

Artifact awareness through screen sharing for distributed groups

Kimberly Tee^a, Saul Greenberg^a, Carl Gutwin^b

^aUniversity of Calgary
2500 University Dr NW
Calgary, AB
Canada T2N 1N4
{tee, saul}@cs.ucalgary.ca
Telephone: +1 (403) 210-9499
Fax: +1 (403)284-4707

^bUniversity of Saskatchewan
57 Campus Drive
Saskatoon, SK
Canada S7N 5A9
gutwin@cs.usask.ca

Abstract

When co-located, people can see the artifacts that others are working on, which in turn enables casual interactions. To help distributed groups maintain mutual awareness of people's electronic work artifacts, we designed and implemented an awareness tool that leverages screen sharing methods. People see portions of others' screens in miniature, can selectively raise larger views of a screen to get more detail, and can engage in remote pointing. People balance awareness with privacy by using several privacy-protection strategies built into the system. An evaluation with two groups using this system shows that people use it to: maintain awareness of what others are doing, project a certain image of themselves, monitor progress, coordinate joint tasks, determine others' availability, and engage in serendipitous conversation and collaboration. While privacy was not a large concern for these groups, a theoretical analysis suggests that privacy risks may differ for other user communities.

Keywords

Artifact awareness, screen sharing, informal awareness, casual interaction, distributed groupware

1. Introduction

Previous studies have shown that *casual interaction* – the brief, unplanned meetings that commonly occur during the day between co-located people – is important for coordinating joint work, tracking progress of joint work, exchanging knowledge and information, and building relationships (Kraut et al., 1988; Whittaker et al., 1994). Casual interaction is made possible by *informal awareness*, the naturally gained understanding of who is around, what tasks they are performing, and whether or not they are available for conversation or collaboration (Kraut et al., 1988).

Informal awareness is easy to maintain in a co-located setting, such as when people inhabit a shared office space. Just by being in the same environment, people naturally accumulate background information about what is going on around them (Bly et al., 1993). People also do walkabouts, where they wander around a shared space just to see what others are up to (Bellotti and Bly, 1996). Because many awareness cues are available in a co-located environment, transitioning to casual interaction is typically effortless.

For distributed groups, however, initiating casual interaction is problematic: people do not see who is around, do not know if others are available for conversation, and lack the awareness cues that naturally lead to serendipitous interaction. Consequently, distributed collaborators must expend a relatively large amount of effort to explicitly coordinate interaction (e.g., by scheduling meetings), or do without this interaction altogether. Kraut et al. (1988) argue that much useful communication between workers in a knowledge-based environment is unplanned and would not occur if it had to be planned, suggesting that distributed groups are missing out on valuable interaction opportunities that naturally occur in co-located groups. This partially explains the explosion and success of low-effort awareness servers and casual interaction systems, such as the widespread adoption of instant messengers by diverse user groups (Nardi et al., 2000), or how email is often used for casual on-going conversations rather than as a formal messaging system. Many existing tools succeed because: (a) they let people know about the activities and thus the approximate availability of their colleagues, and (b) they make initiating conversation extremely easy.

Yet, an important component of informal awareness that is not handled by these awareness servers and instant messengers is *artifact awareness*, defined as one person's up-to-the-moment knowledge of the artifacts and tools that other people are using as they do their work. Artifact awareness is typically easy when collaborators are in a co-located environment 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 on a person's desk catches the eye of someone walking by, that person can stop and discuss that artifact with its owner. This awareness can be very important for exchanging knowledge and building social relationships. For example, consider designers who rely on easy visibility of other people's work. Design 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 (Buxton et al., 2000). Similarly, many command and control situations – such as air traffic control, subway routing, or shipboard navigation – rely on people being able to see each other's artifacts (Heath and Luff, 1992). Artifact visibility also plays a role in how people create common ground in conversations (Clark, 1996).

However, even in a co-located environment, artifact awareness can be difficult to maintain, particularly if the artifact is digital. Most digital artifacts only appear on a screen, which hinders awareness when

people's bodies shield their on-screen work from others or when social etiquette prevents others from taking a closer look at someone's display. Part of the etiquette problem is that the screen contains both semi-public information and highly personal information. A viewer cannot distinguish between these unless he or she is looking closely at the information, and by that time, it is too late to avoid looking at any private information. Another problem with maintaining awareness of digital artifacts is that with current computer windowing systems, artifacts are often overlapped or are not even shown on screen when another artifact is being used. Switching from using one digital artifact to another digital artifact takes little time and can be easily missed by others.

Artifact awareness is even more difficult to maintain in a distributed setting; despite 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. Whittaker et al. (1994) found in their study that a little over half of all casual interactions in an office involved document sharing, strongly suggesting that distributed groups would benefit from being able to more easily share artifacts. Yet, for distributed groups, there is still no real equivalent to the way co-located people can visually share their individual work and maintain artifact awareness. While there are many groupware systems (e.g., shared editors and webcast meeting tools) that do let a distributed group share artifacts, they work only after interaction is initiated. That is, they are intended for focused collaborative work rather than for artifact awareness. They also tend to be heavyweight to set up, e.g., they may involve many interaction steps to get going, which inhibits their use for casual interaction.

Perhaps the closest equivalent for distributed groups are screen sharing applications that let people explicitly share their computer screens, windows, or window fragments with each other. Screen sharing, first demonstrated in 1968 (Engelbart and English, 1968), has been used in practice for over thirty years in real-time distributed collaboration systems, e.g., Farallon Timbuktu (WOS Data Systems, 1987), VNC (Richardson et al., 1998), and NetMeeting [www.microsoft.com]. 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 (similar to 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). While useful, all these systems are designed for focused interaction rather than artifact awareness leading to casual interaction.

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 leverages *screen sharing* techniques 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 article¹ 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 results from end-user evaluations and from a theoretical analysis of privacy, and discuss the ways in which the idea of using screen sharing to support artifact awareness can benefit distributed work groups.

¹ Parts of this article have previously appeared in (Tee et al., 2006).

2. Related Work

2.1 Artifact Awareness in Everyday Coordination and Work

As mentioned, *artifact awareness* is 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 (Kraut et al., 1988; Whittaker et al., 1994):

- *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 public activity, even if it is not part of a joint task, can trigger interest in that activity (Greenberg, 1999).
- *Determining availability*: Knowledge of artifacts in use is yet another source of contextual information that helps people determine how busy others are and when they can be interrupted (Kraut et al., 1988; Whittaker et al., 1994).
- *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 et al. (1994) 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 (1993) found that people opportunistically collaborate over spreadsheets, CAD systems, and other documents over the course of a day.

2.2 Informal Awareness and Casual Interaction Systems

A number of groupware tools have been developed to support informal awareness and casual interaction within small communities of distributed collaborators. These tools purportedly help distributed groups overcome the disadvantages of being distance-separated by providing them with awareness cues and opportunities for informal communication not normally available in a distributed environment.

Instant messaging (IM) systems primarily support real-time text chat across the Internet, though newer systems offer richer communication channels such as internet telephony (VoIP), video, groupware applications, and information exchange via file transfer. IM systems are extremely popular both at home and in the workplace, and are being used by millions of people worldwide for social and work purposes (Isaacs et al., 2002). IM systems typically provide status indicators showing if a person on a contact list is logged on, busy, away, or off-line. Some IM systems also display additional information for each contact on the list, such as a personal message set by that person. This information, as well as the presence information, is valuable to people for creating and maintaining a sense of social connection to those on their lists (Nardi et al., 2000; Smale and Greenberg, 2005). Without even having to interact with them, people are able to get a sense of others, such as how they are feeling, what they are doing, or where they are. While the awareness information provided by IM is basic, it succeeds because it lets people easily establish communication with one another at opportune times. As a result of this, people use IM for a number of purposes, including: coordinating work tasks, asking quick questions, coordinating impromptu social meetings, and keeping in touch with friends and family (Isaacs et al., 2002). However, unless a contact on a list explicitly sets a personal message for others to view, there is no awareness of what that contact's current activities are.

In contrast to IM, which primarily supports casual interaction between personal contacts in mostly dyadic conversations, chat rooms are public places where all can see and post messages. Examples include:

text-based chat rooms such as Internet Relay Chat (IRC), where conversations often happen between people whose real world identities are unknown to each other; graphical chat rooms such as Comic Chat (Kurlander et al., 1996), where people can create avatars to represent themselves while they meet and participate in conversations with others; and virtual worlds, where people can view and manipulate visual artifacts that comprise that virtual world (Bartle, 2004). While there can be strong awareness of the group and/or virtual shared artifacts within chat rooms, it is uncommon for people to use them to have work-related conversations or to share artifacts with their co-workers.

Media spaces link offices and public spaces through networks of audio and video to provide rich awareness of people and their immediate surroundings (Bly et al., 1993). The resulting collection of “always-on” videos can be shown on a personal computer, on a dedicated television monitor (for earlier analogue systems), or even on a video wall placed in a common area. By seeing others through the media space, people get a sense of others’ presence and availability, their social interactions, and sometimes their activities. Media spaces are typically good for seeing other people; however, one problem reported by 90% of the participants in Fish et al.’s study (1992) in the use of the Cruiser media space (Root, 1988) was that it did not allow any way of sharing work artifacts. While Cruiser enabled casual conversation, it did not allow transitions into work-related talk or focused collaboration around work artifacts.

More recent media spaces (e.g., Dourish and Bly, 1992; Greenberg and Rounding, 2001; McEwan and Greenberg, 2005) ‘relax’ the notion of video by instead providing occasionally updated snapshots, sometimes mixed with other media. However, while these media spaces have been successful at building and maintaining the sense of “connection” or “community” between people at different locations (Dourish and Bly, 1992; Greenberg and Rounding, 2001), they still typically do not enable people to easily share dynamic, work-related artifacts such as reports, spreadsheets, code, or other documents. The artifacts such as photos and web pages that some media spaces allow people to share are usually explicitly posted for disseminating information of interest to the group, rather than to provide awareness of ongoing activities. While actual work artifacts can be captured in the video, in practice this rarely happens due to camera angle and resolution limitations.

2.3 Screen Sharing

One way of sharing work-related artifacts that has been used in real-time distributed collaboration is to share one’s screen, window, or window fragment with another person. This emulates 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 of groupware is often called ‘collaboration-transparent’, as the underlying application or screen being shared has no knowledge that multiple people are viewing and/or using it. Although 40 years have passed since its first appearance (Engelbart and English, 1968), little work has been done investigating the use of screen sharing for providing awareness. Instead, current screen sharing systems are designed primarily for focused interaction, rather than artifact awareness leading to casual interaction.

There are a few exceptions. First, SynchronEyes [www.smarttech.com] is a commercial system that lets one person view others’ desktops as thumbnails. However, it is designed for a quite different educational setting where a teacher monitors / controls a co-located class of students, and not as a peripheral awareness tool used by geographically-distant peers. Second, preliminary work by ourselves and/or our colleagues explored the idea of screen sharing for awareness. The Notification Collage Desktop media item (Rounding, 2004) let people post an intermittently updating desktop thumbnail that others could see and select if desired to activate a ‘full’ screen sharing session. Multi-VNC (Gutwin et al., 2005) later followed, in which multiple instances of a screen sharing application were run so that each member of a

group could see the other members' computer desktops. Multi-VNC and the Notification Collage Desktop media item were both proof of concept. We extend their idea, where we examine in detail how to provide awareness using screen sharing.

3. An Artifact Awareness Tool

We designed our artifact awareness tool to use screen sharing to provide intimate collaborators with awareness of each others' artifacts and individual work. Our system is an extension of Community Bar, an augmented media space developed in our lab (McEwan and Greenberg, 2005).

3.1 Community Bar Overview

Figure 1 shows a screen snapshot of Community Bar (CB) in use. As can be seen, CB is 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, similar to a chat room. *Sticky Notes* contain one person's text posting to the group, much like a physical note might. People use these items to broadcast information, ask questions, or move into conversation with others. The public nature of conversations on CB makes *lurking* – where people see or overhear conversations but do not participate – possible. It also enables others to easily and serendipitously join an existing conversation.

Public information such as web pages or photos can be posted to the group through the *Web item* and *Photo item*. Through these items, CB allows people to share some of their artifacts. However, though sharing websites and photos on CB takes little effort, users must still explicitly post each individual artifact that they would like to share. Also, because the Web item and the Photo item each only show one artifact at a time, the bar begins to get cluttered when people share multiple artifacts. CB currently lacks support for artifact awareness and artifact-centered serendipitous interactions.

Media items are presented at three levels of granularity to provide different levels of awareness information and interaction possibilities. The media item's *tile* view is always visible in the sidebar. If something of interest is seen in the tile, individuals can choose to explore and interact with that information in more detail 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*, which displays even more information and interaction capabilities. The easy transition from awareness to interaction is partly what makes CB successful (Romero et al., 2007).

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 available on the tile, tooltip grande, and 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, whereas all people see a Chat item in exactly the same way.

Fundamental to the philosophy of CB is the idea 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. Also, 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. Notifications of new information in a media item are done through a change of background colour in the tile view.




Figure 1. Community Bar.

Finally, CB is an open-ended system based on a plug-in architecture. Using an API and development environment, third-party developers can create and add their own custom media items to CB without recompilation (McEwan et al., 2006).

3.2 The Screen Sharing Item

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. 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* works².

3.2.1 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 labelled with her name (Figure 1, second 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 70 x 60 pixel thumbnail are sufficiently recognizable (Kaasten et al., 2002) so that others can tell if the poster is editing a document, browsing the web, preparing a presentation, etc. While actual content is hard to distinguish, visual landmarks such as photos and text formatting are discernable. For example, the first five thumbnails in Figure 2 show people visiting a web page, editing a Word document with highlighting 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 presentation. The last two thumbnails show particular regions of people's screens: part of a code window in a programming environment, and part of a document.


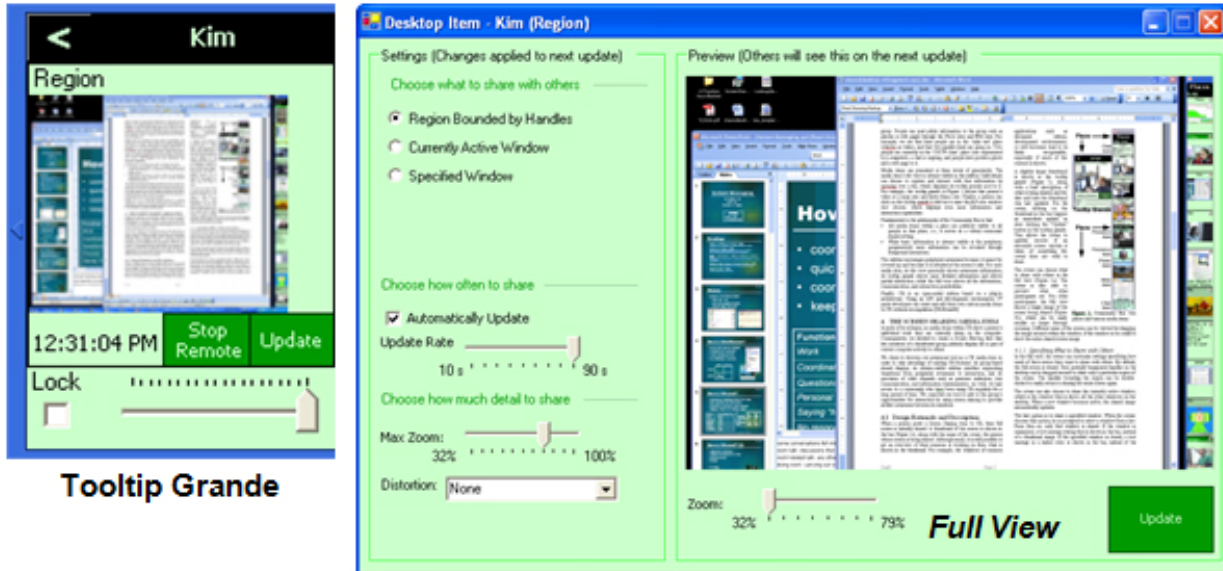


Figure 2. Thumbnails of people's screens.

² Implementation details are not discussed here; they can be found in (Tee, 2007).

Now reconsider the shared screen in the tile from Figure 1, shown again in Figure 3c. 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.



A. The owner's tooltip grande and full view of the Screen Sharing item.



B. The audience's tooltip grande and full view of the Screen Sharing item. C. Tile view, all people

Figure 3. Various views of the Screen Sharing item.

3.2.2 Tooltip Grande

The tooltip grande for both the owner and the audience shows a somewhat larger thumbnail (Figures 3a and 3b, 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. Above 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). Below the thumbnail is a timestamp indicating when the image was last updated. For example, the tooltip grande views of Kim's desktop in Figures 3a and 3b (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. 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 3c is a Screen Sharing item posted by Stephanie, where she is sharing a specific window ("CBC New Brunswick – Mozilla Firefox") 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 3a, left). The audience view of the tooltip grande does not have this button, but the audience has a button in their view titled "Remote Pointing", while the owner has its corollary "Stop Remote". This will be explained later.

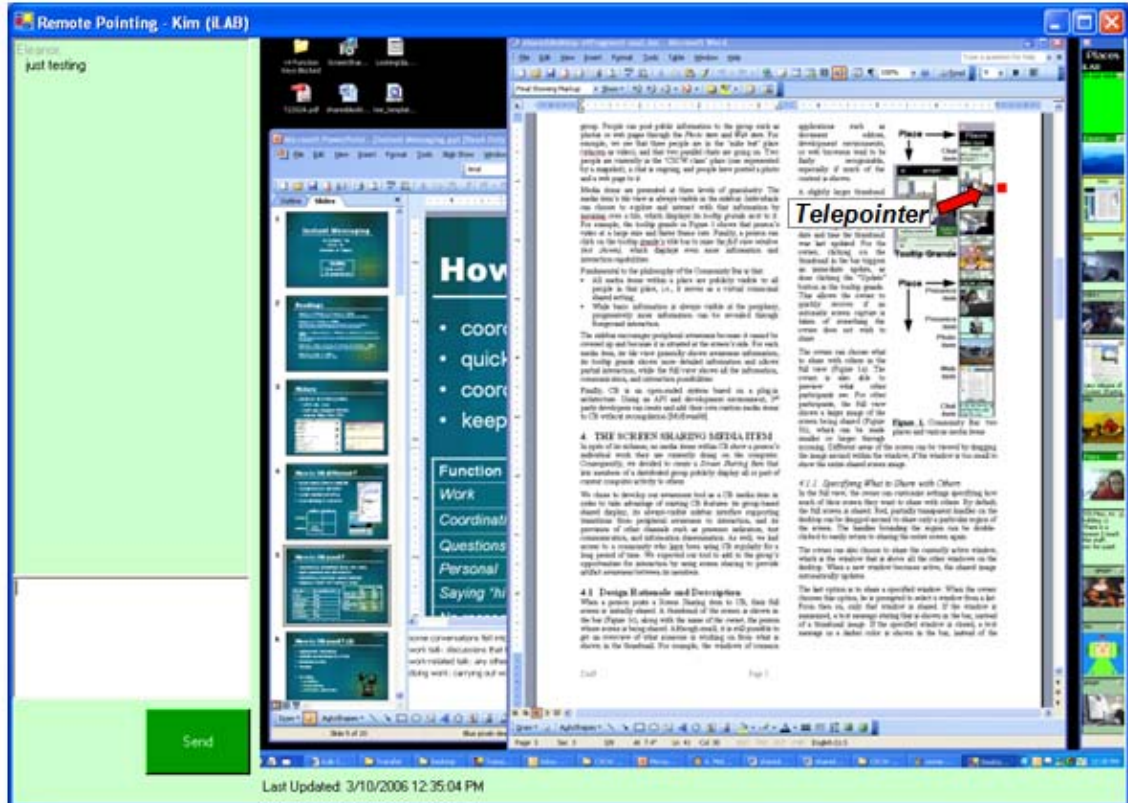
3.2.3 Full View

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 3a and 3b, right). 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 3a, 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 privacy controls are described in Section 4.

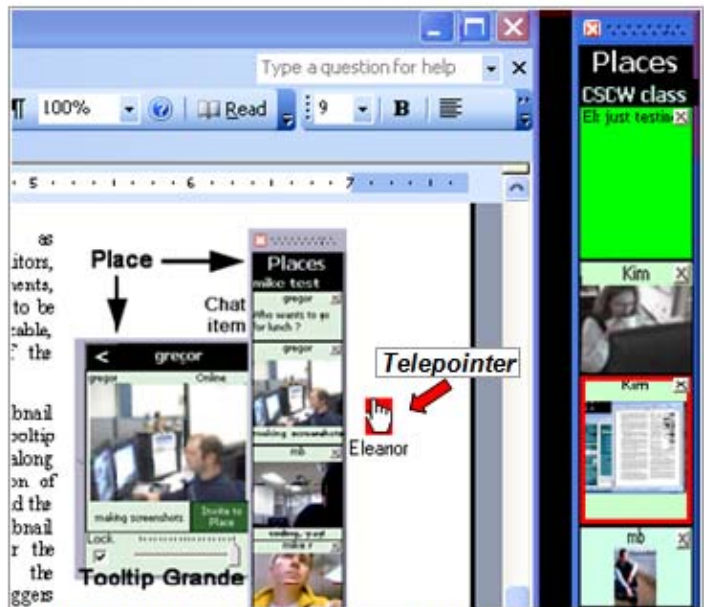
A zoom slider below the image lets the viewer zoom into the image as desired for greater detail. When the zoomed-in image does not fit within the window, the viewer can pan the image by directly selecting and dragging it with the mouse. As will shortly be discussed, for privacy reasons the permissible 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 3a and 3b, it can be seen that Kim has set the maximum zoom level to 79%. The audience member is looking at her screen at this maximum zoom level (Figure 3b, full view), while Kim is looking at it at a 32% zoom so that the entire region fits the full view's window (Figure 3a, full view). It can also be seen that 79% zoom of the true screen resolution produces a fairly legible image (Figure 3b, full view): subsection titles of the paper are easily visible, and the paper text can be read with some effort.

3.2.4 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 denies it, a short message is displayed to the audience member who requested remote pointing notifying them that permission was not granted.



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.

If the owner approves it, a remote pointing window (Figure 4a) appears on the screen of that audience member. This remote pointing window displays the shared screen image at the maximum allowable resolution. It also includes a full-sized chat box (Figure 4a, left) linked to a Chat item in the sidebar (Figure 4a and 4b, top right) so that the audience member can communicate with the owner in the same window as remote pointing, rather than having to switch between the remote pointing window and the CB sidebar or the Chat item's full view window.

The audience member can drag a small red telepointer around the shared screen 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. Remote pointing is currently limited to two participants; if another person tries to request remote pointing while the owner is already involved in a remote pointing session, that person is notified of this and asked to try again later.

Although remote pointing is not as powerful as systems that let people take turns interacting with the application such as VNC (Richardson et al., 1998) or Timbuktu (WOS Data Systems, 1987), remote pointing suffices for most situations. As Whittaker suggests from his observations of casual interactions in offices, "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" (Whittaker et al., 1994).

4. Privacy Controls

Privacy is, of course, a serious consideration in an always-on screen sharing system. For example, 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, note that privacy is not just a technical issue (Boyle and Greenberg, 2005). 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 need and desire to stay 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, people may want some control over what others can see. 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.

4.1 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. Thus, the 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 what to share in the full view (Figure 3a, full view): 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 (primary) screen is shared. Anything displayed in this region is captured: partial and overlapping windows, background wallpaper, dialog boxes, etc. Semi-

transparent red handles define the region being shared (Figure 5), and the owner can easily adjust the bounds of the 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 (primary) screen by double-clicking a handle. For people with two displays, right-double-clicking a handle will share the full secondary screen.




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

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. When the owner resumes working in the window, the thumbnail is displayed again.

4.2 Specifying Update Frequency

The 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 tile 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, this infrequent update should suffice for artifact awareness. These updates still inform others of basic activities while minimizing distraction and privacy concerns that might otherwise arise from real-time movement in the various views.

4.3 Specifying How Much Detail to Share

The 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 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 1280 x 1024 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. 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 easily added. 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. These effects let others roughly see what Kim is working on, while preserving her privacy because the image does not reveal legible detail.




Figure 6. Masking effects applied to the shared screen image.

4.4 Feedback of Image Capture

The 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. Five seconds prior to an auto-update, the owner's tile in the bar is outlined in yellow. If sharing a region, the red handles that bound the region turn yellow as well. Colours 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 (Figure 4b). This outline remains until the full view is closed. However, no identifying information is supplied as to who is looking at the full view; it acts only as an indication that at least one person has the full view open.

4.5 Communal Feedback

The 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.

5. Evaluation

We now report on the experiences that two different groups had while using the Screen Sharing item.

5.1 Internal Research Group

People inhabiting a common space naturally see the screens of their co-workers as they glance around, walk by, or are invited to take a closer look. However, using a shared screen in distributed groupware for artifact awareness is an unusual concept, and such a use will likely have to develop over time as part of a group's everyday practice. Recruiting a group to use Community Bar (CB) and the Screen Sharing item over a long period of time is difficult for pragmatic reasons. Consequently, we decided to first introduce the Screen Sharing item to members of our own research group, who had already been using CB on a daily basis for over a year. While this was perhaps a biased group, there were several good reasons for using them:

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

The group that was introduced to the Screen Sharing item included ten graduate students, faculty, research assistants, and former members of the research laboratory. 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. The 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 and on weekends. 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. Members of this group were comfortable sharing information about themselves with others on CB, and felt that they benefited from sharing.

The Screen Sharing item was distributed to CB users in our research group as an installation package downloadable from a web page [[grouplab.cpsc.ucalgary.ca/ cookbook](http://grouplab.cpsc.ucalgary.ca/cookbook)]. When run, the installation package added the Screen Sharing item to CB. Group members were then free to use or not use the Screen Sharing item as they wished. As this was not a formal study, no usage data was logged by the system. Instead, people were asked to email in descriptions of use and other comments as they used the system over a two week period, and several people were interviewed in depth to discuss details. As CB users, we remained logged into CB much of the time and collected notes whenever we saw people discussing shared artifacts or desktops in CB. The key experiences and reactions people had to the Screen Sharing item, based on both these user reports and our direct observations, are summarized in the next section.

5.1.1 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 people make an informed decision on whether to initiate contact with another 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."

Finally, 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.

5.1.2 Presence Awareness

Besides being used to provide artifact awareness, the Screen Sharing item was also 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 because 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 colour automatically darkened when the computer had been idle for five minutes or more.

5.1.3 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 serendipitous conversations would occasionally transition into remote pointing sessions, which were typically 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. In another example, a group member was working on an initial draft of a paper that her co-author hadn't seen yet. Her co-author noticed the document in the Screen Sharing item, but

the image was blurred, so he asked her to unblur it. They then went into a brief remote pointing session to discuss the title and abstract.

5.1.4 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 that they were the correct ones before she sent them. In a third case, one group member was helping another work on her poster. They used the Screen Sharing item and remote pointing to try out and discuss different layouts.

5.1.5 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 semi-public information (i.e., information that they were willing to share) and private information (i.e., information that they preferred not to share) 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)." This strategy of separating public and private information onto different displays is one example of the 'partitioning' that Grudin (2001) observed when studying how people use multiple monitors.

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]". This perception of digital artifacts as being private when in fact they are semi-public is partly a result of current workplace etiquette, which suggests that people should refrain from looking closely at another's display unless invited to do so. In CB, the act of posting a Screen Sharing item acts as a tacit invitation to look closer if interested, and so people may have the feeling that others are looking closely at what is being shared on their screen regardless of whether anyone actually is. In this regard, the feedback from the Screen Sharing item that indicated to people that someone else was looking at their desktop using the full view was only partially effective. While useful for making people more conscious of what they were sharing, people also wanted to be able to identify who was looking at their desktop without that person having to explicitly tell them. This feedback also 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 (see Section 6).

Second, some people chose to blur the screen image they shared so that text in windows would not be legible to others (Figure 6a). In fact, 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. Even in this case though, 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 particularly 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 colour 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 update was about to occur reminded him that the exam questions should not be publicly available. He then moved the document from the shared region to the private secondary screen before the exam questions could be seen by others.

There was 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' (Boyle and Greenberg, 2005) 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 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". These incidents are examples of some of the privacy maintenance issues discussed by Vaida et al. (2005).

5.1.6 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, such as 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 (Romero et al., 2007).

Contrary to what we expected, no one said that they found the Screen Sharing item or its auto-update warnings distracting. In fact, there were cases when 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 having universal awareness of things that are not initially of interest.

5.1.7 Summary: Internal Research Group's Experiences

Based on observations and discussions, the majority of uses of the Screen Sharing item by the research group 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, e.g., when 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) was not frequently requested.

The people who found the Screen Sharing item most useful were those working together on a joint paper or project. They were able to see which parts their collaborators were working on, and were able to discuss various things about the paper or project, sometimes using remote pointing. Those who weren't directly collaborating with other CB users mainly found the Screen Sharing item useful for determining whether or not someone could be interrupted at a particular point in time.

5.2 External Industry Group

Two months after the Screen Sharing item was deployed to our research group, it came to our attention that members of a commercial development team were also using CB and the Screen Sharing item, which had been introduced to them by a former member of the lab who had gone to work at the company. In this section, we report results from interviews done with several of the people in this industry group. This was again not a formal study; rather, we took the opportunity to get feedback from an outside group of users when the opportunity arose. In particular, we were interested in finding out whether this industry group's use of CB and the Screen Sharing item differed greatly from our own research group's use³.

The development team consisted of seventeen people located in the same building, though they were scattered across the floor and some members were on different floors. Of that group, an estimated (by participants) ten to fifteen people had tried CB, though the core group of daily users was from five to seven people. From this core group of users, we interviewed four volunteers. Three of them were developers, each from a different sub-group, and one was an internal user experience consultant. The internal user experience consultant had been part of our lab several years ago and had used CB before, but was not otherwise involved in the CB project other than as an end-user. Another core CB user, not included in the interviews, was also a former member of our lab, and was the one who suggested to the group that they try CB. Though he was also not directly connected to the CB project, he had been a heavy CB user in the lab (which he had more recently left) and he played a large role in encouraging the development team to try out CB.

Semi-structured interviews were conducted on-site at the company. Each interview consisted of two phases, and only one interview lasted for more than half an hour. The first phase focused on CB and the group's usage of it, as well as background information about the group. The second phase focused specifically on the Screen Sharing item. At the time of the interviews, the group had been using CB for

³ In particular, the company maintained a Bridgit [www.smarttech.com] server that anyone could connect to at any time for desktop conferencing or remote collaboration. While Bridgit has some features similar to CB, it is intended to be used for focused collaboration, not awareness or casual interaction.

about three weeks, and all our participants had been using CB for at least two weeks. While most members of the group had two monitors, only one person had a webcam. The four participants that we interviewed only connected to CB while they were in the office. Their key experiences and reactions to the Screen Sharing item are summarized in the next section.

5.2.1 Artifact Awareness

Two of our four participants shared their screens regularly, one shared his screen occasionally, and the fourth had tried it out once or twice but did not currently share his screen (because he felt that he didn't have anything to share at the moment). People typically shared their full screen or a region of their screen and they usually shared what they were working on, though one participant said that he often focused in on one tiny region of his screen such as his MSN Messenger display picture to share, as a joke. Though an unintended use of the system, this still sometimes led to conversation or banter about what he had shared.

Participants said that they could identify what they saw on others' desktops "surprisingly really well". All reported being able to identify activities such as coding and web surfing, and were able to recognize some programs that others were using just from seeing the outlines of the windows. Again, the added knowledge of what a person was doing helped group members determine whether or not that person was available for interruption. Having this availability information was the most common reason participants said that they liked seeing others' screens.

In contrast to our research group's experience with the Screen Sharing item, the participants from the industry group did not seem to use it to opportunistically coordinate or track progress of each others' work. This may have been because our participants were from different sub-groups of the project and so did not typically work closely together. Also, during the time that the group had been using CB, they were in-between projects and thus felt "no real big push" to get one particular thing done or to act together towards any one particular goal. Rather, they were mainly doing bug fixes and maintenance.

Because participants were from different sub-groups, they didn't usually hear about each others' projects. In combination with being able to overhear conversations that occurred in CB, being able to see what others were working on helped give people a better idea of what was going on with others' projects. This in turn led at least one participant to feel like he knew the other team members a bit better.

5.2.2 Presence Awareness

Because none of the participants interviewed (and only one person in the whole group) had a webcam, the Screen Sharing item was found useful for presence information within the commercial development group more so than within our own research group. The Screen Sharing item acted as a partial replacement for a webcam by showing when people were using their computers. It was only a partial replacement in terms of providing presence awareness because the Screen Sharing item might still show someone as 'away' when they were in fact present but not using their computer (e.g., they were reading at their desk).

However, the Screen Sharing item did indicate useful information about another's presence and availability that would not have been captured by a webcam. For example, a couple of the people who shared their screens on CB also had Macintosh computers that they regularly connected remotely to, e.g., through VNC (Richardson et al., 1998). Because they were sharing their screens, their Mac screens would be captured by the Screen Sharing item when they were connected, indicating to others on CB that email and IM messages to their PC would likely go unnoticed until they disconnected. In other words,

although they might still be available to people walking by or calling in, electronic messages would probably be unanswered while they were connected to their Mac.

5.2.3 Opportunistic Interactions

In the industry group, opportunistic interactions often seemed to be initiated by the person sharing their screen telling another on CB to go look at a shared artifact or region. This was in contrast to our research group's experience, where people tended to notice and ask about shared artifacts without additional encouragement from the person sharing. For example, in the industry group, one group member told a participant to check out a blog entry on his shared screen, and they moved into conversation and remote pointing about it. Another participant used the Screen Sharing item to opportunistically show a team member a bug he had found in what they were working on together. That same participant commented that there was another person that he was working closely with that was not on CB, but he wished that the person was, so that they could share their screens with each other.

Participants also liked how the Screen Sharing item enabled lightweight casual interaction, in which people were able to interact around shared artifacts while remaining at their own desks. For example, one participant felt that the primary benefit of using the Screen Sharing item was that he could easily show people things on his screen without having them come over to join him in person.

Few conversations seemed to transition into remote pointing; participants reported using remote pointing mainly to try it out.

5.2.4 Focused Collaboration

Members of this industry group did not use the Screen Sharing item for focused interactions much, partly because the company maintained a Bridgit server that anyone could connect to at any time. Though there was some overhead for creating or joining a desktop conferencing session, Bridgit allowed full remote control of another's desktop as well as had integrated VoIP. In fact, two participants said that they would like to have a link to Bridgit from the Screen Sharing item.

5.2.5 Privacy Issues

People in the industry group used somewhat similar privacy-protecting strategies as the people in our own research group. In particular, most members of the industry group also had two monitors and reported separating their public and private information onto different regions of their displays. Having two monitors seemed to alleviate many of the privacy concerns that people had; one participant commented, "I can always open up an IDE and make them think that I'm working on something. That's good for me because I have two monitors. If people have [only] one [monitor], probably is different [people would probably feel differently about screen sharing]."

Unlike the participants from our research group, none of the participants from the industry group reported regularly sharing a specific window or their currently active window. Also, none regularly distorted their shared screen image. One person typically shared his screen at the maximum (original) resolution, and another shared his screen at the original resolution about half the time. One participant mentioned that he would like a way of blocking out certain regions of his screen so that they couldn't be seen by others.

In most cases however, privacy did not seem to be a large concern for these participants. None reported having any concerns about screen sharing either before or after trying the Screen Sharing item. This may partly have been because of the current environment or group they worked with; when asked about whether he had any concerns about sharing his screen, one participant responded, "here [at this

company], not really. I guess I'd never thought about... if I switch away from my paper because I'm falling asleep and go check my personal email... people might be able to read it, but then again, what do I get in my personal email that's really all that private anyway." Another participant mentioned that while he only wanted to share his screen at certain times (he didn't like people "over [his] shoulder all the time"), as long as he could get his privacy when he needed it, he had no real concern with using the Screen Sharing item.

5.2.6 Distraction Issues

Again, no one seemed to find the Screen Sharing item or its auto-update warnings distracting. One participant did shrink others' Screen Sharing items when he wasn't interested in what they were doing (he was only interested in what they were doing when he wanted to talk to them).

People in this group seemed to be more careful about trying not to distract others and trying not to clutter up the bar, perhaps because they were working in an industry environment where productivity was important. Two participants mentioned that they only shared their screens when they had something to share, with one commenting that the reason he didn't share his screen all the time was because doing so takes up extra space in the bar. There is a trade-off between awareness and distraction; while sharing screens all the time can lead to opportunistic interactions triggered by shared artifacts, a cluttered bar can make it difficult to find shared screens and artifacts of interest. These issues may be specific to CB and the way CB displays information; alternative designs for screen-sharing awareness tools might address the concerns these group members had about 'frivolously' taking up others' screen space.

5.2.7 Summary: External Industry Group's Experiences

With these interviews, the primary goal was to investigate whether an industry group's use of the Screen Sharing item greatly differed from our own research group's use. While we were only able to interview four of the CB users on the commercial development team, it appears that many of their experiences with the Screen Sharing item were similar to the experiences of our in-house group. People in the industry group also used the Screen Sharing item to maintain awareness of what others were doing, help determine others' availability, project a certain image of themselves, and engage in lightweight casual interaction. Yet, there were also some differences in how the two groups used the Screen Sharing item. For example, there seemed to be fewer opportunistic interactions in the industry group, and the ones that occurred appeared to be more social in nature. We further discuss the similarities and differences we saw between the two groups in the next section.

5.3 Comparing the Research Group with the Industry Group

The internal research group and the external industry group used the Screen Sharing item for many purposes. We examine here the key similarities and differences we saw between the two groups, both in their group characteristics and in their experiences with the Screen Sharing item. We recognize that due to our methods of data collection (interviews and observations with ten participants vs. interviews with four), some of these similarities and differences may not necessarily be significant; however, we describe what we saw to inform further studies that may refine our results.

5.3.1 Possible Factors Affecting Use

There were several factors relating to group characteristics that may have affected how each group used the Screen Sharing item. For example, both the research group and the industry group were composed of people involved in research or development relating to computer science. However, during the deployment periods, the majority of people in the research group were co-authoring papers (working

closely with one or two others) while the majority of people in the industry group were (individually) working on bug fixes and product maintenance.

Also, members of both groups typically had two monitors each. However, while many members of the research group had a webcam, only one member of the industry group did. Considering that every participant in Romero et al.'s study (2007) reported that their primary motivation for using CB was because of the rich awareness they gained from being able to see webcam snapshots of people, the industry group likely had somewhat of a different experience using CB than the research group did. People in the industry group additionally had free access to a Bridgit server that their company maintained for them, where typical uses of Bridgit included desktop conferencing and remote collaboration. While the research group could have used similar systems such as the free application sharing system supported by MSN Messenger, conferencing and collaborating through screen sharing systems was not a part of their work culture.

Finally, the research group had been using CB for over a year, whereas the industry group had only been using CB for several weeks. Still, at the time of each study, the two groups had been using the Screen Sharing item for about the same amount of time.

5.3.2 Main Similarities

From the two groups' observed and reported experiences with the Screen Sharing item, several similarities emerged. First, screen sharing for awareness seemed to be most useful for certain subgroups of people, such as those working closely together on a joint task or project. By being able to see each others' screens, people were able to opportunistically coordinate as well as track the progress of their joint work. People were also able to easily check whether collaborators were interruptible for work-related conversation. As another example, social subgroups of friends also found the Screen Sharing item particularly useful; people would commonly use the Screen Sharing item to check if others were busy working or if they were free to take a break.

Participants from both groups also found the Screen Sharing item particularly valuable for lightweight casual interaction. They liked being able to check whether others were available without having to actually walk over to others' desks, and they liked being able to share artifacts and activities without having to ask others to come over. Using the Screen Sharing item, it was easy for people to tell others to look at their shared screen on CB and then move into conversation about what was being shared. For remote people, the Screen Sharing item enabled them to easily check availability or share artifacts when they normally wouldn't have been able to do so.

Finally, privacy did not seem to be that large a concern for people from either group. Having two monitors seemed to help alleviate much of people's concerns about sharing information since they could then separate public and private information onto different displays, sharing only one. Also, with more screen space, it was possible to 'pretend' to be working in the shared region while actually doing other activities in the non-shared regions. People in both groups recognized that they could project this image of themselves as being productive, and thus were not too concerned about sharing their screen. Further analysis of how the Screen Sharing item potentially affects people's privacy is given in Section 5.

5.3.3 Main Differences

There also appeared to be several differences in how the Screen Sharing item was used by the two groups. First, opportunistic interactions seemed to be initiated quite differently. In the industry group, opportunistic interactions often seemed to be initiated by the person sharing their screen telling another on CB to go look at a shared artifact or region. In contrast, people in the research group tended to notice

and ask about shared artifacts without prompting from the person sharing their screen. This may have been because people in the research group generally shared their screens all the time, whereas people in the industry group tended to share their screens only when they felt that they had something to share. However, Birnholtz et al. (2007) also observed in their studies of privacy in open-plan offices that opportunistic interactions did not happen as a result of what someone saw on another's screen; instead, interactions occurred when people walking by were invited to look at a person's screen, suggesting that other possible reasons for the different ways interactions were initiated were: 1) the etiquette followed by the two groups differed, or 2) the research group had developed a practice of use around CB that differed from their usual practice in the lab, whereas the industry group continued following the same social norms that existed in their workplace.

The industry group's desire to avoid cluttering up the bar was a little surprising considering that there were about half the number of CB users in that group than in the research group. Also, from the interviews, it did not sound like the industry group shared an excessive amount of websites or photos that would take up a lot of space in the bar. Still, it appears that this carefulness to not take up extra space in the bar partly contributed to why the industry group shared their screens less than the participants in the research group did. In turn, sharing their screens less may have contributed to there being fewer conversations transitioning to remote pointing and fewer focused collaborations in the industry group. Another factor may have been the availability of Bridgit to the industry group for desktop conferencing or remote collaboration whenever they wished to use it; full screen-sharing functionality for the Screen Sharing item was more frequently requested from the industry group than from our own research group.

Finally, there seemed to be a difference in how often the people in the two groups collaborated together. There was less focused collaboration through the Screen Sharing item reported in the industry group than in the research group, as well as fewer opportunistic interactions involving coordination or tracking progress of joint work. This may have been a result of the industry group members not working closely together on joint projects, and consequently not needing to do much collaborative work or coordination together overall. Alternatively, the availability of Bridgit (with its enhanced screen sharing features) may have made it easier for the industry group members to use that rather than the Screen Sharing item when they needed to check progress or do focused collaboration. However, despite not using the Screen Sharing item much for focused collaboration, people in the industry group still found the Screen Sharing item useful for awareness and for creating opportunistic interactions. It was also found valuable for determining others' availability and sharing artifacts and computer activities with others. This shows that there is a difference between shared screens for awareness and shared screens for focused collaboration, and highlights the importance of screen sharing for artifact awareness, something not previously promoted or discussed in the product or research literature.

6. Theoretical Analysis of Privacy

In the previous section, we reported on how two different groups used screen sharing for artifact awareness through the Screen Sharing item. While privacy seemed to be less of a concern for both groups than we expected, their use of the system may not predict other communities' concerns about privacy. In this section, we report on a theoretical analysis of how people's privacy could be affected by the Screen Sharing item. As part of this process, we articulate possible privacy-related problems and examine the assumptions made during the design of the system.

6.1 Method

Privacy is a generic term that means different things to different people. To enable the unambiguous description and discussion of privacy-related issues in media spaces, Boyle and Greenberg (2005)

created a framework that collects various terms and concepts related to privacy from the literature. Their framework broadly divides privacy into the three “control modalities” of *confidentiality*, *autonomy*, and *solitude*. Confidentiality relates to a person’s control over what information others can access about him/her; autonomy relates to how a person chooses to present him/herself; and solitude relates to how a person controls interpersonal interactions.

In their framework, Boyle and Greenberg additionally consider the *mechanics of privacy* (low-level details and properties of privacy) and particular issues relating specifically to *computers and privacy* (e.g., security or user interface concerns). Boyle then defines a process where people can use the framework to systematically analyze particular systems (or situations or techniques, etc.). His method of privacy analysis consists of the following four steps (Boyle, 2005): 1) *partition*, where the analyst determines the primary focus of the analysis and distinguishes between main and secondary effects; 2) *describe*, where the analyst uses Boyle and Greenberg’s framework and vocabulary terms to unambiguously describe each aspect of the system; 3) *reveal*, where the analyst identifies assumptions that have been made with respect to the system; and 4) *summarize*, where the analyst summarizes the findings from the previous two steps, including an overview of “merits and demerits” of the system as well as the conditions in which these merits and demerits exist. The reader is referred to (Boyle and Greenberg, 2005) and (Boyle, 2005) for greater detail; in the remainder of this section, we report on the results from our theoretical analysis.

6.2 Results

Boyle’s framework mainly considers privacy from the perspective of how it affects *confidentiality*, *autonomy*, and *solitude*. Within this context, the function of the Screen Sharing item is to share the contents of a person’s screen with other group members. Thus, the Screen Sharing item primarily affects confidentiality, the component of privacy that relates to a person’s control over what information others can access about him/her. For example, much information can potentially be revealed about a person from the contents of their shared screen, such as what they are working on, who they are instant messaging with, or what movie they will be seeing later that evening. Consequently, we primarily examine the Screen Sharing item’s effect on confidentiality in this analysis, as this attribute presents the greatest risk to one’s privacy.

However, it is important to recognize that the Screen Sharing item also secondarily affects autonomy and solitude. Autonomy is the component of privacy that relates to how a person chooses to present him/herself. The “observable manifestations of the self” (Boyle and Greenberg, 2005) such as appearance, identity, and impression that a person uses to present him/herself to others can be affected by the information accessible about that person. For example, seeing someone’s credit card or bank account balance on their shared screen may influence others’ impressions of that person. Solitude is the component of privacy that relates to how a person controls interpersonal interactions. The Screen Sharing item can affect a person’s solitude, for example by showing others that the person is not busy working on their computer, and thus potentially available for interruption. In the following sections, we also consider the Screen Sharing item’s effects on autonomy and solitude, although not to the same depth as confidentiality.

6.2.1 Main Effect: Confidentiality

As a tool for providing artifact awareness through screen sharing, the Screen Sharing item strongly affects confidentiality. It shares what a person is doing on their computer through the *visual information channel*, obtaining this information by *sampling*, i.e., by periodically taking a screenshot of the person’s screen to distribute. From the shared screen image, much information can potentially be seen or inferred; some examples are shown in Table 1. The information being shared can be *sensitive* (e.g., personal

Table 1. Information about a person that can be revealed through the Screen Sharing item.

Information Topic	Examples
<i>Information about the self</i>	a calendar or schedule; what a person is working on or looking at
<i>Personally identifying information</i>	an IM window with the person's name; a unique desktop background; the Screen Sharing item also identifies which person's screen it is sharing
<i>Activities</i>	open application windows can indicate that a person is coding, surfing the web, writing a report, communication through IM, etc.
<i>Whereabouts</i>	windows moving around, etc. indicate a person is at their computer; calendar/schedule information shows where the person is or will be; the Screen Sharing item darkens if a computer becomes idle for 5 minutes
<i>Encounters</i>	text in a group IM; videos of participants in a conferencing call
<i>Utterances</i>	text in IM or email messages
<i>Actions</i>	over time, typing, scrolling of windows, opening or closing of applications
<i>Relationships</i>	names in an IM contact list or email inbox

communications or financial records) or not (e.g., publicly available information). While the expectation is that work-related artifacts or activities will have low sensitivity with respect to other members of the same group, it is possible that a person's work may be highly sensitive if that person is working on a 'secret' project, if viewers are not normally a part of the group, or if a viewer is in a position of authority.

The Screen Sharing item offers a number of controls for mitigating confidentiality concerns; these controls alter the *fidelity* of the information being shared so that some details are obscured. Using the privacy controls, people can change the *capture* region of the screen to include or exclude certain windows or areas. They can change the zoom level / resolution of the shared screen image to make text legible or illegible. They can change the frequency at which the shared screen image is being updated to change the timeliness of the information being shared. They can also distort the shared screen image through blurring or pixelization in an *edit* operation on the image before it is transmitted to others. These privacy controls were described in full detail in Section 3.4. However, these brief descriptions illustrate that the privacy controls in the Screen Sharing item focus specifically on changing the *precision* of the information being shared, rather than the *accuracy*; the shared screen image always shows (with differing amounts of detail) exactly what is on a person's screen in the region being captured.

Still, even though the accuracy of the shared information does not change, *disinformation* is not hard and *misinformation* is possible; for example, the shared screen image could intentionally or unintentionally show that a person is working on writing a paper when in reality they are surfing the web in an area of the screen not being shared. Thus, there is still some *plausible deniability* about what a person is doing on their computer. However, it is possible for the Screen Sharing item to reduce the plausible deniability of IM or email. For example, if others can see a person working on their screen, that person may no longer be able to pretend that they are 'away' from their computer, or if others can see a person read a message or an email, that person can no longer pretend that they haven't received it or seen it yet.

The information shared by the Screen Sharing item typically has low *persistence*; on auto-update, a new shared screen image is captured and replaces the previous image at least once every 90 seconds. When a user logs off of Community Bar (CB), their shared screen image is automatically removed from CB – the images are not stored anywhere for *archival* purposes. However, it is possible for someone to make a persistent copy of a shared screen image by taking a screenshot of it and saving it to file. It is also possible for someone to hack into the underlying distributed data structure in order to automate screen recording. There is no support in the Screen Sharing item for tracking any *use, misuse, or misappropriation* of the information being shared. There is also nothing to prevent the *scrutiny* or

surreptitious surveillance of what one is doing on one's computer, particularly by one's superiors (a common concern). Still, the Screen Sharing item does provide some feedback about what others can see and whether or not others are currently looking closely at one's screen, as well as ways to reduce the amount of information others can gain from a shared screen image (at the cost of also reducing awareness).

6.2.2 Secondary Effects: Autonomy and Solitude

The Screen Sharing item additionally has secondary effects on autonomy and solitude through their relationships to confidentiality. Autonomy is related to confidentiality in that people can present themselves in a particular way by controlling what information about themselves others can access. Thus, the information that is shared through the Screen Sharing item about what a person is doing on their computer can influence others' *impressions* of that person. In particular, if a person works from a remote location where one of the only information channels about them is their shared screen, they may be judged by others primarily based on what their shared screen shows. This can become a problem when *back-stage performances* mistakenly occur on the *front-stage*, such as when someone working from home forgets that they are sharing their screen and starts looking at videos that are inappropriate for a work environment, though acceptable for watching at home. Privacy *harms* such as these that relate to autonomy can be *aesthetic* (e.g., a person often seen watching videos online instead of working being thought of as unproductive) or *strategic* (e.g., that same person being passed over for a promotion because they are perceived as being unproductive).

The Screen Sharing item aims to prevent mistaken back-stage performances from occurring on the front-stage by providing appropriate warning or reminder before a screen capture is taken. Also, since the expectation is that the Screen Sharing item will be used by work peers, it is assumed that *trust* has already been established between the group members who will be sharing their screens with one another and that group members will work to maintain that trust as well as help others preserve their privacy if needed, e.g., audience members can use the Chat item in CB to warn the person sharing their screen about things that the person may not want to be sharing.

Like autonomy, solitude is also related to confidentiality; people can somewhat control their interactions with others by controlling what information about themselves others can see. For example, the Screen Sharing item shows what a person is doing on their computer. This information can be used to determine that person's *availability* and also sometimes that person's *accessibility*, which can help others determine whether that person can be interrupted. Part of how people control their interactions also involves how they focus their attention and how they react to distraction. The Screen Sharing item typically shares information about others' screens through its tile in the bar, which sits on the *periphery* of the user's *attention*. If more information is desired, the user can see additional detail in the tooltip grande or full view, or can transition into *focused* interaction by engaging in a remote pointing session with another user. The information others are sharing may not always be of direct *relevance* to one's work. However, the idea is that valuable interactions can still be triggered by 'unimportant' artifacts on someone's screen.

6.2.3 Other Effects: Mechanics of Privacy and Computers & Privacy

Related to the mechanics of privacy, the Screen Sharing item affects many of the boundaries across which privacy is regulated. It allows regulation of the *disclosure boundary* by offering controls that people can use to reveal or obscure information on their screen. It enables people to transcend *spatial boundaries* by helping them share their individual work with others across distance. Because of the distance however, regulation of the *identity boundary* is more complicated, since the Screen Sharing item makes it possible for problems like back-stage performances on the front stage to occur.

Table 2. Potential privacy problems that might occur through use of the Screen Sharing item.

Privacy Problem	Common Cause	Example
<i>Inadvertent privacy infraction</i>	People forget that their screen is being shared.	Jane forgets that she is sharing her screen when she starts working with confidential data, and inadvertently shares it.
<i>Apprehension</i>	People do not want to make a bad impression in front of others.	John does not want to share his screen because he is afraid others might see him reading comics online during the day and think he doesn't work very hard.
<i>Resentment</i>	People resent a loss of control over their own privacy.	Jane dislikes being required to share her screen so that her boss can monitor her work, an implied lack of trust.
<i>Decontextualisation</i>	People share their screen or a region of their screen without its surrounding context.	John is sharing his web browser, which currently contains comics. Others may think he is taking a break, but he is actually creating a presentation in another window, and looking at comics for graphics to include.
<i>Disembodiment</i>	People share their screens from different locations / environments than their distant colleagues.	Jane, sharing her screen from home, forgets that others on CB are working from the office, and she starts looking at some photos taken at a wild party last weekend, photos inappropriate for a work setting.
<i>Role conflict</i>	People play different roles in different social worlds and sometimes previously non-overlapping social worlds "collide".	
<i>Dissociation</i>	The Screen Sharing item has no explicit way to identify which audience members are looking at a shared screen.	John notices that someone is looking at his shared screen in detail in their full view, but he does not know which of the ten people currently on CB that person is.
<i>Misappropriation</i>	People can be competitive and might take any opportunity to advance their career.	John sees on Jane's shared screen that she has come up with a solution to a problem the team has been working on; he quickly puts together an email to their boss in which he takes credit for the idea.
<i>Misuse</i>	People might take any opportunity to make personal or financial gain.	Jane sees non-public information on John's shared screen and based on that, sells some of her company stock.
<i>Identity theft</i>	People might take any opportunity to make personal or financial gain.	John sees Jane shopping online and copies down her credit card number and expiry date for future use.
<i>Impersonation</i>	People might act maliciously towards others they dislike or might take any opportunity to make personal or financial gain.	Jane has some of John's personal information from previously seeing it on his shared screen, and cancels his home internet service as a prank.

Also, because the Screen Sharing item shows what a person is currently doing on their computer, the *possibility* for various kinds of privacy *violations* of varying *probability* and *severity* exists. One example of a low probability, high severity privacy violation is a *deliberate abuse*, where a staff member from a rival company manages to intercept the shared screen data being sent to or from the shared dictionary in order to gain information for competitive advantage. More examples of possible privacy *problems* are

described in Table 2. While some of these problems are considered in the current version, others are left for future work to address.

To manage these potential privacy problems, people can use a number of strategies. Through *self-appropriation* or self-scrutiny, people monitor what they are sharing and manage the impression that they give to others. Part of self-appropriation can be the use of *disinformation*, for example to make it seem to others that they are working when they are not. However, self-appropriation can be cognitively draining as it takes effort to continually monitor what is being shared and to decide what is appropriate for being shared (Boyle and Greenberg, 2005). The *risk/reward trade-off*, i.e., the balance between privacy and awareness, can be managed through both technological and social controls. Technological controls include the privacy controls that were explicitly built into the Screen Sharing item, which people can use to reduce the amount of information they are sharing at the cost of also reducing awareness. Social controls include *self-policing* by the group. However, note that the Screen Sharing item does not have any *access control* or *authorisation*; anyone who joins the group in CB can see all the screens being shared. The only *content control* the Screen Sharing item provides is the *distortion filtration* of the shared screen images.

The Screen Sharing item does provide several types of *environmental support* for privacy; for example, the Screen Sharing item supports *reflexive interpretability of action* by providing similar tile and tooltip grande views to the owner as to the audience, as well as by showing previews in the owner's full view of what others will see if the owner changes various privacy settings. *Constraints* on what information is being shared can be placed on the Screen Sharing item – people can make a *choice* about how much information they want to show others by configuring the various privacy control settings or even by deciding not to share their screen at all. *Cues* on what information is being shared are provided through *feedback* of what one is sharing and when others are looking at one's screen. *Refuge* can be somewhat gained by putting the Screen Sharing item on manual updates. However, there is no *feedthrough* or natural *transitions* between different levels of privacy, and *reciprocity* is not enforced, which may lead to *risk/reward disparity* since a person revealing more of their screen or more details of their screen (thus increasing the risk for a privacy violation) may not receive proportional benefit / reward since others may not have to share the same level of detail. *Situated action* is also not supported directly by the Screen Sharing item, although over time, people can build up information about the context of someone's work by seeing their screens or through interaction with them.

6.2.4 Design Assumptions

The analysis revealed several assumptions made during the design process about how the Screen Sharing item will be used by a community. By examining these assumptions, we become aware of situations in which privacy problems may be more likely to occur.

- *Sharing with intimate collaborators*: Because the expectation is that groups using the Screen Sharing item will be comprised of people who trust and collaborate with each other, no explicit precautions to prevent privacy problems such as deliberate abuses have been taken. The Screen Sharing item would not work in a competitive environment, except perhaps as a source of disinformation; in fact, it has not been designed for a competitive environment at all.
- *Sharing voluntarily*: Another expectation is that people will be motivated to share parts of their screen useful for others to see. If people are forced to use the Screen Sharing item involuntarily (e.g., through peer pressure or because upper management wants to monitor them), they might potentially share parts of their screen that are not useful to others (e.g., an empty area of their computer desktop), or they might use the privacy controls to obscure information or disinform others, thus

Table 3. Privacy problems seen in the informal evaluation with two groups.

Privacy Problem	Example	Cause
<i>Inadvertent privacy infraction</i>	A participant who was creating exam questions almost accidentally shared them.	The participant forgot that their screen was being shared.
<i>Apprehension</i>	A CB group member did not want to share their screen, though they remained part of the CB community during the study.	The group member did not want to make a bad impression in front of others.
<i>Dissociation</i>	Several participants commented that they would have liked to know the identities of the people looking at their screen in the full view, not just that someone was looking.	The Screen Sharing item currently does not identify which audience members are looking at a shared screen.
<i>Role conflict</i>	A participant working at an industrial site on product development initially shared confidential code at full clarity.	The participant was acting both as a company employee and as a CB member, though the other CB members weren't employees of the same company.

- reducing the effectiveness of the system.
- *Primary privacy concern – inadvertent privacy violations:* In a community of intimate collaborators, inadvertent privacy violations seem more likely to occur than deliberate abuses. Thus, the privacy controls currently built into the Screen Sharing item seek primarily to prevent inadvertent violations such as people accidentally sharing information they meant to keep private.
- *Primary type of sensitive information – high-precision details:* Finally, the Screen Sharing item uses image-manipulation techniques such as blur filtration to obscure high-precision details (such as text from an email or instant messaging window) in a shared screen. However, these details may not always be what is sensitive in a shared screen image; for example, if a person were playing a game or looking at inappropriate videos or pictures at work, this may still be obvious from the overall colours and shapes seen in a shared screen image, even if the image were distorted. In fact, blur filtration has been shown to fail to preserve privacy in some high-risk situations in a home media space (Neustaedter et al., 2006), such as when people are caught by the camera in a state of undress. In other words, blur filtration can work well in cases where the text is the sensitive part of a shared screen, but can fail to preserve privacy when an image or a visual symbol is the part of the shared screen content that is sensitive.

6.3 Revisiting Privacy as Seen in the Informal Evaluations

During the informal evaluations, privacy did not seem to be a large concern for the two groups that used the Screen Sharing item. Of the privacy problems in Table 2, only instances of inadvertent privacy infractions, apprehension, dissociation, and role conflict were reported. These instances are recapped in Table 3, and were the most serious privacy problems encountered by participants (that they were willing to report). Deliberate privacy abuses did not seem to occur, likely because both groups consisted of intimate collaborators who trusted and sometimes worked directly with one another. For example, if the Screen Sharing item were used by employees who were not intimate collaborators but competitors, there would likely be more deliberate abuses of privacy seen.

As new users of the Screen Sharing item, we also suspect that participants were on their 'best' behaviour, and were careful with what they were sharing. Additionally, participation was voluntary, so people who were apprehensive about sharing their screens simply did not share them.

Few participants shared their screens from home when they were not doing work. High-risk privacy situations are more likely to occur in the home, as the home is a place of refuge where privacy is normally assured (Boyle and Greenberg, 2005); thus, by not sharing their screens from home much, people may have been able to avoid many inadvertent privacy infractions. Additionally, the common strategy of blurring shared screen images seemed to work well, as the information considered sensitive was mostly text (e.g., communications, exam questions, code) and not visuals (e.g., pictures, videos). Participants were also only concerned with keeping text illegible – they did not mind others knowing what they were doing.

Finally, participants' privacy strategies included taking advantage of the physical environment and hardware setup, in addition to using the privacy controls built into the Screen Sharing item. For example, one common strategy was to separate semi-public information and private information onto different displays. For groups with only one display, inadvertent privacy violations would likely increase, as there would be less display space available for disinformation.

In summary, note that the assumptions discussed in Section 5.2.4 held true for both groups; use of the Screen Sharing item might change for other groups with differing technology or group dynamics.

7. Future Work

The work in this article is the beginning, rather than the end, of a long-term research program. While our work suggests what could be, there is room for improvement in the system design, in our understanding of what people really want in terms of artifact sharing and awareness, and in evaluating the use and cultural adoption of such systems.

One future direction for this research could include improving on the current approach of using screen sharing in the Community Bar to provide artifact awareness to groups of intimate collaborators. There are several obvious ways that the current design could be improved: first, the system functionality could be expanded to address all aspects of the communications life cycle. For example, the Screen Sharing item could be linked to a full screen sharing system, in order to provide a more complete transition from awareness to full groupware. Second, the system needs to provide better information. One of the most common requests from participants was a way to identify which audience member(s) were looking closely at one's shared screen. In the real world, we see others approach, lean into our workspace, and glance at our artifacts. It is easy to tell who they are, how closely they are looking, where they are looking, and so on. This is not supported well in the current system, in which the Screen Sharing item only indicates that someone is looking closely at a shared screen and not who is looking. One possible approach is to supply additional information, e.g., an equivalent of the glance feature in Montage that shows that people are about to look in and that identifies them (Tang and Rua, 1994). Another approach used in the OpenMessenger system is to use different levels of notifications depending on how much detail others are looking at (Birnholtz et al., 2008). Finally, the omissions revealed by applying Boyle's framework to the Screen Sharing item suggest several additional ways for supporting users' privacy, such as adding controls for distorting the accuracy of information being shared, or adding natural transitions between different levels of privacy (for example, the Screen Sharing item could detect when a person is sharing an email or IM window and automatically blur the shared screen image).

Other possible future directions include exploring alternate ways of using screen sharing for providing artifact awareness, and further evaluation of how such systems can be adopted for use by different groups. Sharing screens within the Community Bar is only one way of providing artifact awareness through screen sharing; different ways of presenting shared screens to people may eliminate some privacy and space-usage concerns, while still providing enough information for people to maintain

awareness of what others doing. Also, longitudinal studies are critical, as this technology falls under what are called 'socio-technical systems' – their acceptance and use is as much about the culture that develops around it as it is about the features that the software provides. Screen sharing of one's desktop is still a strange concept, even though the virtual desktop superficially resembles one's physical desktop and how it is seen by others in an open office. The culture of use that develops around this technology could lead to outright rejection, outright acceptance, or (most likely) something between the two. The expectation is that people will find situations where tools such as these are valuable, and adapt their work behaviours around them. It is this adaption that is extremely interesting but difficult to probe in short-term studies.

8. Conclusion

In this article, we have described the design, implementation, and evaluation of an awareness tool that uses screen sharing within the Community Bar to support artifact awareness between intimate collaborators. With the awareness tool, 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. People balance awareness with privacy by using several privacy-protection strategies built into the system.

Screen sharing was originally created to give collaborators the ability to do focused, joint work across distance; the initial experiences people had reveal that this was one of the ways in which the screen-sharing awareness tool was used. Yet, these experiences also reveal the importance of screen sharing for artifact awareness. People used the screen-sharing awareness tool to maintain awareness of what others were doing, to influence others' impressions of them, to monitor progress and coordinate joint tasks, to help determine when another person could be interrupted, and to engage in serendipitous conversation and collaboration.

Artifact awareness is an important component of informal awareness that has not been well-supported in existing informal awareness and casual interaction systems. We hope that the research described in this article lays a foundation that will motivate others interested in providing informal awareness and casual interaction to groups to also include support for artifact awareness. The research described in this article is a starting point from which they can base their designs and intellectual investigations.

Acknowledgements

This research was partially funded by Alberta Ingenuity, NSERC, and our industrial partners: SMART Technologies and Microsoft.

References

- Bartle, R., 2004. Designing virtual worlds: New Riders.
- Bellotti, V. and Bly, S., 1996. Walking away from the desktop computer: Distributed collaboration and mobility in a product design team. Proc. ACM CSCW, 209-218.
- Berry, L., Bartram, L., and Booth, K., 2005. Role-based control of shared application views. Proc. ACM UIST, 25-32.
- Birnholtz, J. P., Gutwin, C., Hawkey, K., 2007. Privacy in the open: How attention mediates awareness and privacy in open-plan offices. Proc. ACM GROUP, 51-60.
- Birnholtz, J. P., Gutwin, C., Ramos, G., Watson, M., 2008. OpenMessenger: Gradual initiation of interaction for distributed workgroups. Proc. ACM CHI, 1661-1664.

- Bly, S. A., Harrison, S. R., and Irwin, S., 1993. Media spaces: Bringing people together in a video, audio, and computing environment. *Communications of the ACM*, 36(1), ACM Press, 28-46.
- Boyle, M., 2005. Privacy in media spaces. PhD Thesis, University of Calgary, Calgary, Alberta, Canada.
- Boyle, M. and Greenberg, S., 2005. The language of privacy: Learning from video media space analysis and design. *ACM Transactions on Computer-Human Interaction*, 12(2), ACM Press, 328-370.
- Buxton, W., Fitzmaurice, G., Balakrishnan, R., and Kurtenbach, G., 2000. Large displays in automotive design. *Computer Graphics and Applications*, 20(4), IEEE Computer Society, 68-75.
- Clark, H., 1996. *Using language*. Cambridge University Press.
- Dourish, P. and Bly, S., 1992. Portholes: Supporting awareness in a distributed work group. *Proc. ACM CHI*, 541-547.
- Engelbart, D. and English, W., 1968. A research center for augmenting human intellect. *Proc. AFIPS Fall Joint Computer Conference*, 395-410.
- Fish, R. S., Kraut, R. E., Root, R. W., and Rice, R. E., 1992. Evaluating video as a technology for informal communication. *Proc. ACM CHI*, 37-48.
- Frederick, R., 1994. Experiences with real-time software video compression. *Sixth International Workshop on Packet Video*. Portland, Oregon.
- Greenberg, S., 1999. Designing computers as public artifacts. *International Journal of Design Computing: Special Issue on Design Computing on the Net*, University of Sydney, Nov. 30 – Dec. 33.
- Greenberg, S. and Rounding, M., 2001. The Notification Collage: Posting information to public and personal displays. *Proc. ACM CHI*, 514-521.
- Grudin, J., 1994. Groupware and social dynamics: Eight challenges for developers. *Communications of the ACM*, 37(1), ACM Press, 92-105.
- Grudin, J., 2001. Partitioning digital worlds: Focal and peripheral awareness in multiple monitor use. *Proc. ACM CHI*, 458-465.
- Gutwin, C., Greenberg, S., Blum, R., and Dyck, J., Tee, K., and McEwan, G., 2008. Supporting informal collaboration in shared-workspace groupware. *Journal of Universal Computing*, in press.
- Heath, C. and Luff, P., 1992. Collaboration and control: Crisis management and multimedia technology in London underground line control rooms. *Journal of Computer Supported Cooperative Work*, 1(1), Springer, 24-48.
- Isaacs, E., Walendowski, A., Whittaker, S., Schiano, D. J., and Kamm, C., 2002. The character, functions, and styles of instant messaging in the workplace. *Proc. ACM CSCW*, 11-20.
- Kaasten, S., Greenberg, S., and Edwards, C., 2002. How people recognize previously seen web pages from titles, URLs and thumbnails. In X. Faulkner, J. Finlay, and F. Detienne (Eds.), *People and Computers XVI (Proceedings of HCI BCS Conference Series)* (pp. 247-265): Springer Verlag.
- Kraut, R., Egidio, C., and Galegher, J., 1988. Patterns of contact and communication in scientific research collaboration. *Proc. ACM CSCW*, 1-12.
- Kurlander, D., Skelly, T., and Salesin, D., 1996. Comic Chat. *Proc. ACM Computer Graphics and Interactive Techniques*, 225-236.
- McEwan, G. and Greenberg, S., 2005. Supporting social worlds with the Community Bar. *Proc. ACM GROUP*, 21-30.

- McEwan, G., Greenberg, S., Rounding, M., and Boyle, M., 2006. Groupware plug-ins: A case study of extending collaboration functionality through media items. Proc. CollabTech, IPSJ SIG Groupware and Network Services, 42-47.
- Nardi, B. A., 1993. A small matter of programming: Perspectives on end user computing: MIT Press.
- Nardi, B. A., Whittaker, S., and Bradner, E., 2000. Interaction and outercation: Instant messaging in action. Proc. ACM CSCW, 79-88.
- Neustaedter, C., Greenberg, S., and Boyle, M., 2006. Blur filtration fails to preserve privacy for home-based video conferencing. ACM Transactions on Computer-Human Interaction, 13(1), ACM Press, 1-36.
- Richardson, T., Stafford-Fraser, Q., Wood, K. R., and Hopper, A., 1998. Virtual Network Computing. IEEE Internet Computing, 2(1), IEEE Educational Activities Department, 33-38.
- Romero, N., McEwan, G., and Greenberg, S., 2007. A field study of Community Bar: (Mis)-matches between theory and practice. Proc. ACM GROUP, 89-98.
- Root, R. W., 1988. Design of a multi-media vehicle for social browsing. Proc. ACM CSCW, 25-38.
- Rounding, M., 2004. Informal awareness and casual interaction with the Notification Collage. MSc Thesis, University of Calgary, Calgary, Alberta, Canada.
- Sarin, S., 1984. Interactive on-line conferences. PhD Thesis, MIT, Cambridge, Massachusetts, United States.
- Smale, S. and Greenberg, S., 2005. Broadcasting information via display names in instant messaging. Proc. ACM GROUP, 89-98.
- Tang, J. C. and Rua, M., 1994. Montage: Providing teleproximity for distributed groups. Proc. ACM CHI, 37-43.
- Tee, K., Greenberg, S., and Gutwin, C., 2006. Providing artifact awareness to a distributed group through screen sharing. Proc. ACM CSCW, 99-108.
- Tee, K., 2007. Artifact awareness for distributed groups through screen sharing. MSc Thesis, University of Calgary, Calgary, Alberta, Canada.
- Voida, A., Grinter, R. E., Ducheneaut, N., Edwards, W. K., and Newman, M. W., 2005. Listening in: Practices surrounding iTunes music sharing. Proc. ACM CHI, 191-200.
- Whittaker, S., Frohlich, D., and Daly-Jones, O., 1994. Informal workplace communication: What is it like and how might we support it? Proc. ACM CHI, 131-137.
- WOS Data Systems, 1987. Timbuktu: The next best thing to being there, user's guide (version 1).