

Providing Artifact Awareness to a Distributed Group through Screen Sharing

Kimberly Tee
University of Calgary
2500 University Drive NW
Calgary, AB, Canada T2M 1N4
+1 403 210 9499
tee@cs.ucalgary.ca

Saul Greenberg
University of Calgary
2500 University Drive NW
Calgary, AB, Canada T2M 1N4
+1 403 220 6087
saul@cs.ucalgary.ca

Carl Gutwin
University of Saskatchewan
57 Campus Drive
Saskatoon, SK, S7N 5A9
+1 306 966 8646
gutwin@cs.usask.ca

ABSTRACT

Despite the availability of awareness servers and casual interaction systems, distributed groups still cannot maintain *artifact awareness* – the easy awareness of the documents, objects, and tools that other people are using – that is a natural part of co-located work environments. To address this deficiency, we designed an awareness tool that uses screen sharing to provide information about other people’s artifacts. People see others’ screens in miniature at the edge of their display, can selectively raise a larger view of that screen to get more detail, and can engage in remote pointing if desired. Initial experiences show that people use our tool for several purposes: to maintain awareness of what others are doing, to project a certain image of themselves, to monitor progress and coordinate joint tasks, to help determine when another person can be interrupted, and to engage in serendipitous conversation and collaboration. People have also been able to balance awareness with privacy, by using the privacy protection strategies built into our system: restricting what parts of the screen others can see, specifying update frequency, hiding image detail, and getting feedback of when screenshots are taken.

Categories and Subject Descriptors

H.5.3 [Group and Organization Interfaces]: Computer supported cooperative work

General Terms

Design, Human Factors.

Keywords

Screen sharing, artifact awareness, distributed groupware.

1. INTRODUCTION

Casual interaction – the brief, unplanned meetings that commonly occur during the day between co-located people – is important for coordinating and tracking the progress of joint work, for exchanging knowledge and information, and for building relationships [13, 22]. Casual interaction is made possible by *interpersonal awareness*, the understanding of who is around, what they are doing, and whether or not they are available for conversation and collaboration [13]. Interpersonal awareness is easy to maintain in a co-located setting, but casual interaction is problematic for distributed groups, due to the lack of awareness cues. Consequently, distributed collaborators must put a relatively

large amount of effort into coordinating interaction (e.g., by scheduling meetings). This effort is a problem, because many opportunistic interactions would not occur if they had to be planned in advance [13], suggesting that distributed groups are missing out on valuable opportunities for collaboration. This partially explains the explosion of low-effort awareness servers and casual interaction systems, such as the widespread adoption of instant messengers by diverse user groups [17].

An important component of informal awareness that is not handled by awareness servers and instant messengers, however, is *artifact awareness*, defined as one person’s knowledge of the artifacts and tools that other people are working with. For office workers, artifacts include the documents and drawings (both physical and digital) that people work on over the course of a day, the secondary materials that support their tasks, and the tools they use to carry out their work. Being aware of these artifacts is valuable for a variety of reasons:

- *Monitoring and coordinating*: Collaborators who are responsible for different aspects of a joint task can monitor each other’s progress and coordinate their activities.
- *Triggering interest*: Seeing another person’s activity, even if it is not part of a joint task, can trigger interest in that activity. For example, Greenberg [8] presented situations where people initiated interactions when a person’s activity became publicly visible (such as joining a child’s videogame when it was visible on a television set).
- *Determining availability*: Knowledge of artifacts is yet another source of contextual information that helps one determine how busy people are and when they can be interrupted [13, 22].
- *Creating serendipitous opportunities*: Artifact information creates opportunities for people to engage in artifact-oriented conversations, and to move into collaboration over the artifact. For example, Whittaker and colleagues [22] found that over half of all casual interactions in an office involved some form of document sharing, where documents were mostly used as a cue or conversational prop. Similarly, Nardi [16] found that people opportunistically collaborate over spreadsheets, CAD systems, and other documents over the course of a day.

Artifact awareness is easy when people are in a co-located environment – when people inhabit a shared office space, or when one person does a *walkabout* to see what others are doing [2] – because people naturally gather visual and auditory cues about other people’s presence and activities. They notice what artifacts others are working on as they glance into people’s offices, look at

their desks, and see their computer displays. In the real world, for example, if a document or article on another's desk catches someone's eye as he/she is walking by, he/she can stop and discuss the artifact with its owner. This awareness can be very important, especially for designers who rely on easy visibility of other people's work. These groups intentionally work in studio spaces to promote learning, reflection, and discussion about current projects – designs are placed on easels or other semi-public surfaces, and others moving through the studio can monitor and comment on the work as it unfolds over time. Similarly, many command and control situations – such as air traffic control, subway routing, or shipboard navigation – rely on people being able to see one another's artifacts [11]. Artifact visibility also plays a role in how people create common ground in conversations [5].

Our concern is that in spite of the availability of awareness servers and casual interaction systems, distributed groups still lack the easy awareness of others' artifacts that is normally found in a co-located shared environment [22]. While there are many groupware systems that let a distributed group share artifacts (e.g., shared editors and webcast meeting tools), they work only *after* interaction is initiated. That is, they are intended for focused collaborative work rather than artifact awareness.

For distributed groups, there is still no real equivalent to the way that co-located people can visually share their individual work and maintain artifact awareness. Consequently, we set ourselves the following research goal:

Design a system that supports artifact awareness and opportunistic interaction, where the audience is a distributed group with a strong desire to stay in up-to-the-moment contact.

To satisfy this goal, we created an awareness tool that uses *screen sharing* to provide the group with mutual awareness of artifacts and individual work. Our hypothesis is that screen sharing is a critical component of distributed artifact awareness, because the screen contents capture almost all of a person's digital activities.

This paper begins with a brief discussion of previous work on awareness tools and screen sharing. We then describe the system that we built and its design rationale. Finally, we report on the initial experiences our group had using it, and ways that the idea can benefit distributed work groups.

2. RELATED WORK

2.1 Awareness / Casual Interaction Systems

The majority of research on providing informal awareness and casual interaction to distributed groups has focused on instant messaging (IM), virtual worlds, and media space systems. IM [17], which is extremely popular both at home and in the workplace, generally provides status indicators showing if a person on a buddy list is logged on, busy, away, or off-line. One initiates conversation by selecting that person's name. Until recently, IM conversations were restricted to text chat. Newer systems offer richer communication channels such as internet telephony (VoIP) and video, groupware applications such as shared drawings, games and shared screens, and information exchange via file transfer. While the awareness information provided by IM is very basic, it succeeds because it lets people

easily establish communication with one another at opportune times.

Virtual worlds include MUDs, chat rooms and even on-line games [1]. They let people meet other inhabitants of the virtual world, and participate in conversations and activities. In them, people can also view and manipulate visual artifacts that comprise the virtual world. While there is good awareness of virtual shared artifacts, this awareness is limited to artifacts in that world.

Media spaces [7] provide rich awareness by linking offices and public spaces through networks of audio and video. Through video, viewers can easily see who is around and what they are doing. However, actual work artifacts are rarely visible due to camera angle and resolution limitations. While people can shift the camera to focus on work activity, this is usually done after the conversation has started [19].

Some recent media spaces allow group members to augment the space by allowing them to post information publicly seen by all. Examples include editable sticky notes, web page thumbnails, and photos [9, 14, 20]. These contribute to public conversations and shared information in a real-time area that all members can see. However, they do not provide awareness of others' actual individual work, unless this information is explicitly offered during conversation or in a broadcast message.

In summary, these groupware genres succeed because they promote interpersonal awareness leading to casual interaction. However, most do not promote artifact awareness until after people are engaged in interaction, are restricted to artifacts present in the virtual world, or have no means for gathering and displaying awareness of a person's individual work.

2.2 Screen Sharing

Screen-sharing applications let people explicitly share their computer screens, windows, or arbitrary screen regions with each other. These systems emulate over-the-shoulder sharing of a computer: one person can show others what they are working on, and each can take turns interacting with the system. This genre is often called 'collaboration-transparent' groupware as the underlying application being shared has no knowledge that multiple people are viewing and/or using it.

Screen sharing is an old idea. It was first demonstrated in 1968 by Engelbart [6], and has been used in commercial practice for almost twenty years, e.g., Farallon Timbuktu [23] and VNC [18]. It is often a key component of desktop conferencing systems, in which audio/video teleconferencing technologies are integrated with desktop computer application sharing in order to allow individuals to meet, collaborate and work together from their offices. Other common uses of screen sharing include application sharing (desktop conferencing without requiring the use of audio/video) and remote assistance (a system administrator or an expert user can remotely control another's computer to assist them in performing certain tasks).

Screen sharing systems are designed for focused interaction rather than artifact awareness leading to casual interaction. To our knowledge, no previous research has considered screen sharing as an awareness mechanism leading to casual interaction. The exception is preliminary unpublished work done by ourselves and our colleagues. We first raised this idea in the Notification

Collage Bridgit media item [20]: people could post an intermittently updating desktop thumbnail that others could see, and select it if desired to activate a full screen sharing session. Multi-VNC [10] then followed, as a proof-of-concept prototype in which multiple instances of an open-source screen-sharing application were run so that each member of a group could see the other members' computer desktops. Neither examined in detail how to provide awareness using screen sharing.

3. COMMUNITY BAR

We decided to create a screen-sharing tool within the Community Bar (CB), a media space that can be augmented with extra information. To set the scene, we briefly summarize the basic CB functionality here; see [14, 15] for a complete description of the design rationale and interface.

Figure 1 illustrates a screen snapshot of CB in use. As can be seen, it is a sidebar peripheral display, divided into *Places*; two are shown in the figure. Each place represents a sub-group, their communication, their tools, and their information. These are visualized through a number of *media items*, all holding different information and all being publicly visible to the group. As illustrated in the figure, the *Presence item* represents a person as live video, their image, or their name. *Chat items* hold multi-person public conversations, while *Stickies* contain one person's text posting to the group. People can post public information to the group such as photos or web pages through the *Photo item* and *Web item*.

In Figure 1, for example, we see that three people are in the 'mike test' place (shown as video), and that two parallel chats are going on. Two people are currently in the 'CSCW class' place (one represented by a snapshot), a chat is ongoing, and people have posted a photo and a web page to it. One person is also sharing their screen, which will be discussed in a later section.

Media items are presented at three levels of granularity. The media item's *tile* view is always visible in the sidebar. Individuals can choose to explore and interact with that information by mousing over the tile, which displays its *tooltip grande* next to it. For example, the tooltip grande in Figure 1 shows that person's video at a large size and faster frame rate. Finally, a person can click on the tooltip grande's title bar to raise the *full view window* (not shown), which displays even more information and interaction capabilities.

Media items also have the idea of an *owner* (the person who creates the media item) and an *audience* (all others who can see the item). Depending on the media item, the view and controls

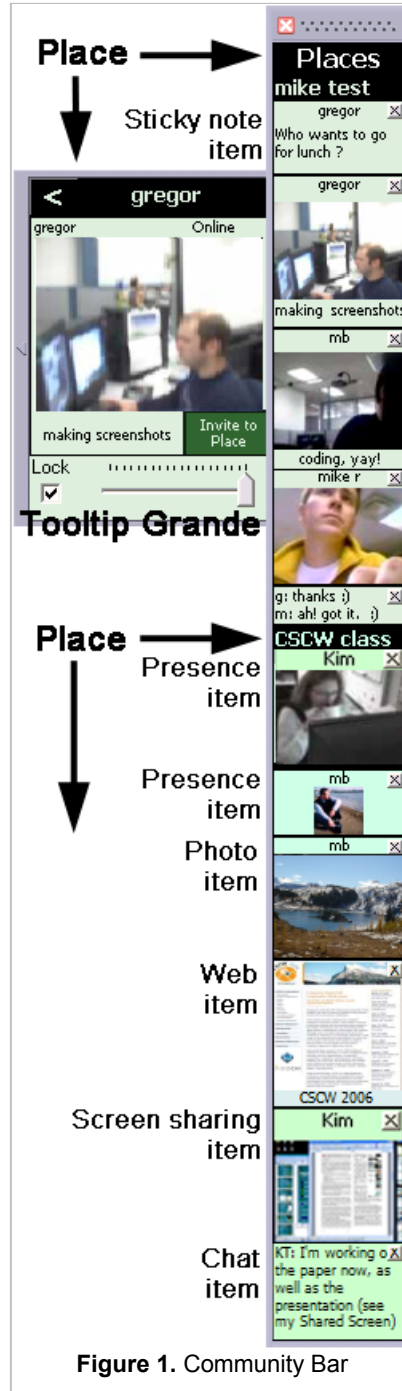


Figure 1. Community Bar

available on the tile / tooltip grande / full view may be different for the owner than for the audience. For example, the owner of a Presence item has additional controls in their full view to change what others see, e.g., a video, an image, or an icon. In contrast, all people see a Chat item in exactly the same way.

Fundamental to the philosophy of the Community Bar is that:

- All the media items within a place are publicly visible to all the people in that place, i.e., it serves as a virtual communal shared setting.
- While basic information is always visible at the periphery, progressively more information can be revealed through focused interaction.

The sidebar encourages peripheral awareness because it cannot be covered up and because it is situated at the screen's side. For each media item, its tile view generally shows awareness information; its tooltip grande shows more detailed information and allows partial interaction; while the full view shows all the information, communication, and interaction possibilities.

Finally, CB is an open-ended system based on a plug-in architecture. Using an API and development environment, 3rd party developers can create and add their own custom media items to CB without recompilation [15].

4. SCREEN SHARING ITEM

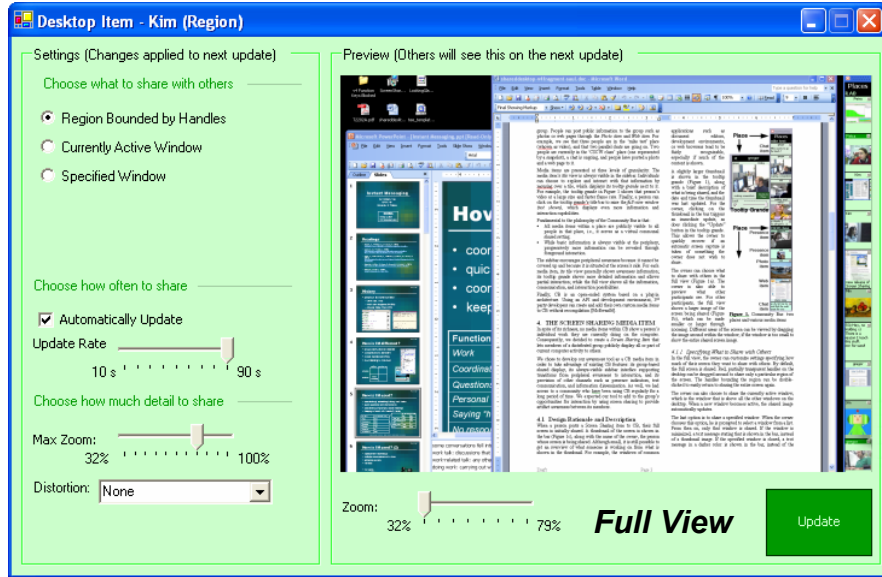
In spite of its richness, no media items within CB show the individual work that a person is currently engaged in on their computer. Consequently, we decided to create a *Screen Sharing item* that lets members of a distributed group publicly display all or part of their current computer screen to others.

We chose to develop our awareness tool as a CB media item in order to take advantage of existing CB features: its group-based public display, its always-visible sidebar interface supporting transitions from peripheral awareness to interaction, and its provision of


other communication and information channels such as presence indicators, text communication, and so on. That is, we expect our *Screen Sharing item* to work within the context of other interpersonal awareness information and public conversations. As well, we had access to a community who were using CB on a daily basis for over a year. By adding our new item to CB as used by this group, we could see how it provided artifact awareness between its members – this will be discussed in a later section. The following subsections describe by scenario how the screen sharing item works.



Tooltip Grande



A. The owner's *Tooltip Grande* and *Full View* of the Screen Sharing item.




Tooltip Grande



screen region

specified window



C. Tile view, all people

B. The audience's *Tooltip Grande* and *Full View* of the Screen Sharing item.

Figure 2. Various views of the Screen Sharing item.

4.1 Tile View

'Kim' (initials KT) is working on a paper of interest to her group, so she decides to share her display with them. Through the CB place's context menu, she invokes the *Screen Sharing* item. A tile view is immediately added to that CB place, which contains a thumbnail of her entire screen labeled with her name (Figure 1, 2nd tile from the bottom). At the same time, she adds a Chat item (Figure 1, bottom tile) saying "I'm working on the paper now, as well as the presentation (see my Shared Screen)" – this gives the group some context to help interpret the image.

By default, this thumbnail is updated once a minute. However, the owner (and only the owner) can trigger an immediate update by clicking the tile. For example, Kim may do this to rapidly replace a screen shot that she did not want others to see, or (more typically) to show others changed screen content in a timely way,

e.g., as part of a discussion of the image that may be occurring in an ongoing text chat.

Although small, the thumbnail and its update frequency suffices to provide all others in that CB place with an overview of what Kim is working on. Typically, the visual characteristics of windows within the thumbnail are sufficiently recognizable [12] so that others can tell if the poster is editing a document, browsing the web, writing a program, preparing a presentation, using an instant messenger, etc. While actual content is hard to distinguish, visual landmarks such as distinctive application graphics, photos and text formatting are discernable. For example, the reasonably recognizable thumbnails in Figure 3 show people visiting a web page, editing a Microsoft Word document with track changes turned on, looking at email through the Outlook email reader and checking MSN Messenger's buddy list, using the Picasso photo viewer, and working on a Powerpoint presentation.




Figure 3. Thumbnails of people's screens.

Now reconsider the shared screen in the tile from Figure 1, shown again in Figure 2c. From the audience's previous knowledge of Kim's work habits and from the contents of her Chat item, they correctly guess that she is currently editing a document in Microsoft Word, where the Powerpoint presentation she mentions is partially visible in the background. The audience can also tell that this is a two-column document typical of most ACM papers, and that a figure is positioned at the top right of the page. If the viewer is a co-author of this paper, then that person could likely guess what page that is from their knowledge of the paper.

4.2 Tooltip Grande

The tooltip grande for both the owner and the audience shows a somewhat larger thumbnail (Figures 2a and 2b, left side). Akin to a glance, people in the CB place may raise this to help them further recognize certain features in the owner's screen. Atop the thumbnail is a brief description of what is being shared, i.e., the full screen, a region of the screen, the active window, or a particular window (to be discussed shortly). It also includes a timestamp indicating when the image was last updated. For example, the tooltip grandes of Kim's desktop in Figures 2a and 2b (left) show that she is sharing a region of the screen, and that it was last updated at 12:31 PM today.

The lock and slider that appear at the bottom of the tooltip grande are common to all CB media items, and are used by the viewer to adjust the size of the tile in the sidebar [14]. When the tile is resized to dimensions that are too small for the thumbnail to be recognizable, the thumbnail is replaced by a descriptive text label. To illustrate, the bottom tile in Figure 2c is a Screen Sharing item posted by Stephanie, where she is sharing a specific window rather than a screen region. Using its tooltip grande, the current viewer has shrunk his view of Stephanie's tile; only a text description is displayed that gives the name of the window being shared. Finally, the arrow at the top left of the tooltip grande is also common to all CB media items, and is used to invoke the full view, discussed shortly.

While the audience and owner view of the tooltip grande are visually similar, there are some differences. As before, the owner (and only the owner) can trigger an immediate update by clicking the thumbnail or by pressing the "Update" button visible at the tooltip grande's bottom right (Figure 2a, left). The audience view of the tooltip grande does not have this button.

The audience also has a button in their view titled 'Remote Pointing', while the owner has its corollary 'Stop Remote'. This will be explained later.

4.3 Full View

The full view as seen by both the audience and the owner gives a larger and much more detailed preview of the captured display (Figures 2a and 2b, center). As with other views, this preview is live: its contents are replaced as updates come in. As visible on

the left side of the full view in Figure 2a, the owner has additional controls that will let him/her adjust and limit how the display is shared, thus providing some balance between awareness and privacy. These are described in Section 5.

A *zoom slider* below the image lets the viewer zoom into the image as desired for greater detail. When the zoomed-in image does not fit within the window, the viewer can pan the image by directly selecting and dragging it with the mouse. As we will shortly discuss, for privacy reasons the permissible level of zoom depends on how the owner has configured sharing, i.e., zooming may be restricted to much less than true screen resolution. For example, in the full views in Figures 2a and 2b, we see that Kim has set the maximum zoom level to 79%. The audience member is looking at her screen at this maximum zoom level (Figure 2b, full view), while Kim is looking at it at a 32% zoom so that the entire region fits the full view's window (Figure 2a, full view). We also see that 79% zoom of the true screen resolution produces a fairly legible image (Figure 2b, full view): subsection titles of the paper are easily visible, and the paper text can be read with some effort.

4.4 Moving to Interaction

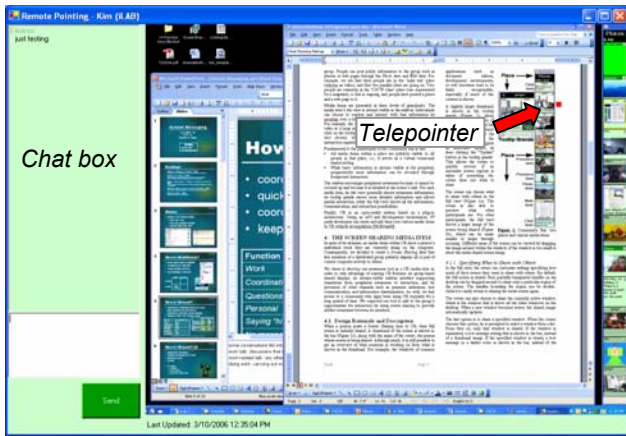
Any audience member can attempt to initiate a real-time remote pointing session. The owner sees a remote pointing request via a dialog box, and can approve or deny it. If the owner approves it, a *remote pointing window* (Figure 4a) appears on the screen of that audience member. This window displays the shared screen image at the maximum allowable resolution. It also includes a full-sized chat box linked to a Chat item in the sidebar (Figure 4a, left and right side). The audience member can then drag a small red telepointer around this image, visible at the top right of Figure 4a. A corresponding telepointer appears and moves around on the owner's actual desktop in the corresponding location, as shown in the partial screenshot of the owner's desktop in Figure 4b. Either the owner or the other participant can terminate the session at any time.

We recognize that remote pointing is not as powerful as systems that let people take turns interacting with the application [6, 18, 23]; this is something we plan to add, perhaps by implementing the VNC protocol [18]. Still, we believe remote pointing suffices for most situations. From observations of casual interactions in offices, Whittaker suggests "Document use indicates a requirement for simple systems rather than full-blown shared editors. A system that allowed mutual viewing of documents, with the ability to point at and possibly make simple annotations, may be all that is required here [22]."

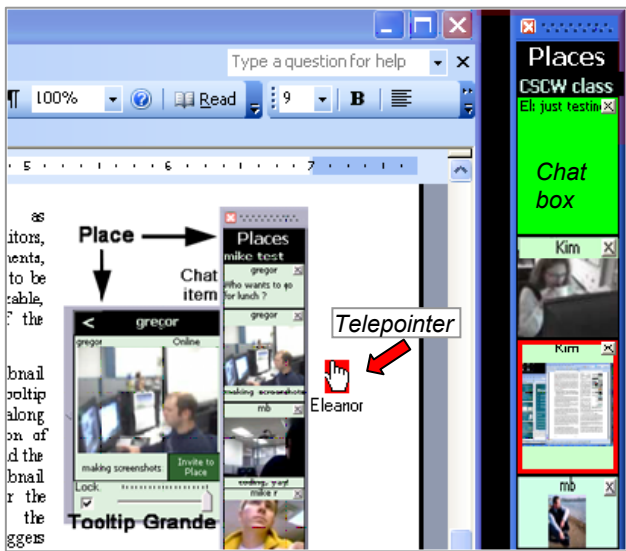
5. PRIVACY CONTROLS

Privacy is, of course, a serious consideration in an always-on screen sharing system. For example, we can easily imagine a situation when Kim inadvertently displays a sensitive email message that others should not be seeing. The challenge is how people can balance the awareness information they want others to have of their work with their own privacy needs.

First and foremost, we stress that privacy is not just a technical issue [4]. Rather, it is heavily dependant on the group culture and the actual practice of use that develops over time. As an always-on media space, Community Bar is designed for a community of intimate collaborators who have a real desire and need to stay



A. Audience member view and telepointer control of another's desktop.



B. Partial screen snapshot of the owner's desktop, illustrating the telepointer. The red-outline in the Screen Sharing tile indicates that someone is looking at a full view of this screen.

Figure 4. A remote pointing session.

connected. This is akin to a shared office of close-knit workers (or close friends, or family members) that are comfortable with seeing each other as they move around the shared space, as well as any information they are working on. Of course, this intended use could be abused by (say) an office manager that insists that all employees use the Screen Sharing media item so that their work can be monitored. However, even in the benign case, we do realize that people may want some control over what others can see as well as feedback of when others are looking. Of course, the most restrictive control is to simply not show the item; this is the default, as the Screen Sharing media item only appears when created explicitly by the owner. Further controls and feedback offered by the Screen Sharing item are discussed in the following subsections.

5.1 Specifying What to Share with Others

Owners have full control over what to 'push' out as artifact awareness; the audience cannot 'pull' any extra information.

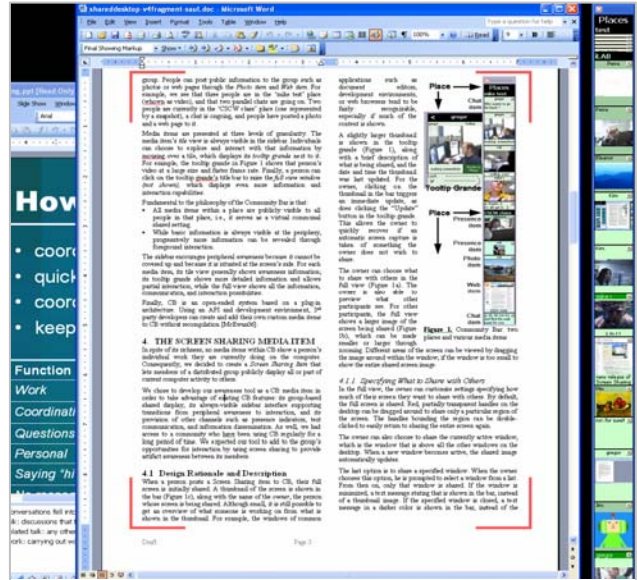
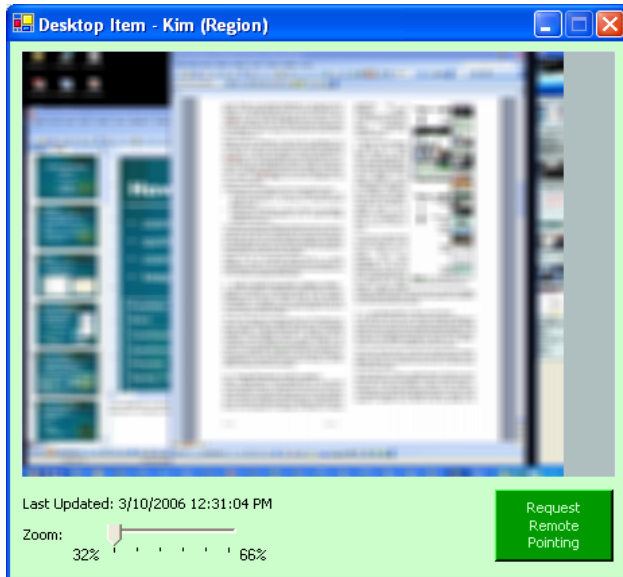


Figure 5. The capture region of the screen is marked by adjustable semi-transparent red handles (enhanced here).

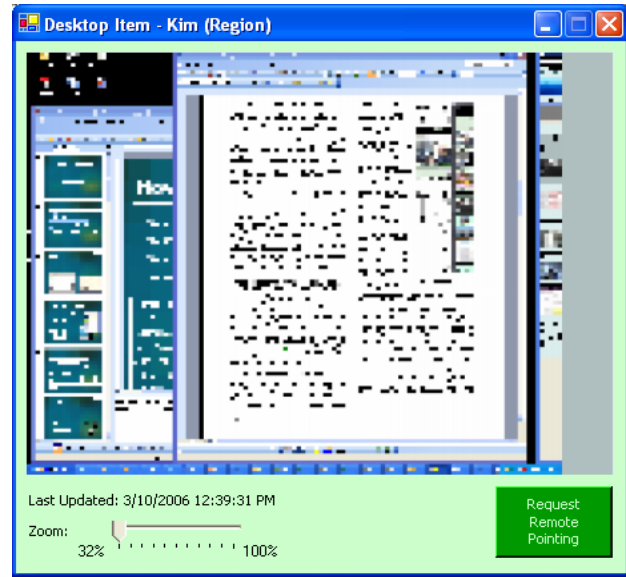
Thus, our first level of privacy control is to let the owner specify how much of the display he or she wishes to share with others. The owner can choose in the full view (Figure 2a, full view) to share: a particular screen region selected by handles (which can include a small area up to and including the entire screen), the currently active window, or a particular user-specified window. The choice restricts what others can see to only those parts of the display the owner wishes to reveal.

By default, a screen *region* encompassing the full screen is shared. Anything displayed in this window is captured: partial and overlapping windows, background wallpaper, dialog boxes. When this option is selected, the owner sees red, partially transparent handles on the actual display that define the bounding box of this region; an example is illustrated in Figure 5. The owner can quickly adjust the bounds of this captured region by dragging these handles around. For example, while Kim can share the entire display, she has repositioned the handles in Figure 5 to restrict sharing to the working area surrounding her text. She returns to sharing the entire screen by double-clicking a handle.

The owner can also share individual windows. First, the owner can selectively share the currently *active window*. This is the window that has the input focus, and that appears atop all the other windows on the desktop. As the user switches to a different window (thus making that one the active window), the shared image automatically updates to replace the old window with this new one. Second, the owner can share a *specified window* from a list of all windows. When selected, only that window is captured and shared (regardless of its position on the screen). If the owner minimizes or closes the window, an appropriate text message comprising the title of the window is shown instead of a thumbnail image, similar to the tile shown at the bottom of Figure 2c. When the owner resumes working in the window, the thumbnail will be displayed again.



A. Blurring with a low level of clarity



B. Pixelization with a medium level of clarity

Figure 6. Masking effects applied to the shared screen image

5.2 Specifying Update Frequency

Our second level of privacy control lets the owner specify how often the display should be captured, and thus how often the audience gets this update. This control allows owners to reduce temporal fidelity.

The owner can also specify if updates are manual or automatic. If manual, the display is only updated when the owner clicks the thumbnail in the file view, or the 'Update' button in the tooltip grande and full view. If automatic, the owner can specify an update frequency interval between 10 and 90 seconds (although clicking the thumbnail/update button will take an immediate snapshot).

Unlike commercial screen sharing systems tailored for real-time interaction, we believe this infrequent update suffices for artifact awareness. These updates still inform others of basic activities while minimizing distraction and privacy concerns that would otherwise arise from real-time movement in the various views.

5.3 Specifying How Much Detail to Share

Our third level of privacy control lets the owner manipulate the image fidelity that others can see. Techniques include adjusting the zoom level and distorting the image through image manipulation techniques. The less detail there is visible, the greater the perceived privacy.

An audience member is allowed to zoom into details in a shared screen image only up to a maximum zoom level set by the owner. Low zoom limits transform the image into a low resolution image. For example, if Kim set a low level zoom of around 33% and her captured region encompasses 1280x1024 pixels, the shared image is visually compressed to about 1/9 of the original area (~426 x 341). Alternatively, she can set an increasingly higher zoom limit, so that others can zoom in and view the shared image up to the original resolution. Depending on the zoom level, for example, Kim could set the level so that a viewer can read large-font

section headings in a text document, but not the actual text contents in paragraphs.

Alternatively or in combination with zoom limits, the owner can mask and distort the image by selecting one of several image masking effects. Current options include image blurring, pixelization, and image randomization; others could be added easily. These distortion techniques offer people a high degree of control of image fidelity not only in the thumbnails, but in the larger zoomed-in full views as well.

For example, Figure 6 illustrates what people would see when Kim uses the blur effect (6a), or the pixelate effect (6b) at 32% zoom. Both let others roughly see what Kim is working on, while preserving her privacy because the image does not reveal legible detail.

Another image masking possibility, which we did not implement, is to use the role-specific view-masking technique advocated by Berry, Bartram and Booth [3]. The idea is that the screen's owner can selectively mask portions of his/her screen image on a per person basis, where others may see different things depending on the settings.

5.4 Feedback of Image Capture

Our fourth level of privacy control is for the system to provide sufficient feedback to the owner about what others can see. First, the owner can always see exactly what the audience can see, because the Screen Sharing item is visible on both owners and audience members' sidebars. Similarly, if the owner raises the tooltip grande or the full view, they see the same image as the audience member.

Second, several mechanisms warn the owner just before an auto-update happens. About five seconds before the auto-update begins, the Screen Sharing item in the owner's bar is outlined in yellow. If sharing a region, the red handles that specify the image region turn yellow as well. Colors revert back to normal after the update is completed. This feedback aims to be a reasonable

compromise that reminds the owner of what is going on without being overly distracting.

Third, when an audience member opens the full view belonging to the owner, the Screen Sharing item in the owner's bar is outlined in red. This outline remains until the full view is closed. However, no identifying information is supplied as to who is looking at the full view.

5.5 Communal Feedback

Our fifth level of privacy control is social, as defined by the CB group. Screen Sharing items are visible only to the other people in the CB place. Because all people logged on in a CB place are visible, the owner of an item knows who can see. As well, because people in a CB place are part of a social group, one can reasonably expect – security violations aside – that only socially appropriate people can see it. Finally, because these people are expected to be colleagues, the viewers themselves can use the other facilities in the Community Bar to warn the owner about inappropriate things that are being shared.

6. INITIAL EXPERIENCES

People inhabiting a common space naturally see the screens of their co-workers as they glance around, or walk by, or when they are invited in for side-by-side work. However, the notion of using a shared screen in distributed groupware for artifact awareness is an unusual concept. We recognize that such a use will likely have to develop over time as part of a group's everyday practice.

Consequently, we decided to introduce the Screen Sharing item to our research group and collect their initial experiences and reaction to it. While this is a biased group, there are several good reasons for using them:

- the group has voluntarily used Community Bar for their own use on a daily basis for over a year;
- as a group with an established culture of use of a distributed system, they had already established a practice of balancing awareness provision with privacy;
- they were willing to aggressively use the Screen Sharing item for the initial deployment period;
- they were willing to report on their usage;
- as most were knowledgeable about awareness systems, they could provide not only comments of their personal use but reflective comments as well.

The group using CB included ten research assistants, graduate students, faculty, and former members of the lab. People were both co-located and distributed. Most lab members primarily worked in one of three connected laboratory spaces, while faculty was located in separate offices. This laboratory space was large enough that a person in one of the rooms would not normally see what a person in an adjoining room was doing. Group members were not always in the laboratory, as some tele-commuted when working at home. More people tele-commuted in the evenings. Former lab members also connected to CB from their distant work offices, one in the same city as the lab, one in a different city. Most members of this group had a webcam and two monitors connected to their computers, and regularly used CB as a peripheral display on one of them.

This was not a formal study; rather, we solicited comments as people used the system over a two week period, and we

interviewed several people to discuss details with them. We summarize their key experiences and reactions to the Screen Sharing item below.

6.1 Artifact Awareness

Most people shared their active window or the region of their screen that they were working in. From seeing these shared images over a period of time, people said they were typically able to identify what group members were working on. This information was used for several purposes.

First, the added knowledge of what a person was doing helped group members determine whether or not that person was interruptible. This adds to the other information available on CB (e.g., video) to help one make an informed decision on whether to initiate contact with this person.

Second, the Screen Sharing item helped people track the progress of joint work. For example, several members of the group were co-authoring papers during this deployment period. One member reported that because his co-author was using change tracking while editing their document, he was able to tell from the amount of red (changed) text seen in the Screen Sharing item that his co-author had been busy writing, and thus his own personal copy of the document was “stale”. His co-author had “definitely taken the ‘lock’ on this version.”

The Screen Sharing item was also used as an asynchronous awareness tool. One group member had been working on a paper and had shared the document window in CB. She then left the lab for a few hours, but kept her item active. While she was gone, her co-author logged onto CB. He noticed that the document was visible, and through looking at the document in the full view, he was able to see where she had left off working on the last page, and that she had not yet revised the text in the final section.

6.2 Presence Awareness

Besides being used to provide artifact awareness, we also saw the Screen Sharing item being used to provide presence information in the form of computer activity. One of the members of the group who did not have a webcam (and so only had a static image to indicate his presence on CB) used his Screen Sharing item as a replacement for his Presence item. Since he used a frequent auto-update rate for sharing his screen, it was easy to see when he was at his computer: windows would be scrolled up and down or be moved about. It was also possible to tell when he had been away from his computer for some time, as the Screen Sharing item's background color automatically darkened when the computer had been idle for five minutes or more.

6.3 Opportunistic Interactions

Conversations would sometimes arise as a consequence of people seeing artifacts in the Screen Sharing item. For example, one member saw his co-author working on their paper, and asked how it was going. His co-author responded “It's going ok – I've got some inspiration about how to proceed for a bit.” They then proceeded to coordinate when each would work on it, deciding that the co-author would continue writing for the day, and then pass the draft on to the other author. In another instance, one member of the group saw some interesting-looking design images on another's desktop. When asked what they were for, she was told that they were t-shirt designs. This led to a brief conversation

about that individual's extra-curricular activities outside the lab, which were not widely known before.

These conversations would occasionally transition into remote pointing sessions, which were often used to discuss joint work between two people. Most of these sessions occurred when at least one of the participants was working from home for the day. For example, a group member noticed that his co-author was working on a figure for their paper after he had sent her an email with some suggestions for improving it. They used remote pointing to discuss which parts of the figure should be changed. Afterwards, the group member was able to peripherally see his co-author making refinements to the figure, and she would intermittently ask him to check his view of her shared screen in order to get feedback on the image.

6.4 Focused Collaboration

People also reported using the Screen Sharing item for focused interactions after they had already begun a conversation or a meeting. For example, two co-authors were discussing a paper via VoIP and had to look at an image. They started a remote pointing session and used the telepointer to make sure they were talking about the same parts of the image. In another instance, one group member had asked for some visualizations from another member on CB. She wasn't sure exactly which ones he wanted, so she shared them on her screen in order to confirm with him they were the correct ones.

6.5 Privacy Issues

People adopted different strategies to protect their privacy while sharing their screens.

First, because most members of the group had two monitors, some chose the strategy of separating public and private information onto different displays. Some chose to share their entire public screen, while others chose to share only a specific region of it. For example, one group member reported "I have a two screen system, where I normally read email on the right screen, and do work on the left. I decided that I am happy to share my work (left) screen, so I set the region to the top half of that (using the idea that things above the fold are more relevant)." One person questioned this public / private separation as it differed from real life activities. He liked having a screen where things weren't publicly visible, but he wasn't sure why, "because anyone can walk by [in the lab] and see [it]".

Second, some people chose to blur the screen image they shared so that text in windows would not be legible to others. In fact, we saw that almost everyone who chose to share their active window opted to blur it. This is because unlike the private/public separation strategy for screen sharing mentioned above, sharing the active window is indiscriminate in what it displays. We stress that even in this case, people didn't have a problem with others being able to see the basic tasks they were working on. Rather, they were not always comfortable with sharing the details. This was true of activities involving personal communication, such as checking email or instant messaging chats. People who had to work with confidential information such as study data (protected by ethics reporting) or product source code (e.g., people working offsite in industry) also blurred their shared screen images to obscure the details, but still felt comfortable giving others an idea of what they were doing.

Third, people generally limited the amount others could zoom into their full view to less than the original resolution. When people moved into interaction, such as during discussions about shared artifacts visible in the full view or during remote pointing sessions, people would increase the maximum zoom. After these discussions, they would then decrease it.

Fourth, people reported that the auto-update feedback, where the handles and tile changed color to indicate an impending update, was particularly effective as it served as a constant reminder that the region was being shared. There were several cases reported where private windows were almost shared when they should not have been. For example, one faculty member began setting examination questions on his public screen, where he normally did his work. Fortunately, the visible warning from the Screen Sharing item that an automatic update was about to occur reminded him that the exam questions should not be publicly available. He then moved the document from the shared screen region to the private secondary screen before these exam questions could be seen by others.

When feedback from the Screen Sharing item indicated to people that someone else was looking at their desktop using the full view, they felt more conscious of what they were sharing, especially since they were not able to identify who was looking at their desktop unless that person explicitly told them. This feedback unintentionally discouraged people from looking at others' shared screens in the full view; one group member reported that there were times when he had wanted to look at another's desktop using the full view, but was slightly hesitant to do so because the other person would then know that someone was looking. This suggests that additional information should be supplied, e.g., the equivalent of the glance feature in Montage that shows that people are about to look in and that identifies them [21].

There was also some concern from audience members that people could see too much of others' desktops. After observing on CB that one person was composing email, and that another was reading sports news online, one group member commented "So here I am perusing people's desktops [in the full view]... Hmm, am I seeing too much?" This concern was surprising; we expected that people sharing their desktops would be concerned about sharing too much information, but we did not expect that audience members would feel uncomfortable seeing too much of someone else's desktop. This idea of 'reveal' [4] can actually heighten privacy, as it allows one person to warn others when they are unintentionally revealing something. For example, one person noticed that a colleague working at a distant industrial site was working on code development, where full details were visible. He used the CB to start a discussion with that person, where he asked if there would be concerns about proprietary code being revealed outside that site. He then taught the person how to use blurring, where levels could be set to reveal coding activity without revealing contents.

Even with the privacy controls and feedback, there were some members of the lab who were not comfortable sharing their desktop using the Screen Sharing item. One member of the lab was concerned that if others did not see him working on his computer, they would think that he was "slacking off". In contrast, another member of the group who used the Screen Sharing item commented that one reason he liked it was because

it “lets me project a certain image of myself. I can use it to indicate I’m working, or pretend that I’m working”.

6.6 Distraction Issues

When many desktops were being shared on CB, people found it difficult to find the ones they wanted to see. Most people were only interested in a subset of the desktops being shared. Examples included task-oriented subsets that included only the desktops of people working on different aspects of a collaborative task (e.g., paper writing), or social subsets that included the desktops of close friends. This could be easily resolved using the Place feature in the Community Bar to create a more focused sub-group, but this practice had not yet been established by this larger group.

In spite of distraction concerns, we saw cases where an artifact on someone’s ‘unimportant’ screen caught the eye of another person. This sometimes resulted in a purely serendipitous and opportunistic conversation, which is one of the benefits of universal awareness.

7. CONCLUSION

Screen sharing was originally created to give collaborators the ability to do focused joint work across distances [6, 18]. Our initial experiences reveal that this was one of the ways people used the CB Screen Sharing item. Yet, these experiences also reveal the importance of screen sharing for artifact awareness, something not previously promoted or discussed in the product or research literature.

The majority of uses of the Screen Sharing item were to maintain awareness of what others were doing, and to glance at details of this activity through either the tooltip grande or the full view. This sometimes led to brief interactions outside of the Screen Sharing item, i.e., where people would use the Chat item to discuss things they could see. For work-oriented activities, these chats were often part of a broader longitudinal discussion of how the joint work was progressing as a whole. Discussions only sometimes proceeded to remote pointing, usually when focused interaction over the artifact was required. Full screen-sharing functionality (e.g., application sharing [18]) was not frequently requested, although we plan to add that capability in order to provide a more complete transition from awareness to full groupware.

These results thus point to a broader use of screen sharing, ranging from artifact awareness, to monitoring activities, to brief discussions, and sometimes to focused work. Future directions include deploying the system to other groups and observing how they adopt it for their own use.

Acknowledgements. This research was partially funded by Alberta Ingenuity, NSERC NECTAR Research Networks, and our major industrial partners: Smart Technologies Inc., and Microsoft.

Software. The Shared Screen media item and the Community Bar is available at: grouplab.cpsc.ucalgary.ca/cookbook/ and follow the links to Community Bar and to CB Media Items

8. REFERENCES

- [1] Bartle, R. *Designing Virtual Worlds*. New Riders, 2004.
- [2] Bellotti, V. and Bly, S., Walking Away from the Desktop Computer: Distributed Collaboration and Mobility in a Product Design Team, *Proc. ACM CSCW* 1996, 209-218.
- [3] Berry, L., Bartram, L., and Booth, K. S. Role-based control of shared application views. *Proc ACM UIST* 2005. 23-32.
- [4] Boyle, M. and Greenberg, S. The Language of Privacy: Learning from Video Media Space Analysis and Design. *ACM TOCHI*. 12 (2), June, 2005, 328-370.
- [5] Clark, H. *Using Language*. Cambridge Univ. Press, 1996.
- [6] Engelbart, D. and English, W. A Research Center for Augmenting Human Intellect. *AFIPS Proc Fall Joint Computer Conference*. 1968.
- [7] Finn, K., Sellen, A., Wilbur, S. (Eds) *Video-Mediated Communication*. LEA Press, 1997.
- [8] Greenberg, S. (1999). Designing Computers As Public Artifacts. *Int. J. Design Computing: (DCNet'99)*, U. Sydney.
- [9] Greenberg, S. and Rounding, M. The Notification Collage: Posting Information to Public and Personal Displays. *Proc ACM CHI*, 2001, 515-521.
- [10] Gutwin, C., Greenberg, S., Blum, R. and Dyck, J. Supporting Informal Collaboration in Shared Workspace Groupware. *HCI Report 2005-01*, U. Saskatchewan, Canada. 2005.
- [11] Heath, C. and Luff, P. Collaboration and Control: Crisis Management and Multimedia Technology in London Underground Line Control Rooms. *J. CSCW*, 1(1). 1992.
- [12] Kaasten, S. Greenberg, S. and Edwards, C. How People Recognize Previously Seen WWW Pages from Titles, URLs and Thumbnails. X. Faulkner, J. Finlay, F. Detienne (Eds) *People and Computers XVI* (Proc HCI BCS Conference Series), Springer Verlag, 2002 247-265.
- [13] Kraut, R., Egidio, C., and Galegher, J. Patterns of Contact and Communication in Scientific Research Collaboration. In *Intellectual Teamwork: Social & Technological Foundations of Cooperative Work*. LEA Press, 1990, 149-181.
- [14] McEwan, G. and Greenberg, S. Supporting Social Worlds with the Community Bar. *Proc ACM Group*, 2005.
- [15] McEwan, G., Greenberg, S., Rounding, M. and Boyle, M. Groupware Plug-ins: A Case Study of Extending Collaboration Functionality through Media Items. *Report 2006-822-15*, Dept Comp. Science, U. Calgary, Canada 2006
- [16] Nardi, B. A. *A Small Matter of Programming - Perspectives on End User Computing*. MIT Press. 1993.
- [17] Nardi, B., Whittaker, S. and Bradner, E. Interaction and Outeraction: Instant Messaging in Action. *Proc ACM CSCW*, 2000, 79-88.
- [18] Richardson, T., Stafford-Fraser, Q., Wood, K. and Hopper, A. Virtual Network Computing, *IEEE Internet Computing*, 2 (1), Jan/Feb 1998. 33-38.
- [19] Root, R.W. (1988). Design of a multimedia vehicle for social browsing. *Proc ACM CSCW*, 1988.
- [20] Rounding, M. *Informal Awareness and Casual Interaction with the Notification Collage*. MSc Thesis, Dept Comp. Science, U. Calgary, Canada, April 2004.
- [21] Tang, J. and Rua, M. Montage: providing teleproximity for distributed groups. *Proc ACM CHI*. 1994. 37-43.
- [22] Whittaker, S., Frolich, D., and Daly-Jones, O. Informal workplace communication: What is it like and how might we support it? *Proc ACM CSCW*, 1994, 131-138.
- [23] WOS Data Systems. *Timbuktu: The Next Best Thing to Being There*, User's Guide (Version 1). 1987.