

Artifact awareness through screen sharing for distributed groups

Kimberly Tee^{a,*}, Saul Greenberg^a, Carl Gutwin^b

^aUniversity of Calgary, 2500 University Drive NW, Calgary, Alta., Canada T2N 1N4

^bUniversity of Saskatchewan, 57 Campus Drive, Saskatoon, Sask., Canada S7N 5A9

Received 4 June 2008; received in revised form 6 April 2009; accepted 7 April 2009

Communicated by A. Sears

Available online 18 April 2009

Abstract

Co-located collaborators can see the artifacts that others are working on, which in turn enables casual interactions. To help distributed collaborators 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. A preliminary 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.

© 2009 Elsevier Ltd. All rights reserved.

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 with shared interests – is important for seeding collaborations, 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 miss 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 messaging (IM) systems by diverse user groups (Nardi et al., 2000), or how

*Corresponding author. Tel.: +1 403 210 9499; fax: +1 403 284 4707.

E-mail addresses: tee@cs.ucalgary.ca (K. Tee), saul@cs.ucalgary.ca (S. Greenberg), gutwin@cs.usask.ca (C. Gutwin).

email is often used for casual ongoing conversations rather than as a formal messaging system. Many existing electronic 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 distributed people are using as they perform their individual, ongoing work’. By individual work, we mean the activities that one is pursuing at the moment by oneself, even though this work may be something that others are involved in over time. This awareness contributes to all parts of interpersonal interaction—it helps people discover opportunities for initiating casual interaction, it informs knowledge exchange, it is vital for helping people coordinate their tightly coupled work, and it helps people build social relationships. For example, consider collaborators who inhabit a co-located environment such as a shared office space. Artifact awareness is typically easy and tacit. As collaborators move around the space, they naturally gather visual and auditory cues about other people’s presence and activities. They notice what artifacts others are working on as they glance at their desks and see their computer displays. This creates interaction opportunities, where they can easily move into more tightly coupled work. 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, discuss, and even work over that artifact with its owner. That is, artifact awareness moves between information noticed in the periphery (or background) of one’s attention, through to information that is in one’s active (or foreground) attention.

Artifact awareness is crucial in some co-located settings (and indeed in various theories of group behaviour), where work environments are explicitly crafted to encourage such awareness between collaborators. First, consider the *studio space* used by groups of designers, artists, and/or architects. In these studios, people craft their work on easels or other semi-public surfaces strategically placed to be within the flow of traffic, so that others moving through the space can monitor, comment, and critique the work as it unfolds over time (Buxton et al., 2000). The studio space promotes easy visibility of people’s work, which in turn leads to casual interaction. Indeed, the hallmark of a good studio is whether it has a continual ‘buzz’, where its inhabitants are constantly discussing and reflecting on current projects and, in the process, are learning from one another. Second, consider *supervisory control environments*, e.g., for subway routing, power plants, factories, and so on. A good supervisory control environment provides easy visibility of the actors, their actions, and their artifacts (Heath and Luff, 1992). Through this visibility, the team is aware of what individuals are doing, the effect of these actions, and the global state of the system. It creates implicit

opportunities where people can micro-coordinate and interleave actions to control the system as a whole. That is, a good control room enables distributed cognition, a cognitive organization where the team seamlessly works together (Hutchins, 1995). Third, consider the *everyday conversational environment* and how artifact visibility plays a role in supporting conversations by the way it contributes to joint activities and common ground (Clark, 1996). If the conversation involves the artifacts around them, the artifact helps participants form their initial common ground (i.e., the presupposed facts, assumptions, and beliefs recognized and shared between them), and then informs the current state of the joint activity (i.e., how they perceive, coordinate, and synchronize their individual and joint activities). In part, the verbal conversation often confirms, clarifies, and corrects people’s mutual understanding and actions surrounding these artifacts (Clark, 1996).

The various theories of casual interaction, distributed cognition, common ground, and studio design all presuppose, to some extent, a degree of artifact awareness. However, even in a co-located environment, artifact awareness can be difficult to maintain, particularly if the artifact is digital. Most digital artifacts appear only 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 collaborators 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.

Given all the above, our first research goal is to recreate artifact awareness for distributed groups:

Goal 1. Design a system that supports artifact awareness and opportunistic interaction, where the audience is a small distributed group with a strong desire to stay in up-to-the-moment contact.

We have a specific audience for this system. We primarily consider the well-formed group somewhat similar to a co-located work group comprising trusted people with both a working and social relationship, and who share at least a modest degree of mutual interest in the artifacts under consideration. That is, we target *intimate collaborators* who have a real need and desire to stay connected. In sharp contrast, we are not suggesting such a system for (say) an Internet café, where strangers frequently come together.

To satisfy this goal, we created an awareness tool that leverages *screen-sharing* techniques to provide the distributed group of intimate collaborators with mutual awareness of artifacts and individual work by letting each other continually see some or all of each other's screen contents. 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'.

At this point, many readers may become concerned with our goal due to its privacy implications (Boyle and Greenberg, 2005). As mentioned, while people's screens may contain semi-public information relevant to the group, they may also contain highly personal information that one would likely not want to share. The existing culture of screen use, where screens are considered personal vs. public spaces, means that even people in a closely knit group may be unwilling to share their screens and/or view other people's screens. Thus our second research goal is to balance the awareness information people want others to have of their work with their own privacy needs:

Goal 2. Design privacy controls that allow individuals to balance the artifact awareness they want to provide to others with their own particular privacy needs.

We can now restate our hypothesis as: 'screen sharing is a critical component of distributed artifact awareness, because one's screen contents, as mediated by privacy controls, let a person selectively indicate what digital activities they want their collaborators to see'. Yet this is just a hypothesis, and we need to see if this hypothesis holds in real situations, at least in principle. Thus our third goal is to evaluate the system and privacy mechanisms.

Goal 3. Evaluate how groups use screen sharing for artifact awareness, and the effectiveness of its privacy controls.

This article¹ begins with a brief discussion of previous work on awareness tools and screen sharing (Section 2).

We then describe and provide the design justification of the artifact-awareness system that we built (Section 3) and its privacy mechanisms (Section 4). Finally, we report preliminary results from end-user evaluations (two field trials; Section 5). We close with a discussion of the ways in which the idea of using screen sharing to support artifact awareness can benefit distributed work groups.

2. Related work

Artifact awareness is not new, in the sense that many existing collaboration theories and groupware systems contain some notion of artifact sharing and support how group members coordinate their actions over these artifacts by maintaining some degree of awareness of them. To provide context, a sampling of these theories and systems is briefly described below. What sets our work somewhat apart from this prior work is that we are primarily interested in how artifact awareness leads into casual interaction (particularly through screen sharing), and the nuances of supporting such awareness in light of privacy concerns.

2.1. Artifact awareness in everyday coordination and work

As mentioned, artifact awareness is one person's up-to-the-moment knowledge of the artifacts and tools that other distributed people are using as they perform their individual, ongoing work. 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. More generally, artifact awareness is valuable – sometimes crucial – for team work. Indeed, various theories of group activity incorporate artifact awareness as a key element. Several examples are listed below.

Casual interaction theory (Kraut et al., 1988). Artifact information informs many aspects of casual interaction. First, it creates serendipitous opportunities for people to engage in artifact-oriented conversations, and to move into collaboration over the artifact. We previously mentioned that 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. In some cases, this interaction helps the group accomplish a goal, e.g., Nardi (1993) found that people opportunistically collaborate over spreadsheets, CAD systems, and other documents over the course of a day. In other cases, the interaction is more social, e.g., Greenberg (1999) describes how seeing another person's public activity, even if it is not part of a joint task, can trigger interest and conversation around that activity. Second, knowledge of artifacts in use is yet another source of contextual information that helps people determine

¹While parts of this article have appeared previously in (Tee et al., 2006), this paper expands considerably on it.

others' availability, e.g., how busy others are and when they can be interrupted (Kraut et al., 1988; Whittaker et al., 1994).

Workspace awareness theory (Gutwin and Greenberg, 2002). Workspace awareness is defined as 'the up-to-the-moment understanding of another person's interaction with a bounded shared visual workspace' (e.g., a tabletop surface, a whiteboard, a shared display). It is awareness of people and how they interact with the workspace and its contained artifacts, rather than just awareness of the workspace itself. It is limited to events happening in the workspace—inside the temporal and physical bounds of the task that the group is carrying out. Similar to situated awareness theory, this awareness informs the group's perception of relevant elements of the environment, their comprehension of those elements, and the prediction of the states of those elements in the near future (Endsley, 1995). The artifacts provide this awareness information in how they manifest themselves within the space, and in how they provide the group with feedthrough, i.e., when artifacts are manipulated, they give off information that informs others who are watching (Dix et al., 1993). While workspace awareness includes some degree of artifact awareness, the two are not the same. Workspace awareness describes how people monitor their focused interaction over a small visual area, while artifact awareness (as used in this paper) is focused on how artifact visibility within the environment both informs participants and creates opportunities for casual interaction that then leads into this focused interaction over these artifacts.

Common ground theory (Clark, 1996). Artifact visibility supports conversations by the way it contributes to people's understanding of each other's utterances and joint activities. When artifacts are not visible, people have to expend greater effort in understanding, confirming, and repairing verbal conversations concerning these artifacts. As with workspace awareness, common ground mostly concerns itself with focused interaction.

Distributed cognition theory (Hutchins, 1995). The artifact state and people's actions around artifacts inform how the team works and think together as a whole (Heath and Luff, 1992; Hutchins, 1995). For example, collaborators who are responsible for different aspects of a joint task can monitor each other's progress and coordinate their activities. Distributed cognition is a more global perspective of group activity, where casual interaction, focused work, workspace awareness, artifact awareness, and common ground all inform the group's cognition.

2.2. Systems supporting informal awareness, casual interaction, and joint work

A large number of groupware tools have been developed to support informal awareness and casual interaction

leading to joint work within small communities of distributed collaborators (e.g., the ACM CSCW and CHI Proceedings document the evolution of many such systems over the years). 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. A sampling is included below.

Instant messaging 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 share artifacts with their co-workers.

Shared workspaces provide a distributed group with a common digital work area. Ideally, all participants can work simultaneously, where all artifacts and people's fine-grained actions with them are immediately reflected on

other people's displays. Examples include graphical drawing systems (e.g., groupware sketchpads and CAD systems) and joint editing systems (e.g., groupware text editors). Early systems typically implemented 'what you see is what I see (WYSIWIS)', where all saw exactly the same thing on their displays (Stefik et al., 1987). While this worked well for tightly coupled work, it proved too constraining for loosely coupled work. Thus later systems relaxed WYSIWIS. People could now scroll to different regions of the workspace (Stefik et al., 1987), or even have different view representations of the artifacts within it (Greenberg et al., 1996). Yet workspace awareness in these relaxed WYSIWIS settings was still considered crucial for group activity. Consequently, researchers created and evaluated a broad variety of awareness widgets (e.g., Gutwin et al., 1995; Gutwin and Greenberg, 2002; Greenberg et al., 1996). A variant of the shared workspace is screen sharing, which will be discussed shortly in its own subsection.

Collaborative virtual environments (CVEs) expand on the notion of a shared workspace, where participants move about a (usually 3-D) spatial world (e.g., Sohlenkamp and Chwelos, 1994). Within this world, people see its stationary aspects (e.g., buildings, landscape), the other people who inhabit the world, and the artifacts within the world that can be manipulated. CVEs have become very successful as social virtual worlds (e.g., Second Life), and in massively multiplayer online games (as found in many PC and console network game systems). While applying CVEs to work settings have been a constant theme over the years (e.g., Sohlenkamp and Chwelos, 1994; Greenberg and Roseman, 2003), they have yet to gain more than niche acceptance.

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 (1992) study 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.

The social web – web-based systems and sites that let people link and socialize with one another – have caught the world by storm (e.g., Porter, 2008). Systems vary considerably. Some sites, such as MySpace, Facebook, and LinkedIn, are centered around how people identify themselves on the web, and how they form relationships with others on the web. These include not only strong relationships (collaborators, close friends), but weak ties as well. Artifact sharing happens as a side effect and consequence of these relationships, where artifacts (such as personal photos or activity announcements) are broadcast and available to others. Other social web sites are focused primarily on particular styles of artifact sharing, such as Flickr for photos, SlideShare for PowerPoint slides, and YouTube for videos. Artifact awareness happens two ways. First, people can subscribe to events (such as posting of photos) produced by others, where they receive asynchronous notifications that give an overview of the event as well as a way to quickly access details of that event via a web link. Second, they can opportunistically browse the web sites of their relations – often done as a side effect of receiving a notification – and explore the various links. While this lets people know what others are doing, it rarely leads to real-time interaction over the artifact.

Going beyond individual systems, the need for large distributed teams to share information and artifacts has long been recognized by the research community. For example, in 1989 the US National Science Foundation promoted the idea of a *national collaboratory*, a nationwide structure that provides tools enabling researchers to perform their work and collaborations without regard to geographical locations (Lederberg and Uncapher, 1989). Collaboratories were expected to provide many functions, such as a community knowledge base, the ability to find collaborators, project management, digital libraries, and so on. Most relevant to artifact awareness is the idea of supporting remote and collaborative experimentation, where people can share scarce resources and physical instruments. An example is the UARC Collaboratory, where various tools were crafted to provide the distributed upper atmosphere research community with 'real-time access to remote instrumentation and ... collaborative tools that would allow them to interact with each other over real-time data' (Olson et al., 1998).

2.3. Screen sharing

Screen-sharing applications let people explicitly share their computer screens, windows, or window fragments

with each other. 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 (Lauwers and Lantz, 1990). Because these applications were usually designed for just a single person to use at a time, the shared view is usually implemented by intercepting the graphical output and sending it to multiple screens. Similarly, input is usually done through turn taking, where one person is given control of the system, and their input actions are funnelled to the application. While free-for-all input is possible, the application can become quite confused unless participants take care in how they interleave their actions.

Because of the above, screen sharing has many known disadvantages. People cannot work in parallel, they cannot easily see different parts of the display, they must resort to turn taking, and so on. However, its primary advantage outweighs these many disadvantages—they work over the unaltered single-user applications, which make up the vast majority of software systems in use today. This is why we use screen sharing for the basic artifact-awareness system we report in this paper, rather than a true groupware system.

Screen sharing was first demonstrated in 1968 (Engelbart and English, 1968), replicated in more detail in 1984 (Sarin, 1984), and then gained considerable research interest in the late 1980s (e.g., Lauwers and Lantz, 1990). It became available commercially and has been used in real-time distributed collaboration for around 20 years, e.g., Farallon Timbuktu (WOS Data Systems, 1987), VNC (Richardson et al., 1998), and NetMeeting (Summers, 1998). Screen sharing 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). In essence, screen sharing emulates over-the-shoulder, face-to-face sharing of a computer—one person can show others what they are working on, and each can take turns interacting with the system.

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 leading to casual interaction. There are a few exceptions. First, the SMART Sync classroom management software – formerly called SynchronEyes – is a commercial system that lets one person view others’ desktops as thumbnails (SMART Technologies, Inc., 2009). 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 us and 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., 2008) 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

Goal 1. Design a system that supports artifact awareness and opportunistic interaction, where the audience is a small distributed group with a strong desire to stay in up-to-the-moment contact.

Basic design justification: Mimic the real-life way co-located collaborators maintain artifact awareness—they glance at each other’s work as they move about the shared space, and have the opportunity to move into casual interaction if desired.

As we will see, we created an artifact-awareness tool based on screen sharing to provide a small group of intimate collaborators with awareness of each others’ artifacts and their individual work. Because we wanted people to be able to move into casual interaction, we embedded this tool into the Community Bar (CB), an augmented media space developed in our laboratory that provides a group with both awareness and communication capabilities (McEwan and Greenberg, 2005).

3.1. Community Bar overview

We begin by briefly describing the Community Bar, a groupware environment that provides basic facilities for group awareness and communication. Additional details about the system and its design justification are found in McEwan and Greenberg (2005) and McEwan et al. (2006).

Fig. 1 shows a screen snapshot of CB in use. As can be seen, CB is divided into *Places*; two are shown in the figure—‘mike test’ and ‘CSCW class’. Each place represents a sub-group of people and the communications, tools, and information they are sharing. These are visualized through a number of *media items*, all holding different information and all being publicly visible to the sub-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, initiate opportunistic conversations with others, or to easily and serendipitously join an existing conversation. CB allows

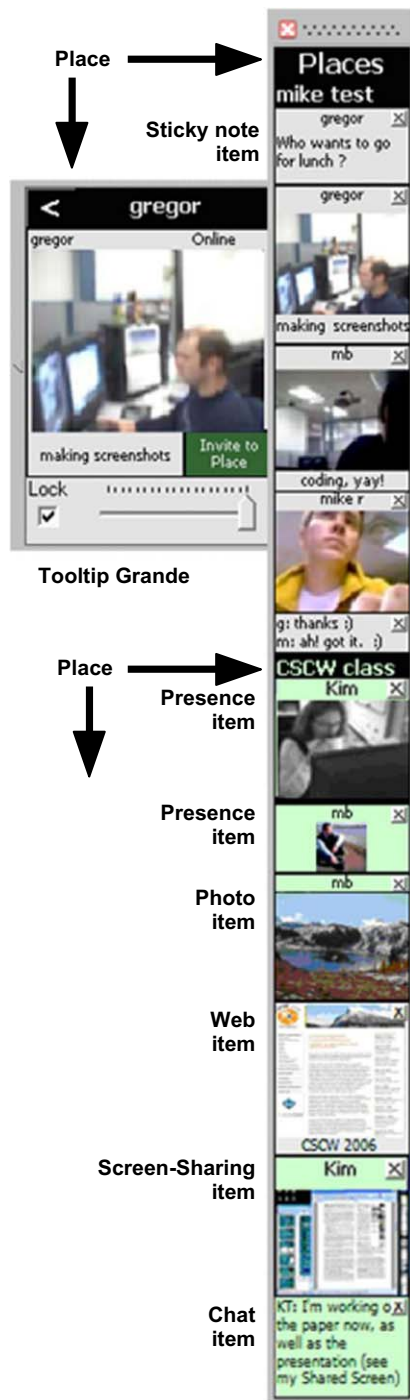


Fig. 1. Community Bar.

people to share some of their artifacts, where one person can explicitly post information they believe to be of interest to others in the group. For example, one can share web pages, photos, and slide shows with others through the *Web item* and *Photo item*.

Media items are presented within a sidebar that displays three levels of granularity, each providing different levels of awareness information and interaction possibilities (Cadiz et al., 2002). The media item's *tile* view is always visible in

the sidebar; these are the many small items seen in Fig. 1. 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 Fig. 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 (not shown in Fig. 1). This 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, owners of a Presence item have 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. We note, however, that each person can have a somewhat different view of this community. For example, each person can resize particular tiles to increase or decrease his or her own level of awareness of that particular activity. Similarly, a person raising a tooltip grande and full view are local actions, i.e., others' displays are not affected. Next, while basic information is always visible at the periphery, progressively more information can be revealed through focused interaction. The sidebar design encourages peripheral awareness because it cannot be covered up and because it is situated at the screen's side (Cadiz et al., 2002).

As currently described, CB lacks support for artifact awareness and artifact-centered serendipitous interactions. That is, people do not have the ability to show their actual ongoing work to others, to monitor other people's work, or to move into interaction over that work. To enable this, we created the Screen-Sharing media item.

3.2. The Screen-Sharing item

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). We chose to develop our screen-sharing awareness tool as a CB media item plug-in. By doing so, we could immediately leverage 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 which had been using CB on a daily basis for



Fig. 2. Thumbnails of people's screens.

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 Section 5. The following subsections describe by scenario how the *Screen-Sharing item* works,² as well as our design justification for each component.

3.2.1. Tile view

Design justification: The ‘owner’ of the work artifacts should be able to provide others with a coarse overview of their ongoing work by broadcasting infrequently updated miniatures of their screen. Similarly, the audience members should see this overview at the periphery of their attention.

‘Kim’ (initials K.T.) 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 in the sidebar, which contains a thumbnail of her entire screen labelled with her name (Fig. 1, second tile from the bottom). At the same time, she adds a *Chat item* (Fig. 1, bottom tile) saying “I’m working on the paper now, as well as the presentation (see my Shared Screen)” – this serves the same role as a verbal aloud (Heath and Luff, 1992), where it 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 a coarse overview (i.e., basic awareness) of what Kim is working on. Typically, the visual characteristics of windows within the 70×60 pixel thumbnail are sufficiently recognizable (Kaasten et al., 2002) so that others can tell if the person who posted their screen 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 Fig. 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 Picasa 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.

Now reconsider the shared screen in the tile from Fig. 1. 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 it is from their knowledge of the paper.

3.2.2. Tooltip grande

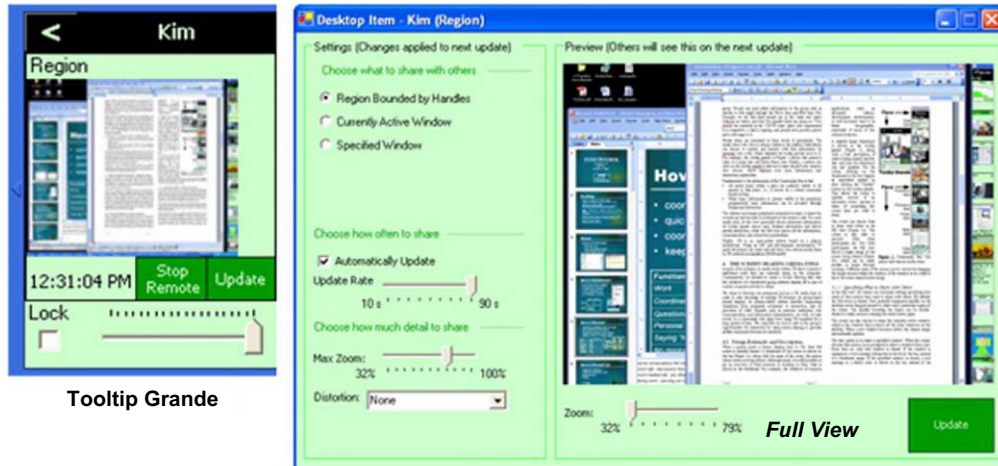
Design justification: An audience member should, if interested, be able to quickly see additional details of another person’s work artifacts by raising a higher fidelity, annotated view of the shared artifact. This act is somewhat equivalent to an extended glance.

The tooltip grande for both the owner and the audience shows a somewhat larger thumbnail (Figs. 3a and b, 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 Figs. 3a and b (left) show that she is sharing a region of the screen, and that it was last updated at 12:31 PM today. The arrow at the top left of the tooltip grande, common to all CB media items, is used to invoke an even more detailed full view, discussed shortly. 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 (Fig. 3a, left).

Design justification: An audience member should be able to reduce or increase the visibility and detail of the shared artifact to correspond with their interest in it, e.g., by adjusting the size of the tile.

The lock and slider that appear at the bottom of the tooltip grande are common to all CB media items, and are

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



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



The audience's tooltip grande and full view of the Screen-Sharing item.

Tile view, all people

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 Fig. 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.

3.2.3. Full view

Design justification: An audience member should, if interested, be able to see full details of the shared artifact, equivalent to directly attending the artifact after being attracted to the information contained by it. As well, the owner should be able to prescribe limits on what the viewer can see to safeguard privacy.

The full view as seen by both the audience and the owner gives a larger and much more detailed preview of the captured display (Figs. 3a and b, 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 Fig. 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 Figs. 3a and b, 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 (Fig. 3b, full view), while Kim is looking at it at a 32% zoom so that the entire region fits the full view's window (Fig. 3a, full view). It can also be seen that 79% zoom of the true screen resolution produces a fairly legible image (Fig. 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

Design justification: As appropriate opportunities arise, an audience member should be able to engage the artifact owner by moving into conversation and direct interaction over the shared artifact.

Any audience member can initiate a real-time remote pointing session by pressing the ‘Remote Pointing’ button in the tooltip grande or the full view (Fig. 3b). 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.

If the owner approves it, a remote pointing window (Fig. 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 (Fig. 4a, left) linked to a Chat item in the sidebar (Figs. 4a and b, 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 Fig. 4a. A corresponding named 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 Fig. 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

Goal 2. Design privacy controls that allow individuals to balance the artifact awareness they want to provide to others with their own particular privacy needs.

Basic design justification: Give artifact owners full control of what artifacts they wish to share with others, as well as how viewable those artifacts are.

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 dependent on the group culture and the actual practice of use that develops over time. As an always-on media space, CB 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) who 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 who insists that all employees use the Screen-Sharing media item so that their work can be monitored (see Section 4.4). However, even in the benign case, people may want some control over what others can see.

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*. In the framework, each modality contains additional terms that detail aspects of it. The modalities and a sampling of terms are described below. Full details are available in Boyle and Greenberg (2005) and Boyle et al. (2009).

Confidentiality relates to a person’s control over what information others can access about oneself. It directly regulates the outward flow of information to others. Example information topics include identity, activity whereabouts, encounters, utterances, actions, and relationships. Other terms include information sensitivity, security, authorization, information fidelity (precision, accuracy, temporal, misinformation, and disinformation), ambiguity, surreptitious surveillance, misuse, and misappropriation. Confidentiality is also influenced by properties of the information channel medium, i.e., whether it is aural, visual, numeric, or textual, and how that information is obtained, e.g., by sampling, interpolation, aggregation, and inference.

Autonomy relates to how a person chooses to present him/herself to others. That is, it is the freedom to choose how one acts and interacts in the world. Example terms include social construction of self, identity, digital persona, appearance, impression, social relationships, back- and front-stage performances, and aesthetic and strategic harms. It also includes norms about self, such as expectations, preferences, deviance, conformance, and social acceptability.

Solitude relates to how a person controls interpersonal interactions with oneself, i.e., how a person can be ‘apart’ from others. Example defining terms include

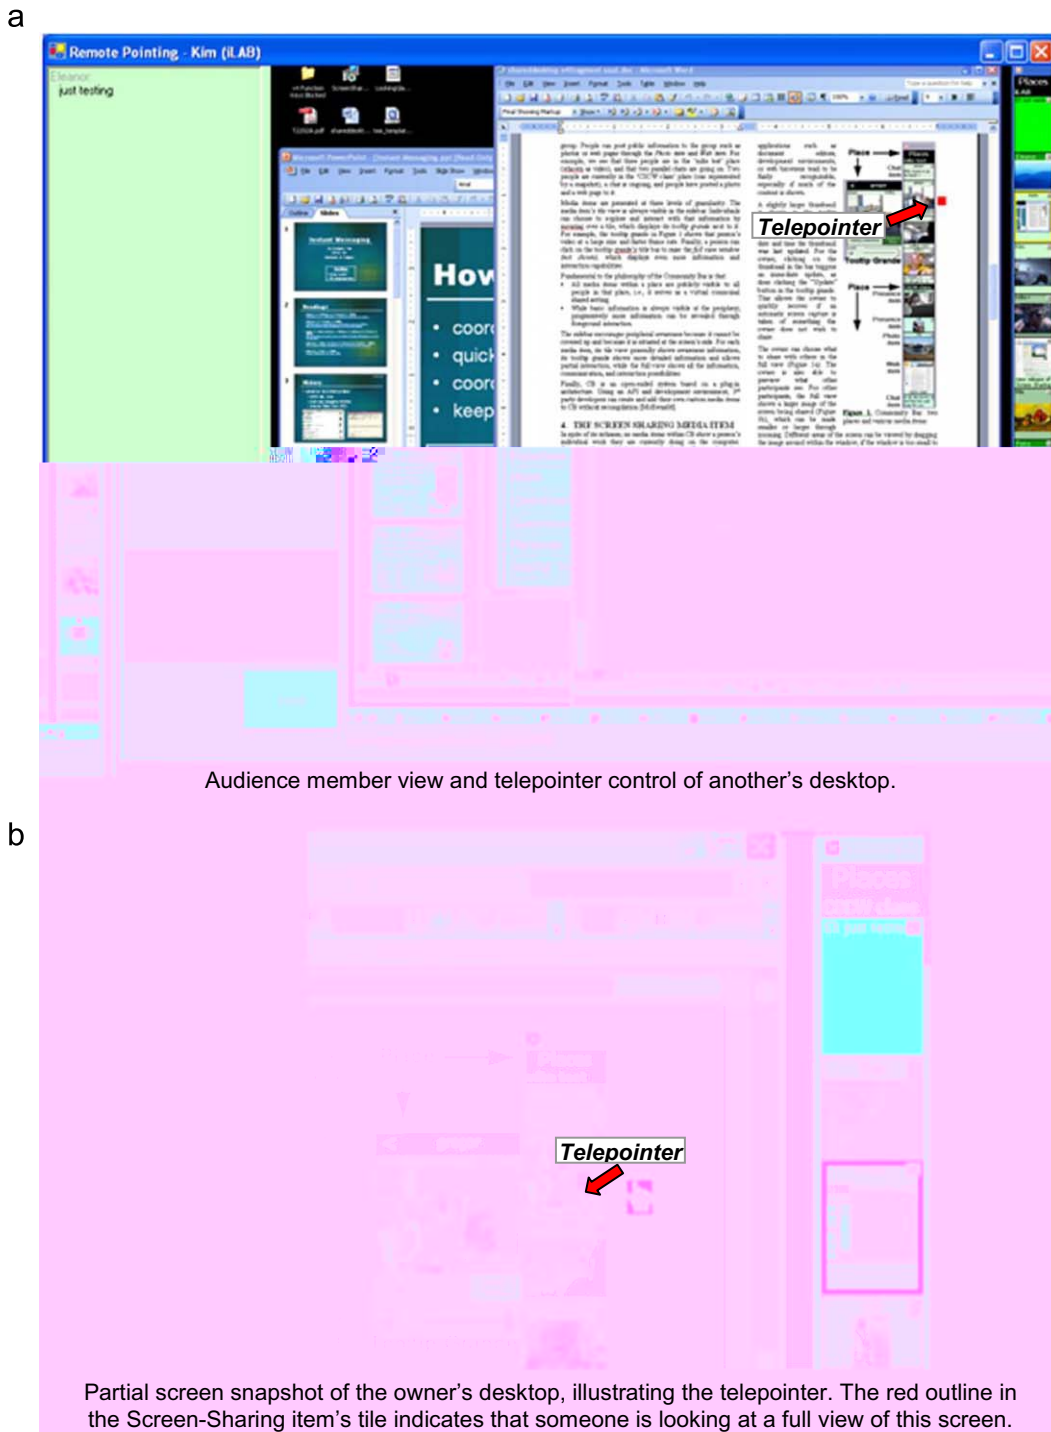


Fig. 4. A remote pointing session.

crowding, isolation, intimacy, anonymity, attention, refuge, awareness regulation, distraction, availability, and accessibility.

In the following subsections, we use these three modalities and some of their terms (shown italicized) to describe our design justification for privacy-related controls and feedback in the Screen-Sharing item. We also discuss some of the limitations of these privacy controls,

which were revealed by a *post hoc* theoretical analysis of privacy using Boyle's (2005) technique of systematically applying his framework to the Screen-Sharing item.

4.1. Specifying what to share with others

Design justification: The primary concern of artifact sharing within the CB is confidentiality, and the owner of

the artifact should be able to control one's information confidentiality.

The CB targets a small community of intimate collaborators, i.e., people who know each other well and have a desire to stay connected. Thus a person's *autonomy* is not particularly threatened, as the CB community should extend existing social practices. Similarly, because each member can opt to enter or leave a CB place at any time, or remove CB from their display, *solitude* is also not largely threatened. Of course, solitude and autonomy threats can become serious if CB is not used as expected. For example, a dictatorial manager can insist that his staff join CB with their video media item always on (affecting *autonomy*), where they are expected to immediately engage in all conversations initiated by that manager (affecting *solitude*).

Turning specifically to the Screen-Sharing item, its function is to share the contents of a person's screen with other group members through the *visual information channel*, obtaining this information by *sampling* the bitmapped contents of one's screen. Thus, within the context of Boyle and Greenberg's framework, 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. From the shared screen image, much information can potentially be seen or inferred. The information being shared can be *sensitive* (e.g., personal 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. For example, people often pursue *personal activities* on their computer that they may not wish others to see: playing games, reading and composing personal email, online shopping, etc. Even if the information contents are not sensitive, it may reveal *information about self* (e.g., an online gambling addiction) that one may not want others to see. Alternately, they may be working on projects where information should not *persist*, e.g., a legal report where drafts are discoverable. While CB does not persist this information, there is nothing to stop others from capturing and storing it through an image-grabbing tool. As another example, one group member may be viewing information that she – and only she – is *authorized* to see; accidentally revealing this to others undermines the authorization control.

Consequently, the Screen-Sharing item offers a number of privacy controls primarily for mitigating confidentiality concerns.

4.1.1. Specifying the capture region

Design justification: Owners should be able to restrict what visual information is pushed out – and thus is viewable – to others.

Owners have full control over what *information visuals* 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: 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 owner-specified window. The choice restricts what others can see to only those parts of the display the owner wishes to reveal. The bitmap image explicitly reveals only the visual information being worked on. However, viewers will likely be able to infer other information, such as the owner's *activities* (from changes made) and one's whereabouts (if the information is changing then the owner is likely there working on it). Because viewers can see what the owner is doing (or not doing), the owner becomes *accountable* to their actions, and *plausible deniability* is threatened.

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. This can be readjusted at anytime. Semi-transparent red handles define the region being shared (Fig. 5), and the owner can adjust the bounds of the region by dragging these handles around. The owner can quickly return to sharing the entire (primary) screen by double clicking a handle. For people with two displays, right double clicking a handle shares the full secondary screen.

The owner can opt to share individual windows instead of screen regions. 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. This does not safeguard privacy very well, for the owner can inadvertently share a window containing sensitive information. 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.

4.1.2. Specifying update frequency

Design justification: Owners should be able to adjust the sampling of the visual information pushed out, which controls the temporal fidelity and persistence of information that others can see.

Our second level of privacy control lets the owner specify how often the display should be captured, and thus how often the audience gets this update. This control allows owners to manage *temporal fidelity*. Specifically, the owner can specify if updates are manual or automatic. If manual, the display is updated only when the owner explicitly clicks the thumbnail in the tile view, or the 'Update' button in the tooltip grande and full view. If automatic, the owner can

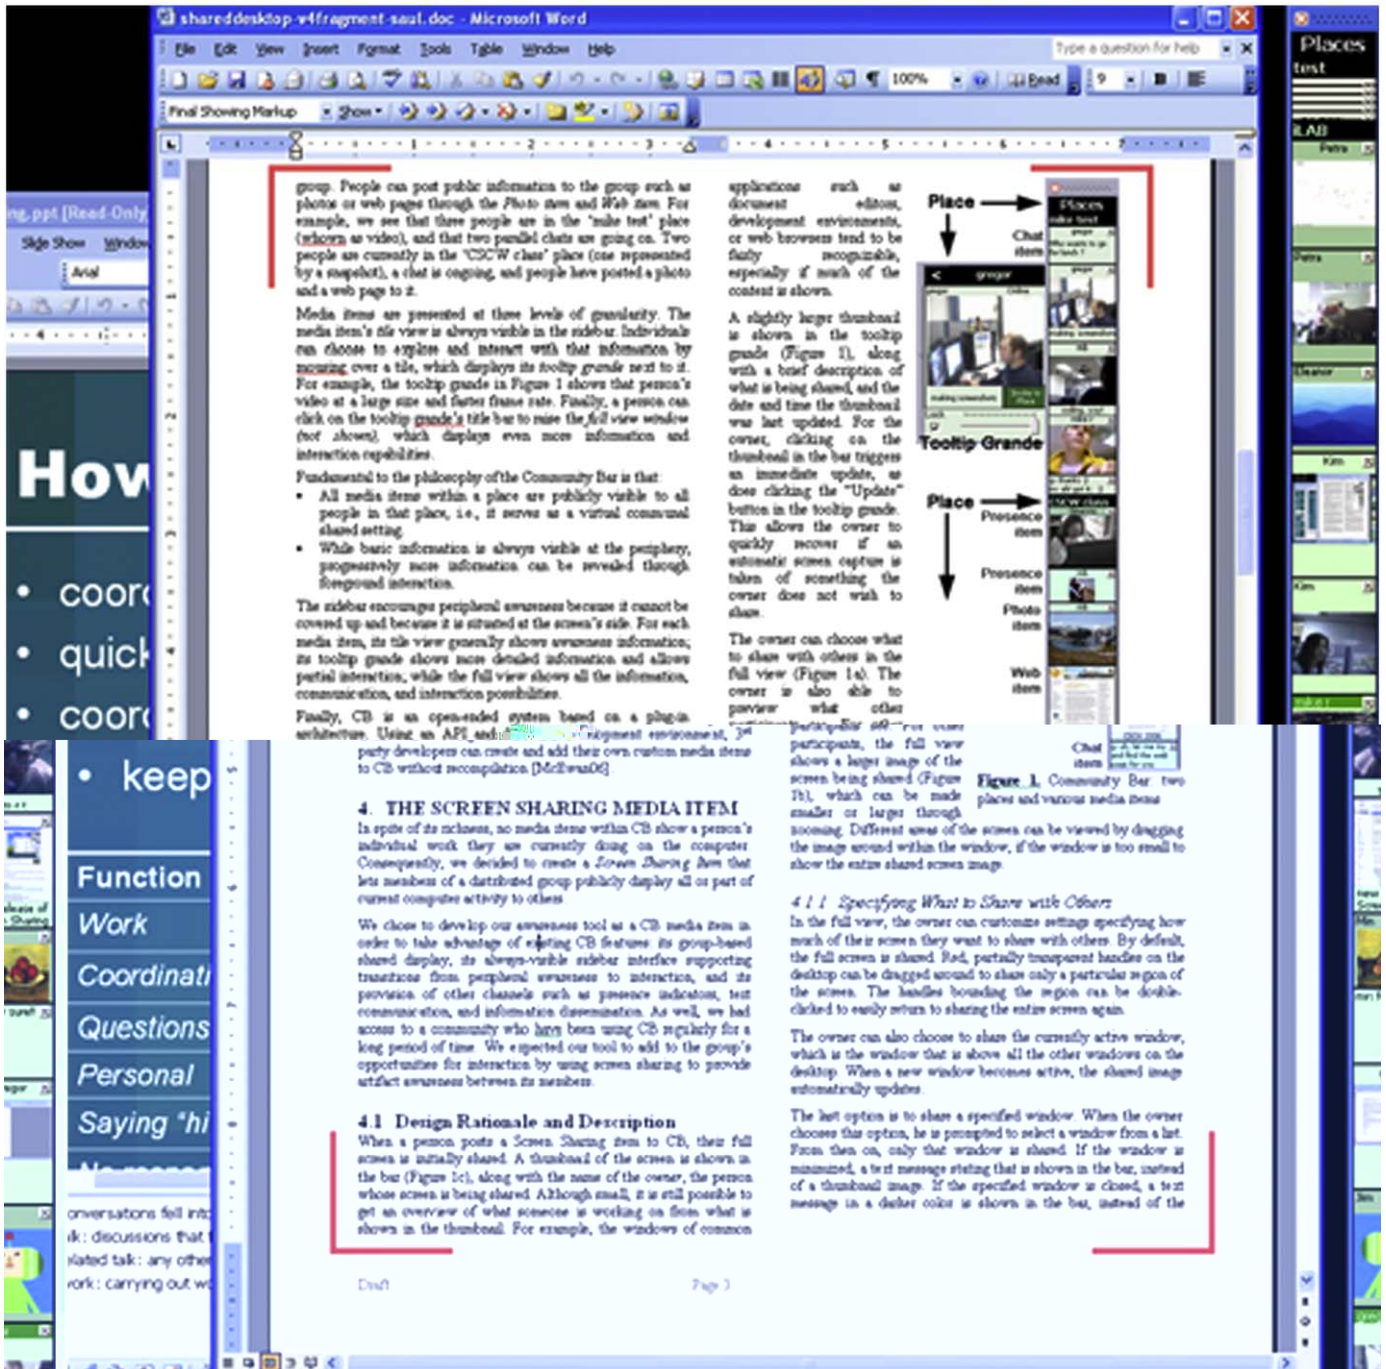


Fig. 5. The capture region of the screen is marked by adjustable semi-transparent red handles.

specify an update frequency interval between 10 and 90 s, where they can still opt to transmit an immediate snapshot at any time by clicking the thumbnail/update button. Unlike commercial Screen-Sharing systems tailored for real-time interaction, this infrequent auto-update should suffice for artifact awareness, as they help inform others of basic activities while minimizing *distraction* that might otherwise arise from real-time movement in the various views.

Because screen contents are largely *transient*, the shared information should reflect that attribute. The information shared by the Screen-Sharing item typically has low *persistence*; on auto-update, a new shared screen image replaces the previous one at least once every 90 s. Images are not stored anywhere for *archival* purposes; when an owner logs off of Community Bar, their Screen-Sharing item (and thus their shared screen image) is automatically removed.

4.1.3. Specifying how much detail to share

Design justification: Owners should be able to adjust the visual fidelity of the information pushed out, which limits the details that others can see.

The third level of privacy control lets the owner manipulate the *image fidelity* that others can see through *distortion filtration*. One of our techniques reduces image resolution by letting the owner adjust the zoom level and/or distort 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 zoom level of around 33% and her captured region encompasses 1280×1024 pixels, the shared image is visually compressed to about 1/9 of the original area ($\sim 426 \times 341$). 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, Fig. 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.

Of course, image fidelity is not guaranteed to fully protect confidentiality. For example, text legibility is influenced by font size and word shape. Some images,

especially if large and of high contrast, may be easily recognizable even when heavily blurred. Still, it should be a reasonable safeguard for typical sensitive information, such as email and detailed document contents.

4.2. Providing feedback

The Screen-Sharing item has secondary effects on *autonomy* through its relationships to confidentiality. Autonomy is the component of privacy that relates to how a person chooses to present himself/herself through ‘observable manifestations of the self’ (Boyle and Greenberg, 2005) such as appearance, identity, and impression. 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). One’s behaviour can also be judged in terms of how well they follow the norms of the work culture, i.e., where actions *conform* or *deviate* from *expectations*.

Our strategy to safeguard these threats is to give people feedback about what they are transmitting, so that they can adjust either their behaviours or what is being transmitted as needed.

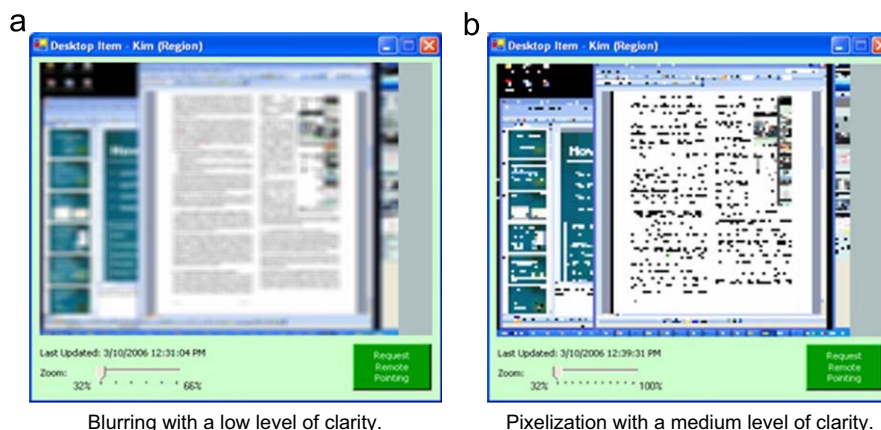


Fig. 6. Masking effects applied to the shared screen image.

4.2.1. Feedback of image capture

Design justification: Owners need feedback about what they are transmitting, so they can safeguard threats to their autonomy.

The Screen-Sharing item aims to prevent mistaken *back-stage performances* from occurring on the *front stage* by providing appropriate warnings or reminders before a screen capture is taken. That is, the fourth level of privacy control is for the system to provide sufficient *feedback* to the owner about what others can see. First, the Screen-Sharing item supports *reflexive interpretability of action*—the owner can always see exactly what the audience can see, because their 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 always see exactly the same transmitted 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 (Fig. 4b). This outline remains until the full view is closed, indicating to the owner that at least one person has their full view open. However, no identifying information is supplied as to who is looking at the full view; we initially wanted to encourage opportunistic interaction and we felt that adding identifying information would discourage others from 'peeking' at someone's screen in the large view if they knew that they could be identified. However, in practice, users commonly requested that people looking at their full view be identified; this is discussed more in a later section.

4.2.2. Communal feedback

Design justification: Owners need feedback from their community, i.e., who is present and any concerns they may have about what they can see.

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 it. That is, *reciprocity* of presence is enforced. As well, because people in a CB place are part of a social group, one can reasonably expect – security and social violations aside (there is no *access control* or *authorization* in CB) – 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 CB to warn the owners about inappropriate things that are being shared to help the owners preserve their privacy if needed, e.g., the group can *self-police*.

4.3. Managing interruptions and distraction

Finally, the Screen-Sharing item has secondary effects on *solitude*. Solitude is the component of privacy that relates to how a person controls interpersonal interactions. 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 to determine whether that person can be interrupted.

Part of how people control their interactions 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.

4.4. Discussion

Following our system design, we performed a theoretical analysis of privacy using Boyle's technique (2005) of systematically applying his framework to the Screen-Sharing item. Boyle's (2005) work includes an analytic method that helps reveal (perhaps tacit) privacy-related assumptions made during the design process of a collaborative technology. While it is beyond the scope of this paper to describe his methods and how we applied them to the Screen-Sharing item, we summarize several situations in which privacy problems can occur in spite of our safeguards. Additional examples of possible privacy problems are included in Table 1.

Sharing with intimate collaborators: Our assumption is that groups using the Screen-Sharing item will comprise people who trust and collaborate with each other. Thus no explicit precautions are included to prevent privacy problems such as deliberate abuses. We do not expect the Screen-Sharing item to work well in a competitive or malignant environment, except perhaps as a source of disinformation.

Sharing voluntarily: Another assumption 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.

Table 1
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 some-times 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

They may also disinform others (e.g., by raising and posting a seemingly work-oriented window, when they are actually pursuing personal activities), thus reducing the effectiveness of the system.

Primary privacy concern—inadvertent privacy violations: In a community of intimate collaborators, we assume that inadvertent privacy violations are 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. With the Screen-Sharing item, 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 CB internal server (that stores and transmits information) in order to gain information for competitive advantage.

Primary type of sensitive information—high-precision details: The Screen-Sharing item primarily uses image-manipulation techniques such as blur filtration, where we assume this is sufficient to obscure high-precision details (such as text from an email or IM 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 modest-sized text is the sensitive part of a shared screen, but can fail to preserve privacy when sensitive images or large-font text banners or visual symbols or even graphical layout remains identifiable. Perhaps a better approach would be to change the *accuracy* of the information being shared (e.g., by replacing a graphic with an icon, and text with dummy lorem ipsum text), rather than the image *precision*. Another option is to completely mask selected parts of the display, as done by Berry et al. (2005).

It is possible for someone to make a persistent copy of a shared screen image. We assume that information in our screen-sharing item is transient, but this may not be the case. Taking and saving screen snapshots of what is currently shown in the shared screen item is trivial. More sophisticated people could also 'hack' into the underlying distributed data structure (which is currently not protected) held by the server in order to automate screen recording. Thus people can have access to the information record for a far longer time than intended. They can even pass it on to

others to abuse, e.g., to calculate one person's performance (and thus promotion potential) by analyzing the temporal record for work efficiency and productivity over time. 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).

Reciprocity of artifact sharing is not enforced, i.e., one person can see other's shared screens while choosing not to share one's own. Our assumption is that the community would all be willing to share their displays, but this may not be the case. This 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 from those others who may not have to share the same level of detail.

Access control of artifact sharing is all or none. Our assumption is that if people choose to share their screens, they would be willing for all in the trusted CB group to see it. Real life is not generally so black and white. One may want to *reveal* information to only a subset of the group, either because they may not want others to see it (confidentiality) or because they know that others may not be interested (*solitude and distraction*).

Disinformation (intentionally inaccurate information) is not hard and *misinformation* (unintentionally inaccurate information) 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 have not received it or seen it yet.

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).

People can, of course, completely safeguard their privacy by deciding not show the Screen-Sharing item, or not to be involved in a CB group. To see what actually happens in practice, the next section describes how the Screen-Sharing item was actually used by two communities.

5. Preliminary evaluation

Goal 3. Evaluate how groups use screen sharing for artifact awareness, and the effectiveness of its privacy controls.

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 have to likely develop over time as part of a group's everyday practice. While the privacy analysis suggests some issues, it does not reveal what actually happens (and what becomes issues or successes) in the everyday world. Thus we decided to perform preliminary evaluations of Community Bar and the Screen-Sharing item through a field deployment to two groups.

Recruiting a group to use Community Bar and the Screen-Sharing item over a long period of time was difficult for pragmatic reasons; so 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. Shortly afterwards, we interviewed members of an outside commercial development team who were also using CB and the Screen-Sharing item.

In the following sections, we report on the experiences that the two different groups had while using the Screen-Sharing item. We stress that this is a preliminary evaluation, where results are suggestive rather than prescriptive. While our methods use qualitative approaches (observations, interviews, etc.), they are not highly rigorous. That is, we consider this evaluation a first exploration of system use, where we are looking for big effects rather than small nuances. Still, our findings are suggestive of what works and what does not, and what should be done in the next round of work on artifact awareness via screen sharing.

5.1. Internal research group

The group that was introduced to the Screen-Sharing item had already been voluntarily using CB for their own use on a daily basis for over a year. This group included ten graduate students, faculty, research assistants, and former members of the research laboratory. People were both co-located and distributed. Most laboratory 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 laboratory members also connected to CB from their distant work offices, one in the same city as the laboratory, 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; as a group with an established culture of CB use, they had already established a practice of balancing awareness provision with privacy for other media items.

The Screen-Sharing item was distributed to the CB users in our research group as an installation package downloadable from a web page. 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 (seven people appeared to share their screen on a regular basis). No usage data were logged by the system; instead, people were asked to email in descriptions of use and other comments as they used the system over a 2-week period, and three 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 laboratory 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 5 min 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 laboratory, 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 had not 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 was not 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 were not publicly visible, but he was not 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 (Fig. 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 did not 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 IM 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 through *self-policing*. 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 laboratory who were not comfortable sharing their desktop using the Screen-Sharing item. One member of the laboratory 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), as well as Boyle and Greenberg (2005) concerns about autonomy and social norms.

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 CB 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.

In summary, what is important about this group is that they had a long and successful culture of CB use. Critical mass already existed, and they saw the Screen-Sharing item as just another feature in the CB repertoire. There were few surprises in how this group adopted the Screen-Sharing item, for it complemented their existing activities. The validity of our design justification, expectations, and concerns were largely confirmed by the group’s actual practice.

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 laboratory 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 use.³

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. One core CB user, not included in the interviews, was a former member of our laboratory. He was the CB champion, where he suggested to the group that they try CB. Though he was also not directly connected to the CB or artifact-awareness project, he had been a heavy CB user in the laboratory (which he had more recently left).

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 as a whole had been using CB for about 3 weeks, and all our participants had been using CB for at least 2 weeks. While most members of the group had two monitors, only one person had a webcam. The four participants who we interviewed connected only 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 did not 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

³In particular, the company maintained a Bridgit server that anyone could connect to at any time for desktop conferencing or remote collaboration (see www2.smarttech.com/st/en-US/Products/Bridgit/). While Bridgit has some features similar to CB, it is intended to be used for focused collaboration, not awareness or casual interaction.

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 with 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 did not 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 who he was working closely with who 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. However, 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 could not be seen by others (see Berry et al., 2005 for how this technique can be done).

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 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 did not 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 was not 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 shared their screens only when they had something to share, with one commenting that the reason he did not 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.

5.3. *Use by the two groups*

While we were able to interview only 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 both groups used the Screen-Sharing item to help determine others’ availability, project a certain image of themselves, and engage in lightweight casual interaction. In particular, people in both groups found it valuable to be able to check what others were doing without actually having to walk over to their desks. They also liked being

able to share artifacts and activities without having to ask others to come over. For remote people, the Screen-Sharing item enabled them to easily check availability or share artifacts when they normally would not have been able to do so. Privacy did not seem to be that large a concern for either group (discussed more in Section 5.3.2).

5.3.1. *Key differences in use*

Yet, several differences in use between the two groups did emerge. We recognize that due to our methods of data collection (interviews and observations with ten participants vs. interviews with four), some of these differences may not necessarily be significant. We also realize that some of these differences could be due to the dominant tasks done by each group, e.g., researchers writing papers vs. programming. Still, we discuss the differences below to inform further studies that may refine our preliminary results

First, opportunistic interactions seemed to be initiated quite differently in the two groups. 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 seems to fit with the way IM is sometimes used in the workspace, albeit extended so that people could now collaborate over artifacts. 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 laboratory, whereas the industry group continued following the same social norms that existed in their workplace.

Second, 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 and 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. Also, only one member of the industry group had a webcam, whereas many members of the research group had one. Considering that every participant in Romero et al.'s (2007) study 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.

Third, the industry group reported fewer instances of focused collaboration through the Screen-Sharing item than the research group, as well as fewer opportunistic interactions involving coordinating 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 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. 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.

Finally, the two groups differed considerably in their culture of CB use over time. Our internal group was reasonably cohesive, had a long history of CB use, and had evolved their cultural practices of CB use for over a year. Incorporating the Screen-Sharing item was just a matter of appropriating it to fit within the ecology of these practices. The outside group, on the other hand, was less cohesive and was still evolving its CB practices. Thus the introduction of the Screen-Sharing item was somewhat more variable and idiosyncratic. The relationships between some group members were not as strong, and different people had different views of how CB was used. All this emphasizes the nature of CB and its screen-sharing item as a socio-technical system: its uses, successes, and failures are highly dependent on the sociality and cultural norms of the group and how it appropriates the technology over time for its own needs.

5.3.2. Revisiting privacy

Only four kinds of privacy problems across the groups were reported as a result of the Screen-Sharing item. The problems and most serious example incidents are listed below, along with how these problems would be considered in terms of Boyle and Greenberg (2005) and Boyle et al.'s (2009) privacy framework.

Inadvertent privacy infraction. A participant who was creating exam questions almost accidentally shared them, because he forgot that his screen was being shared. This is an example of inadvertent *confidentiality* violation due to *information sensitivity*.

Apprehension. A CB group member did not want to share his screen, though he remained part of the CB community during the study. The reason was that he did not want to make a bad impression in front of others. This is an example of a trade-off; on one hand a person is fulfilling *obligations* that are part of their *autonomy*, while balancing that with a desire for *refuge* that is part of one's *solitude*.

Dissociation. 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. Here we see a tension between *anonymity* and how it respects one's *solitude* vs. a desire for *reciprocity* that comprises the *mechanics of privacy*.

Role conflict. A participant working at an industrial site on product development initially shared confidential code at full clarity with others offsite. This is an example of *disembodiment* – a person becomes cut off from the (multiple) contexts of those people viewing him – which confounds self-appropriation and can lead to inadvertent privacy infractions (Boyle et al., 2009).

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. 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, few participants shared their screens from home when they were not doing work. Why did they

addition to using the privacy controls built into the Screen-Sharing item. In particular, having two monitors seemed to alleviate many concerns about screen sharing, as they could more easily segregate their public from private work. This suggests that people with one display may be more apprehensive about sharing their screen, and they may also experience more inadvertent privacy infractions due to less available display space.

6. Future work

To recap, our hypothesis was: ‘screen sharing is a critical component of distributed artifact awareness, because one’s screen contents, as mediated by privacy controls, let a person selectively indicate what digital activities they want their collaborators to see’. We implemented a single type of screen-sharing system for artifact awareness and privacy control (many other designs are possible) and we showed – at least in two cases – that the hypothesis holds in real situations. This serves as an existence proof that screen sharing for artifact awareness can work at least some of the time, although our claims are limited to favourable groups and work contexts.

Thus 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 CB 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 and Greenberg’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 CB is only one way of providing artifact awareness; different ways of presenting shared screens to people may eliminate some privacy and space-usage concerns (e.g., Berry et al., 2005) while still providing enough information for people to maintain awareness of what others are 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 adaptation that is extremely interesting but difficult to probe in short-term studies.

Finally, we recognize that screen sharing as a method for capturing and presenting artifact-awareness information has a limiting design bias—we are imitating the real-world affordances of visual work artifacts viewable by people inhabiting a common area. This offers a reasonable ‘first cut’ solution, as we know what mechanisms work in our everyday world. Yet it is also a pale imitation of what really happens. In the everyday world, it is easy, immediate, and intuitive for one to see who else in the group can see their work and at what level of detail. Similarly, it is easy to negotiate this—people greet each other through body language (e.g., by making eye contact) as well as tacit gestures that invite others into their space. People can quickly protect or reveal selected parts of their work from view, e.g., they protect by how they are shielding it with their bodies, by putting their hands over it, by covering it up, or by keeping viewers at a distance. In contrast, a shared screen is an all-or-none exposure of a single region for public display. Another problem is that much more is being shared in the real setting than just work information (e.g., the sights, sounds, and smells of the physical environment). All these nuances and subtleties are certainly not captured by CB and our Screen-Sharing items, and indeed it would be very difficult to produce a design that incorporates such social and environmental cues. Perhaps adding context-aware interruptability cues

could help somewhat (e.g., on phone, do not disturb, etc.). Yet another issue is that our Screen-Sharing item assumes that the artifacts valuable for awareness are those on the personal screen, which in turn reflects an individual's work efforts and products. However, for some groups it may be the common work product being crafted by the group that is more valuable, e.g., how one's personal contributions affect the current state of a command and control system, or how it affects a commonly shared dataset. In addition, artifacts not on the screen – such as paper – are excluded from view.

At the same time, we need not be shackled by this real-world 'being-there' view. As Hollan and Stornetta (1992) note, we can go to 'beyond being there', where a deep understanding of the role artifact awareness plays in serving the communicative needs of people can lead to quite different media and mechanisms than mere screen sharing. For example, our approach works only if participants are tele-workers at their computers using standard GUIs, with the screen-sharing mechanism on continual display. Yet people are becoming increasingly mobile, where they are using alternate technologies and media to stay in touch: cell phones, video broadcasts, social networks, twitter feeds, and so on. The question – and another avenue for future research – is how we can apply the notion of artifact awareness to these new settings and new media. Undoubtedly the solution will go beyond screen sharing.

7. Conclusion

In this article, we have described the design, implementation, and preliminary evaluation of an awareness tool that uses screen sharing within the CB 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, iCORE, and our industrial partners: SMART Technologies and Microsoft. We thank our CB participants, who let us interview them about their CB practices. We also thank our anonymous referees, who provided insightful comments that helped us improve upon an earlier draft of this work.

References

- Bartle, R., 2004. *Designing Virtual Worlds*. New Riders.
- Bellotti, V., Bly, S., 1996. Walking away from the desktop computer: distributed collaboration and mobility in a product design team. In: *Proceedings of ACM CSCW*, pp. 209–218.
- Berry, L., Bartram, L., Booth, K., 2005. Role-based control of shared application views. In: *Proceedings of ACM UIST*, pp. 25–32.
- Birnholtz, J.P., Gutwin, C., Hawkey, K., 2007. Privacy in the open: how attention mediates awareness and privacy in open-plan offices. In: *Proceedings of ACM GROUP*, pp. 51–60.
- Birnholtz, J.P., Gutwin, C., Ramos, G., Watson, M., 2008. Open-Messenger: gradual initiation of interaction for distributed workgroups. In: *Proceedings of ACM CHI*, pp. 1661–1664.
- Bly, S.A., Harrison, S.R., Irwin, S., 1993. Media spaces: bringing people together in a video, audio, and computing environment. *Communications of the ACM* 36 (1), 28–46.
- Boyle, M., 2005. *Privacy in media spaces*. Ph.D. Thesis, University of Calgary, Calgary, Alberta, Canada.
- Boyle, M., Greenberg, S., 2005. The language of privacy: learning from video media space analysis and design. *ACM Transactions on Computer-Human Interaction* 12 (2), 328–370.
- Boyle, M., Neustaedter, C., Greenberg, S., 2009. Privacy factors in video-based media spaces. In: Harrison, S. (Ed.), *Media Space: 20+ Years of Mediated Life*. Springer, New York, pp. 99–124.
- Buxton, W., Fitzmaurice, G., Balakrishnan, R., Kurtenbach, G., 2000. Large displays in automotive design. *Computer Graphics and Applications, IEEE Computer Society* 20 (4), 68–75.
- Cadiz, J.J., Venolia, G., Jancke, G., Gupta, A., 2002. Designing and deploying an information awareness interface. In: *Proceedings of ACM CSCW*, pp. 314–323.
- Clark, H., 1996. *Using Language*. Cambridge University Press, Cambridge.
- Dourish, P., Bly, S., 1992. Portholes: supporting awareness in a distributed work group. In: *Proceedings of ACM CHI*, pp. 541–547.
- Dix, A., Finlay, J., Abowd, G., Beale, R., 1993. *Human-Computer Interaction*. Prentice-Hall, Englewood Cliffs, NJ.
- Endsley, M., 1995. Toward a theory of situation awareness in dynamic systems. *Human Factors* 37 (1), 32–64.
- Engelbart, D., English, W., 1968. A research center for augmenting human intellect. In: *Proceedings of AFIPS Fall Joint Computer Conference*, pp. 395–410.
- Fish, R.S., Kraut, R.E., Root, R.W., Rice, R.E., 1992. Evaluating video as a technology for informal communication. In: *Proceedings of ACM CHI*, pp. 37–48.
- Greenberg, S., 1999. Designing computers as public artifacts. *International Journal of Design Computing, Special Issue on Design*

- Computing on the Net, University of Sydney, November 30–December 33.
- Greenberg, S., Gutwin, C., Roseman, M., 1996. Semantic telepointers for groupware. In: Proceedings of the Sixth Australian Conference on Computer–Human Interaction—OZCHI'96, November 24–27, Hamilton, New Zealand, IEEE Computer Society Press, pp. 54–61.
- Greenberg, S., Rounding, M., 2001. The notification collage: posting information to public and personal displays. In: Proceedings of ACM CHI, pp. 514–521.
- Greenberg, S., Roseman, M., 2003. Using a room metaphor to ease transitions in groupware. In: Ackerman, M., Pipek, V., Wulf, V. (Eds.), *Sharing Expertise: Beyond Knowledge Management*. MIT Press, Cambridge, MA, pp. 203–256 January.
- Grudin, J., 2001. Partitioning digital worlds: focal and peripheral awareness in multiple monitor use. In: Proceedings of ACM CHI, pp. 458–465.
- Gutwin, C., Greenberg, S., 2002. A descriptive framework of workspace awareness for real-time groupware. *Computer Supported Cooperative Work. The Journal of Collaborative Computing* 11 (3–4), 411–446 (Special Issue on Awareness in CSCW, Kluwer Academic Press).
- Gutwin, C., Greenberg, S., Blum, R., Dyck, J., Tee, K., McEwan, G., 2008. Supporting informal collaboration in shared-workspace groupware. *Journal of Universal Computing* 14 (9), 1411–1434.
- Gutwin, C., Stark, G., Greenberg, S., 1995. Support for workspace awareness in educational groupware. In: Proceedings of the First Conference on Computer Supported Collaborative Learning (CSCL'95). LEA Press, Indiana, USA, pp. 147–156 October 17–20.
- Heath, C., 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), 24–48.
- Hollan, J., Stornetta, S., 1992. Beyond being there. In: Proceedings of ACM CHI, pp. 119–125.
- Hutchins, E., 1995. *Cognition in the Wild*. MIT Press, Cambridge, MA.
- Isaacs, E., Walendowski, A., Whittaker, S., Schiano, D.J., Kamm, C., 2002. The character, functions, and styles of instant messaging in the workplace. In: Proceedings of ACM CSCW, pp. 11–20.
- Kaasten, S., Greenberg, S., Edwards, C., 2002. How people recognize previously seen web pages from titles, URLs and thumbnails. In: Faulkner, X., Finlay, J., Detienne, F. (Eds.), *People and Computers XVI* (Proceedings of HCI BCS Conference Series). Springer, New York, pp. 247–265.
- Kraut, R., Egido, C., Galegher, J., 1988. Patterns of contact and communication in scientific research collaboration. In: Proceedings of ACM CSCW, pp. 1–12.
- Kurlander, D., Skelly, T., Salesin, D., 1996. Comic chat. In: Proceedings of ACM Computer Graphics and Interactive Techniques, pp. 225–236.
- Lauwers, J.C., Lantz, K.A., 1990. Collaboration awareness in support of collaboration transparency: requirements for the next generation of shared window systems. In: Proceedings of ACM CHI, April, pp. 303–311.
- Lederberg, J., Uncapher, K., 1989. Towards a national collaboratory. The Rockefeller University, March 17–18. Also available from the National Science Foundation.
- McEwan, G., Greenberg, S., 2005. Supporting social worlds with the Community Bar. In: Proceedings of ACM GROUP, pp. 21–30.
- McEwan, G., Greenberg, S., Rounding, M., Boyle, M., 2006. Groupware plug-ins: a case study of extending collaboration functionality through media items. In: Proceedings of CollabTech, IPSJ SIG Groupware and Network Services, pp. 42–47.
- Nardi, B.A., 1993. *A Small Matter of Programming: Perspectives on End User Computing*. MIT Press, Cambridge, MA.
- Nardi, B.A., Whittaker, S., Bradner, E., 2000. Interaction and out-eraction: instant messaging in action. In: Proceedings of ACM CSCW, pp. 79–88.
- Neustaedter, C., Greenberg, S., Boyle, M., 2006. Blur filtration fails to preserve privacy for home-based video conferencing. *ACM Transactions on Computer–Human Interaction* 13 (1), 1–36.
- Olson, G.M., Atkins, D.E., Clauer, R., Finholt, T.A., Jahanian, F., Killeen, T.L., Prakash, A., Weymouth, T., 1998. The Upper Atmospheric Research Collaboratory (UARC). *Interactions* 5 (3 May), 48–55.
- Porter, J., 2008. *Designing for the Social Web*. New Riders.
- Richardson, T., Stafford-Fraser, Q., Wood, K.R., Hopper, A., 1998. Virtual network computing. *IEEE Internet Computing, IEEE Educational Activities Department* 2 (1), 33–38.
- Romero, N., McEwan, G., Greenberg, S., 2007. A field study of Community Bar: (mis)-matches between theory and practice. In: Proceedings of ACM GROUP, pp. 89–98.
- Root, R.W., 1988. Design of a multi-media vehicle for social browsing. In: Proceedings of ACM CSCW, pp. 25–38.
- Rounding, M., 2004. *Informal awareness and casual interaction with the Notification Collage*. M.Sc. Thesis, University of Calgary, Calgary, Alberta, Canada.
- Sarin, S., 1984. *Interactive on-line conferences*. Ph.D. Thesis, MIT, Cambridge, Massachusetts, United States.
- Smale, S., Greenberg, S., 2005. Broadcasting information via display names in instant messaging. In: Proceedings of ACM GROUP, pp. 89–98.
- SMART Technologies, Inc., 2009. SMART Sync classroom management software. <www.smarttech.com/sync> (Retrieved March, 2009).
- Sohlenkamp, M., Chwelos, G., 1994. Integrating communication, cooperation, and awareness: the DIVA virtual office environment. In: Proceedings of ACM CSCW, pp. 331–343.
- Stefik, M., Bobrow, D.G., Foster, G., Lanning, S., Tatar, D., 1987. WYSIWIS revised: early experiences with multiuser interfaces. *ACM Transactions on Office Information Systems* 5 (2), 147–167.
- Summers, R., 1998. *Official Microsoft NetMeeting Book*. Microsoft Press.
- Tang, J.C., Rua, M., 1994. Montage: providing teleproximity for distributed groups. In: Proceedings of ACM CHI, pp. 37–43.
- Tee, K., Greenberg, S., Gutwin, C., 2006. Providing artifact awareness to a distributed group through screen sharing. In: Proceedings of ACM CSCW, pp. 99–108.
- Tee, K., 2007. *Artifact awareness for distributed groups through screen sharing*. M.Sc. Thesis, University of Calgary, Calgary, Alberta, Canada.
- Voida, A., Grinter, R.E., Ducheneaut, N., Edwards, W.K., Newman, M.W., 2005. Listening in practices surrounding iTunes music sharing. In: Proceedings of ACM CHI, pp. 191–200.
- Whittaker, S., Frohlich, D., Daly-Jones, O., 1994. Informal workplace communication: what is it like and how might we support it? In: Proceedings of ACM CHI, pp. 131–137.
- WOS Data Systems, 1987. *Timbuktu: The Next Best Thing To Being There, User's Guide (Version 1)*.