Bowling Alone: The Collapse and Revival of American Community.* Simon and Schuster, 2000.

[94] Resnick, P. Beyond Bowling Together: SocioTechnical Capital. To appear in Carroll, J. (ed.) *HCI in the New Millenium*. Addison-Wesley, 2001.

[95] Rogers, J. D. Theoretical Consideration of Collaboration in Scientific Research. In Hauger, J. S. and McEnaney, C. (eds.) *Strategies for Competitiveness in Academic Research*. AAAS, 2000.

[96] Roseman, M. and Greenberg, S. TeamRooms: Network Places for Collaboration. In *CSCW'96 Conference Proceeding*, 325-333. ACM, 1996.

[97] Rossi, P. H., Freeman, H. E., and Lipsey, M. W. *Evaluation: A Systematic Approach*, 6th edition. Sage Publications, 1999.

[98] Sawhney, N., Wheeler, S., and Schmandt, C. Aware Community Portals: Shared Information Appliances for Transitional Spaces. In *Personal and Ubiquitous Technologies*, vol. 5(1). Springer-Verlag, 2001.

[99] Schwab, I., Pohl, W., and Koychev, I. Learning to Recommend from Positive Evidence. In IUI 2000 Conference Proceedings, 241-247. ACM, 2000.

[100] Short, J., Williams, E., and Christie, B. Visual Communication and Social Interaction. In *The Social Psychology of Telecommunications*, 43-60. John Wiley & Sons, 1976.

[101] Sibert, L. E. and Jacob, R. J. K. Evaluation of Eye Gaze Interaction. In *CHI 2000 Conference Proceedings*, 281-288. ACM, 2000.

[102] Slashdot. News for Nerds, Stuff that Matters. Available at `http://slashdot.org/`.

[103] Smith, M., Cadiz, J. J., Burkhalter, B. Conversation Trees and Threaded Chats. In *CSCW 2000 Conference Proceeding*, 97 - 105. ACM, 2000.

[104] Smith, M. A., Farnham, S. D., and Drucker, S. M. The Social Life of Small Graphical Chat Spaces. In *CHI 2000 Conference Proceeding*, 462-469. ACM, 2000.

[105] Smith, M. K. Community: An Essay Prepared for the Informal Education Homepage. Available at `http://www.infed.org/community/community.htm`.

[106] Sohlenkamp, M. and Chwelos, G. Integrating Communication, Cooperation, and Awareness: The DIVA Virtual Office Environment. In *CSCW'94 Conference Proceeding*, 331-343. ACM, 1994.

[107] Sproull, L. and Kiesler, S. *Connections: New Ways of Working in the Networked Organization*. MIT Press, Cambridge, MA, 1991.

[108] Squeaker's. At `http://www.cc.gatech.edu/projects/squeakers/`.

[109] Steinfield, C., Jang, C. Y., and Pfaff, B. Supporting Virtual Team Collaboration: The TeamSCOPE System. In *Group'99 Conference Proceeding*, 81-89. ACM, 1999.

[110] Tang, J. C., Isaacs, E., and Rua, M. Supporting Distributed Groups with a Montage of Lightweight Interactions. In *CSCW'94 Conference Proceeding*, 23-34. ACM, 1994.

[111] Thompson, T. N. Organizational Creation and Evolution. In Byers, P. Y. (ed.), *Organizational Communication: Theory and Behavior*, 235-256. Allyn and Bacon, 1997.

[112] Thorngate, W. More Than We Can Know: The Attentional Economics of Internet Use. In Kiesler, S. (ed.) *Culture of the Internet*. Lawrence Erlbaum Associates, 1997.

[113] Tollmar, K., Sandor, O., and Schömer, A. Supporting Social Awareness @Work, Design, and Experience. In *CSCW'96 Conference Proceeding*, 298-307. ACM, 1996.

[114] Trethewey, A. Organizational Culture. In Byers, P. Y. (ed.), *Organizational Communication: Theory and Behavior*, 203-234. Allyn and Bacon, 1997.

[115] Van Laar, C. What is Sense of Community? In *Increasing a Sense of Community in the Military: The Role of Personnel Support Programs*, RAND Report MR-1071-OSD, 5-8. RAND, 1999.

[116] Viegas, F. B., and Donath, J. S. Chat Circles. In *CHI'99 Conference Proceeding*, 9-16. ACM, 1999.

[117] Want, R., Hopper, A., Falcão, V., and Gibbons, J. The Active Badge Location System. *ACM Transactions on Information Systems*, vol. 10(1), 91-102. ACM, 1992.

[118] Weather Channel, The. At `http://www.weather.com/`.

[119] Weiser, M. and Brown, J. Designing Calm Technology. *PowerGrid Journal*, v1.01, July 1996. Available at `http://powergrid.electriciti.com/1.01`.

[120] Whittaker, S., Frohlich, D., and Daly-Jones, O. Informal Workplace Communication: What is it like and how might we support it? In *CHI'94 Conference Proceeding*, 130-137. ACM, 1994.

[121] Whittaker, S., Swanson, J., Kucan, J., and Sidner, C. TeleNotes: Managing Lightweight Interactions in the Desktop. *ACM Transactions on Computer-Human Interaction* vol. 4(2), 137-168. ACM, 1997.

[122] Wired News. At http://www.wired.com/.

[123] Wolf, T. Creating Virtual Places for Community Involvement. In *Personal and Ubiquitous Technologies*, vol. 5(1). Springer-Verlag, 2001.

[124] World Wide Web consortium, The. Extensible Markup Language (XML). Available at http://www.w3.org/XML/.

[125] Xiong, R. and Donath, J. PeopleGarden: Creating Data Portraits for Users. In *UIST'99 Conference Proceeding*, 37-44. ACM, 1999.

[126] Yahoo!. At http://www.yahoo.com/.

[127] You, Y. A Survey for the Study of Awareness in Co-operative Systems. In Information Systems Research Seminar in Scandinavia (IRIS) 32 Conference Proceeding. 2000.

[128] Zhao, Q. A. and Stasko, J. T. Evaluating Image Filtering Based Techniques in Media Space Applications. In *CSCW'98 Conference Proceedings*, 11-18. ACM, 1998.

[129] Zhao, Q. A. and Stasko, J. T. The Awareness-Privacy Tradeoff in Video Supported Informal Awareness: A Study of Image-Filtering Based Techniques. Graphics, Visualization, and Usability Center, Georgia Institute of Technology, Atlanta, GA, Technical Report GIT-GVU-98-16, June 1998.

[130] Zhao, Q. A. and Stasko, J. T. What's Happening? The Community Awareness Application. In *CHI 2000 Extended Abstracts*, 253-254. ACM, 2000.

Vita

Qiang Alex Zhao was born and raised in Beijing, People's Republic of China. He entered Tsinghua University in 1987 to pursue a Bachelor of Engineering degree as a Computer Science and Technology major. After an unforgettable college life, he came to the United States in 1991 and graduated with a Master of Science degree in Computer Science from the University of Arizona in 1993. He moved to Atlanta afterwards to attend Georgia Institute of Technology and will receive his Doctor of Philosophy degree in Computer Science with a minor in Social Psychology in December 2001.