Neustaedter, C., Greenberg, S. and Boyle, M. (2006). Blur Filtration Fails to Preserve Privacy for Home-Based Video Conferencing. ACM Transactions on Computer Human Interactions (TOCHI), 13, 1, March, p1-36.

Blur Filtration Fails to Preserve Privacy for Home-Based Video Conferencing

CARMAN NEUSTAEDTER, SAUL GREENBERG, and MICHAEL BOYLE
University of Calgary

Always-on video provides rich awareness for distance-separated coworkers. Yet video can threaten privacy, especially when it captures telecommuters working at home. We evaluated video blurring, an image masking method long touted to balance privacy and awareness. Results show that video blurring is unable to balance privacy with awareness for risky situations. Reactions by participants suggest that other popular image masking techniques will be problematic as well. The design implication is that image masking techniques will not suffice for privacy protection in video-based telecommuting situations. Other context-aware privacy-protecting strategies are required, as illustrated in our prototype context-aware home media space.

Categories and Subject Descriptors: H.5.1 [**Information Interfaces and Presentation**]: Multimedia Information System—*Evaluation/methodology; Video*; H.5.3 [**Information Interfaces and Presentation**]: Group and Organization Interfaces—*Collaborative computing; computer-supported cooperative work; evaluation/methodology; synchronous interaction*; J.4 [**Social and Behavioral Sciences**]: Psychology; Sociology; K.4.1 [**Computers and Society**]: Public Policy Issues—*Privacy; Regulation*

General Terms: Design, Experimentation, Human Factors, Security

Additional Key Words and Phrases: Telecommuting, awareness, video conferencing

1. INTRODUCTION

Throughout a typical day, coworkers naturally converse and interact among each other in what is known as *casual interaction*—the frequent and informal encounters that occur when people meet serendipitously or that are initiated when one person seeks another [Fish et al. 1993; Hudson and Smith 1996]. Casual interactions are important, for they foster knowledge and help individuals accomplish both individual and group work [Kraut et al. 1988; Fish et al. 1993]. *Informal awareness*—an understanding of who is around and available for interaction—holds casual interaction together by helping people decide if and when to smoothly move into and out of conversation and collaboration

We are grateful to NSERC and Microsoft Research for their partial funding of this research.
Authors' address: University of Calgary, 2500 University Drive NW, Calgary, AB, Canada; email: {carman,saul,boylem}@cpsc.ucalgary.ca.

[Kraut et al. 1988; Bellotti and Sellen 1993]. Informal awareness is easily gained when people are in close physical proximity, but deteriorates over distance [Kraut et al. 1988; Greenberg 1996]. As a result, casual interaction suffers when coworkers are distributed.

One possible (and popular) solution for providing awareness between distance-separated collaborators is to use an always-on video link to connect remote locations [Mantei et al. 1991; Fish et al. 1993; Bly et al. 1993; Tang et al. 1994; Bellotti 1996; Greenberg 1996; Lee et al. 1997; Greenberg and Kuzuoka 2000]. Always-on video within a media space provides rich awareness, yet it also broadcasts information that individuals may judge as privacy sensitive [Bellotti and Sellen 1993; Bly et al. 1993; Bellotti, 1996, 1998; Hudson and Smith 1996; Greenberg and Kuzuoka 2000; Boyle et al. 2000; Palen and Dourish 2003]. The direct benefit of the video is awareness, as well as an easy and natural way to move from awareness into opportunistic conversation and interaction over the same media channel. As with many other researchers our concern in this article is about the awareness that precedes conversation, especially how awareness competes with privacy. Thus, a common goal of many researchers in computer supported cooperative work [CSCW] over the last decade has been to find a balance between the rich awareness provided by video-based media spaces and the privacy concerns they raise.

In practice, video media spaces have found some limited success in office situations, primarily at research sites [Fish et al. 1993; Mantei et al. 1991; Jancke et al. 2001]. Most installations simply ignore privacy issues: risks are fairly low in office settings, and simple privacy safeguards often suffice, for example, people can explicitly switch off the video channel, or turn the camera around to face a wall. Yet the situation becomes complicated when people choose to work from home as telecommuters while still desiring close contact with colleagues at work. The important challenge facing telecommuters' use of such a *home media space*—defined as an always-on video media space used within the home to connect to the office—is that privacy risks increase drastically for the telecommuter. The home is not the office; activities and appearances appropriate in the home are often inappropriate when viewed in an office environment by a colleague. For example, consider these following home media space situations—all derived from actual events reported to us by telecommuters—where the telecommuter habitually uses always-on video to provide a colleague at work with awareness. These situations are all very realistic and threaten the privacy of the telecommuter, as well as others within the home.

First, imagine yourself as a telecommuter living the dual role as worker and as home occupant. It is a hot day, and you are going shirtless or wearing a very revealing top, such as a bra. Forgetting you are shirtless (because this is not a problem at home), you enter your home office to quickly check your email. Your colleagues—at the office, wearing typical business attire in their air conditioned environment—see you shirtless and you are left feeling apprehensive about whether they perceive your appearance as being inappropriate.

Our second example results from a telecommuter's unconscious acts, their ease of forgetting that their distant colleague is (virtually) sitting right across from them, and from the lack of feedback that they are actually in a public

setting. Imagine yourself again as the telecommuter working at your home computer when you suddenly sneeze. Naturally, you proceed to blow your nose, forgetting that a camera on top of your monitor captures this at a very close range. You next begin to pick your nose at great length. Your colleague views the scene over the video link and is disgusted at how inconsiderate you are being. While this could also happen at work, the close proximity of other office colleagues in an open office, or an open door to a hallway would subconsciously remind one that they are visible to others and they would likely be more surreptitious or subtle at these overt acts.

Third, family members and friends in the home likely gain no benefit from the video link yet still incur its privacy threat. You, as the telecommuter, are working in your home office in the early morning when your spouse (who has just woken up) enters the room wearing only pajamas (or perhaps even less) and gives you a passionate kiss. All this is captured on camera, and seen by your colleague. Your spouse realizes this and becomes infuriated, telling you never to use the camera again.

Our fourth example results from the dual purposes typical of most home offices. As with many home offices, your home office is also your spare bedroom. One hot day you take a shower in the bathroom next door. You towel off, and then go into the spare bedroom to retrieve a bathrobe in the closet. A few moments after entering the room, you realize that the camera is directed at you. You drop to the floor, crawl to the camera, and knock it off the computer.

The final example results again from the dual purposes typical of most home offices. You have a guest from out of town staying in the spare bedroom. You were using the media space during the day in this room, yet forgot to turn it off after finishing work. Your guest enters the room in the evening and undresses for bed. She is completely oblivious to the capturing camera and does not realize that your colleague has just seen (and continues to see) her in a precarious situation.

These issues are very real. Video cameras and media spaces are rapidly moving out of the hands of office workers and researchers, and into everyday use in many settings—home and otherwise. With the declining cost of PC cameras and a number of companies offering free video conferencing or webcam software (e.g., Webcam for MSN Messenger, Yahoo! Messenger), video is increasingly being used in homes. Its prevalence can easily be seen by simply surfing the web for sites showing live webcam footage. While these home uses of video may not be the specific case we are looking at, it demonstrates a need for techniques to mitigate privacy concerns for home users of video conferencing technology. Whether home users of PC cameras know it or not, as soon as they turn on their camera, they are placing not only their own privacy at risk, but the privacy of others in the home as well.

In an effort to help mitigate privacy concerns over video links, other researchers have studied video obfuscation techniques such as *distortion filters*: algorithmic reduction of image fidelity to hide sensitive details in a video image [Hudson and Smith 1996; Zhao and Stasko 1998; Greenberg and Kuzuoka 2000; Boyle et al. 2000; Crowley et al. 2000]. Example filtration techniques include *pixelize filters* that produce a mosaic of solid rectangles; *blur filters* that

average pixel values to produce a blurred effect; *subtraction filters* that remove the static background of a scene; and, *eigen-space filters* that reconstruct a scene with benign information.

Underlying these techniques is an assumption that privacy and awareness trade off against each other [Hudson and Smith 1996; Boyle and Greenberg 2005]. For example, if two individuals are using a media space and one uses a particular technique to gain more privacy, there is an assumption that the other will gain less awareness. For the most part this is the case, however, there are situations where this is not a straight trade-off. At times revealing more information, thus providing more awareness for another, can increase one's privacy because the chance of being interrupted at an inappropriate time diminishes. For example, if you are reading intently and a colleague knows this, she may be more inclined to not interrupt you while you are concentrating. This privacy-awareness trade-off motivates many of the video obfuscation techniques; while privacy and awareness are not always competing measures, video obfuscation is still widely put into practice and thought to be effective for balancing privacy and awareness.

In spite of the proliferation of video obfuscation techniques, very few of them have been studied to see whether they do indeed safeguard privacy. Zhao and Stasko [1998] examined what awareness information people can see in video filtered by several different techniques, but did not question how well they safeguarded privacy. More recently, Boyle et al. [2000] specifically examined the awareness/privacy tradeoff. In particular, they studied the blur and pixelize filters at different levels of obfuscation; they wanted to see how each filter at various levels balanced privacy and awareness in mundane and benign office situations, for example, people working or reading, people chatting, people eating lunch. They found that each filter offered an obfuscation level that adequately preserved privacy and still provided awareness for these benign office situations. The blur filter, however, was found to balance privacy and awareness over a wider range of filtration levels than the pixelize filter.

However, Boyle et al. [2000] did not study the effects of their distortion filters on situations where one's privacy may be at moderate to extreme risk, such as those typified in our home media space telecommuting examples. Intuitively, we expect to find that equivalent or higher blur levels are needed as risk increases. Yet it is not clear if there will still be blur levels that provide an overlap between privacy protection and awareness. This is an important question, for no overlap implies that filtration techniques, in spite of their popularity, may be a questionable method for safeguarding privacy. Consequently, we set ourselves two research goals.

1. Determine how well video-blurring safeguards privacy in always-on video links that connect home-based telecommuters with office colleagues.
2. Based on the above outcome, revisit how privacy and awareness can be balanced.

This article mostly concerns the first goal. We study people's perceptions and reactions to blurred video scenes, and we gather their feedback about video and

video obfuscation in general. Our results allow us to draw conclusions about the effectiveness of video blurring to safeguard privacy, which we then generalize to other video obfuscation techniques.

To foreshadow what is to come, we constructed a controlled experiment to test blur filtration with a set of scenes that typify home telecommuting, where scenes range greatly in perceived privacy risk. Scenes included: mundane situations such as working at a computer, moderately risky situations such as the telecommuter kissing her partner, and extremely threatening situations such as being shown completely naked. Since awareness is the direct benefit of a home media space as substantiated by many researchers [Kraut et al. 1988; Hudson and Smith 1996; Boyle et al. 2000], we look at both privacy and awareness in our experiment. Thus, we assume that privacy and awareness are competing measures that must be balanced. While our first goal is an incremental study that uses a similar methodology to Boyle et al. [2000], there is a major difference between our study and theirs. Specifically,

- we are testing video usage in a home media space setting rather than an office setting, and
- we explore scenes that are much more threatening to one's privacy than everyday mundane situations that occur at an office.

This inclusion of risk is critical. Even if we assume that accidental high-risk exposures in front of a home telecommuting camera may be rare (such as nudity, sexual encounters), even a single occurrence or the threat of exposure is enough to undermine current and future use of a home media space. As we will see, our study outcome strongly suggests that:

- blur filtration *is not sufficient* for balancing privacy and awareness in home media space situations, which in turn raises serious question about other video obfuscation techniques.

This is an important contribution, for it questions the premise behind much prior work in CSCW [Hudson and Smith 1996; Zhao and Stasko 1998; Crowley et al. 2000; Boyle et al. 2000].

Based on this result, our second goal is to identify a set of design implications that articulate the difficulties with designing privacy-protecting strategies for home-based media spaces, and suggest alternate strategies that may be appropriate for balancing privacy and awareness. While this goal is not the major focus of this article, we will contribute several critical design considerations to help guide practitioners in future designs of home-based media spaces. From this, we then briefly outline the design of our own context-aware home media space that uses sensing-based technology as a tool for regulating privacy.

The next section of this article outlines the study's methodology and includes the specific research questions the study answers. Following this, we discuss the study results and its implications for the design of home-based media spaces.

2. METHODOLOGY

Because there is as yet no common technique for observing and analyzing the tradeoffs between awareness and privacy, this section includes sufficient detail of our methodology to allow others to replicate, vary, or critique it.

In our study, participants are asked to adopt the role of an office worker where they are the close-working colleague of a telecommuter. Participants then view a series of five video scenes—each blurred at ten different levels of blur—containing the telecommuter in a situation that varies by perceived privacy risk. Each level of blur, or blur level, is characterized by how much an image is blurred at that particular level, described in detail in Section 2.6. For each blur level, participants answer privacy and awareness questions, described in Section 2.7. We first describe our research questions and variables, followed by our materials and procedure.

2.1 Research Questions

The study answers three main research questions.

***Question* 1:** At what blur levels are participants able to identify who is in the scene, what they are doing, and what they are wearing?

***Question* 2:** At what blur levels is privacy adequately preserved for the telecommuter and others in the home?

***Question* 3:** What blur levels do participants choose in order to make a given scene appropriate for a colleague to view?

The first two questions evaluate blur filtration's effectiveness at balancing privacy and awareness: Question 1 identifies the blur levels that provide adequate levels of awareness, while Question 2 identifies the blur levels that provide adequate levels of privacy. For Question 2, participants must decide for a given scene and blur level if too much or too little information is being revealed. For a particular scene, if the blur levels that preserve privacy (Question 2) fall in the range of blur levels that provide awareness (Question 1), then blur filtration is able to balance privacy and awareness using the found blur levels. The third question looks at what blur level people would predicatively choose to use for a given situation. Participants are also given the option to choose no blur levels, where they can simply opt to turn the camera off. For Question 3, participants must decide for a particular scene what blur level they are comfortable in having a colleague see them at.

2.2 Independent and Dependent Variables

The independent variables for the study are scene type (5) $\times$ blur filter levels (10), and are discussed in detail in Sections 2.3 and 2.6 respectively. The dependent variables recorded in a during-test questionnaire (described in Section 2.7) are:

- a participant's ability to correctly identify awareness cues,
- a participant's confidence in identifying awareness cues,

- a participant's perception of the videos' level of privacy threat, and
- the chosen blur level for safeguarding each video.

2.3 Materials: Video Scenes

We recorded five video scenes that vary in the level of risk presented, from scenes we judged to have no risk to those with very high risk. We base risk on the severity of the consequences imagined by those apprehensive of the situation. Each scene shows a telecommuter performing a different activity or shows the telecommuter with a different appearance. All scenes are recorded from the same point of view: behind the computer monitor at the same angle. Each scene was recorded twice, once with a paid male, and again with a paid female acting as the telecommuter. As illustrated in Figure 1 and described below, scenes are sorted by our own expectations of privacy risk—from low risk to high risk—that we later validate with our subjects' actual perception of privacy risk, as detailed in Sections 2.5 and 3.1.

1. ***Working at a computer*** (Figure 1a): The telecommuter is working at a computer while wearing clothes appropriate for both home and the office (Low risk).
2. ***Picking one's nose*** (Figure 1b): The telecommuter is working at a computer wearing clothes appropriate for both home and the office when he/she begins to pick his/her nose (Moderate risk).
3. ***Working with no shirt on*** (Figure 1c): The telecommuter is working at a computer with no shirt on (Moderate risk).
4. ***Kissing a partner*** (Figure 1d): The telecommuter is working at a computer when his/her partner enters the room, kisses the telecommuter intimately, and leads him/her out of the room (Moderate risk).
5. ***Changing clothes / Naked*** (Figure 1e): The telecommuter enters the room in a robe, is shown completely naked, and then puts on underwear (High risk).

We made most scenes quite sensitive to risk, as Boyle et al.'s study [2000] had already evaluated nonrisky scenes. However, we did not want to fall into the trap of giving our subjects scenes that could be interpreted as pornography, nor did we feel it was appropriate to show illegal or morally poignant acts. Our goal was to emulate what people normally see over a video link: ordinary people who we could imagine were our colleagues and would not want to see in compromising situations. Thus, the actors were deliberately chosen to be middle-aged individuals with the appearance of working professionals doing otherwise "normal" home activities.

Videos were captured using high resolution: 720×480 pixels at 30 frames per second (fps) DV-format. While this is better than today's PC cameras, we expect at least this level of quality in future hardware. We used normal lighting as this mirrors what would happen in the home. By way of comparison, Boyle et al.'s [2000] study used Intel Indeo^{TM} compressed videos at 176×144 pixels and 24 fps.

Fig. 1. The five (unfiltered) video scenes typifying home situations facing a telecommuter. Participants did not see the black bars for the Changing scene—the telecommuter was shown completely nude.

2.4 Materials: Scenarios Provided to Participants

In the everyday world, peoples' perception of what comprises a privacy violation in a particular home media space situation will vary because of individual differences. For example, we expect people will differ in their willingness to give out personal information to others, and in their desire for more (or less) privacy. These differences can be extreme; while some people are adamantly against video, others host freely accessible Internet sites allowing the public a video glimpse of their home activities at any time of the day or night. Similarly, peoples' reactions will depend heavily on the relationship they have to the distant telecommuter, for example, a lover *vs.* a family member *vs.* a co-worker *vs*. a stranger. To normalize this, we asked participants to imagine they were in a telecommuting scenario, described below, that set the context of how they would use the video link. In the scenario, they are a colleague of a telecommuter, either Larry (for males) or Linda (for females), and have a real need and desire to work closely with this individual:

> *"Here is a picture of your work colleague Larry [or Linda—a picture is shown of only their face]. You have known Larry for more than a year now and have a close working relationship with him. It is easy to see when Larry is around and working because he is in the office next to you. Throughout the day you talk to Larry very frequently and often you will be working together on a project. To better manage his family, Larry has decided to work from home two days of the week. You both still really want to work closely together so you and Larry decide to set up a video link between Larry's home and your office. The video link mostly captures you both working, but occasionally it captures Larry doing other things at home and sometimes you see his family members because the room doubles as a spare bedroom. Today you are working at your office and Larry is working from home. You have a question to ask Larry and decide to look at the video link to see if he is busy. . ."*

2.5 Materials: Assessing the Risk of Each Scene

The level of risk presented in each scene was assessed prior to the study using Boyle's theory of privacy in video media spaces [Boyle and Greenberg 2005] as applied to our telecommuting scenario. Boyle's theory states that privacy can be violated in one or more of three fundamental ways: a breech of confidentiality, invasion of solitude, or loss of autonomy.

- ***Solitude*** is freedom from *interruption* and *distraction*. Solitude can be invaded if someone attempts to interact with another at an inappropriate time or simply causes unwanted distraction.
- ***Confidentiality*** is control over who knows *what* about you and at what level of detail. Confidentiality is breeched when media space participants lose this control or when someone learns more about the person than is desired.
- ***Autonomy*** is the control over defining oneself and can be lost when a media space participant is no longer able to choose *how* and *when* he/she participates in the space.

As Boyle and Greenberg [2005] discuss, these three aspects of privacy are tightly coupled and a violation of one often leads to violations of another. While these privacy violations may occur because of the media space's design, participants in a media space may also fail to *appropriate* themselves correctly for their current situation [Bellotti 1998], for example, create a socially acceptable behavior or appearance. We now use these potential violations to describe each scene and assess its privacy risk. These assessments are validated with a post-test questionnaire, to be discussed in Section 3.1.

In scene 1, Linda (or Larry) is working at a computer while wearing clothes appropriate for both home and the office, for example, clean and casual clothes (Figure 1a). This scene is representative of a mundane activity that one would typically perform when working at home, for example, reading the newspaper, checking email. There are no immediately obvious ways for Linda's autonomy, confidentiality, or solitude to be violated; we assess the scenes as little to no risk.

In scene 2, Larry (or Linda) is working at a computer wearing clothes appropriate for both home and the office when he begins to pick his nose (Figure 1b). This scene is representative of an unconscious act, for example, scratching yourself, blowing your nose, or other grooming activities. Larry would have liked to pick his nose without others seeing, yet he failed to receive feedback of the media space capturing his act and as a result, he failed to correctly appropriate himself for others to see. The media space has threatened his autonomy because he has lost control over how he is being recorded. One can expect that all people pick their nose at some point or another, yet now viewers have actually seen Larry perform this act at close quarters. We feel this is a mild breech of confidentiality. Many would tolerate seeing this accidentally at a distance, yet now they are seeing a close-up, somewhat disgusting view of it. We assess this situation as a moderate risk.

In scene 3, the telecommuter is working at a computer wearing no shirt (Figure 1c)—Larry is shown bare-chested for this scene, while Linda is shown in a bra. This scene is representative of situations where one may be working on a hot day or had to quickly check email while lounging around the home in "comfortable" attire. Perhaps Linda thinks she is alone when in fact she is not, causing her to have a misconception of her solitude. This solitude violation causes a breech of confidentiality as viewers now know what type of bra she wears and perhaps even the size of her chest (Larry: amount of chest hair, level of muscular build). The media space is also threatening Linda's autonomy because it forces her to choose between being comfortable in her own home and collaborating with her colleagues. Linda is appropriately dressed for the privacy of her home, but definitely not for an office. We assess this scene as a moderate risk.

In scene 4, Larry is working at a computer when his spouse, Linda, enters the room (roles are reversed for the alternate video scene). The two intimately kiss and Linda leads Larry out of the room (Figure 1d). This scene is representative of intimate interpersonal activities between inhabitants of a home, for example, disciplining your children, arguing with your spouse, or showing affection. Larry likely wishes to engage in this activity, but there is no way he wants others to see. The media space fails to give Larry adequate feedback that this activity

is being captured and Larry's solitude is violated as a result. His confidentiality is also being breeched because he does not want his coworkers to know the details of his personal life, yet now they know that he is about to partake in sexual activities with Linda. This leads to an autonomy violation as Larry is forced to choose between dealing with his spouse's needs for affection and collaborating with his colleague. Linda is also experiencing similar threats to her solitude, autonomy, and confidentiality as she becomes subject to the video media space, although these are even worse because Linda gains no benefit from the video connection. While both Larry and Linda are appropriate in appearance, this behavior is not appropriate for office environments and would not normally be seen by others (albeit such affairs may occur behind closed office doors). We assess this scene as a moderate risk, yet now it is more severe than the previous scene because people other than just the telecommuter are affected.

In scene 5, Linda (or Larry) walks into the room wearing only a housecoat. She takes off the housecoat, reveals full-frontal nudity, and then begins to dress (Figure 1e). This scene is representative of situations where one clearly does not want colleagues looking. While Linda wants to dress, she does not want others to view this act nor her naked body—clearly her autonomy is violated. Linda is facing the same solitude, confidentiality, and autonomy violations as the telecommuter in Scene 3. In this case, it is even worse: by changing and by being naked there is no way she is appropriated correctly for an office environment. As such, this scene is much more severe and constitutes a level of high risk.

While not previously mentioned, each scene does have the potential to violate the solitude of the telecommuter if the viewer distracts or interrupts the telecommuter at an inappropriate time, for example, if the telecommuter is busy working, or is interrupted at an embarrassing time. We feel, however, people viewing each scene at full fidelity are capable of determining if the time is appropriate to move into interaction. Our scene assessment did not look at background information in the scene as we hypothesize this will have little effect on privacy violations.

2.6 Materials: Blurred Video Scenes

The ten video scenes (five male, five female) were preprocessed to create a set of videos at each of the ten different blur levels to be evaluated (Figure 2). We used the same algorithm to blur our images as Boyle et al. [2000] and our blur levels are roughly equivalent. Our distortion algorithm computes a filtered pixel's value as the unweighted average of itself and its neighbors. For a given image, the larger the neighborhood, the greater the perceived distortion. For example, blur level 2 uses a neighborhood of 230 x 153 pixels for blurring (Figure 2). This means that for a video blurred at level 2, each pixel is averaged with the neighbouring 230×153 pixels. This is a typical method for smoothing (blurring) an image.

The blur levels are shown in Figure 2, yet the images do not accurately portray what people see, due to reproduction quality, size reduction, and most importantly because no motion is visible in these still images. Readers should not second guess what is visible from these photos; the videos themselves are available by contacting the authors.

Fig. 2. The ten blur levels evaluated in the study (currently showing the Working scene with the male actor) and the corresponding size of the pixel neighbourhood used for blurring. Reproduction quality and a lack of motion may cause these images to appear blurrier than the videos used in the study.

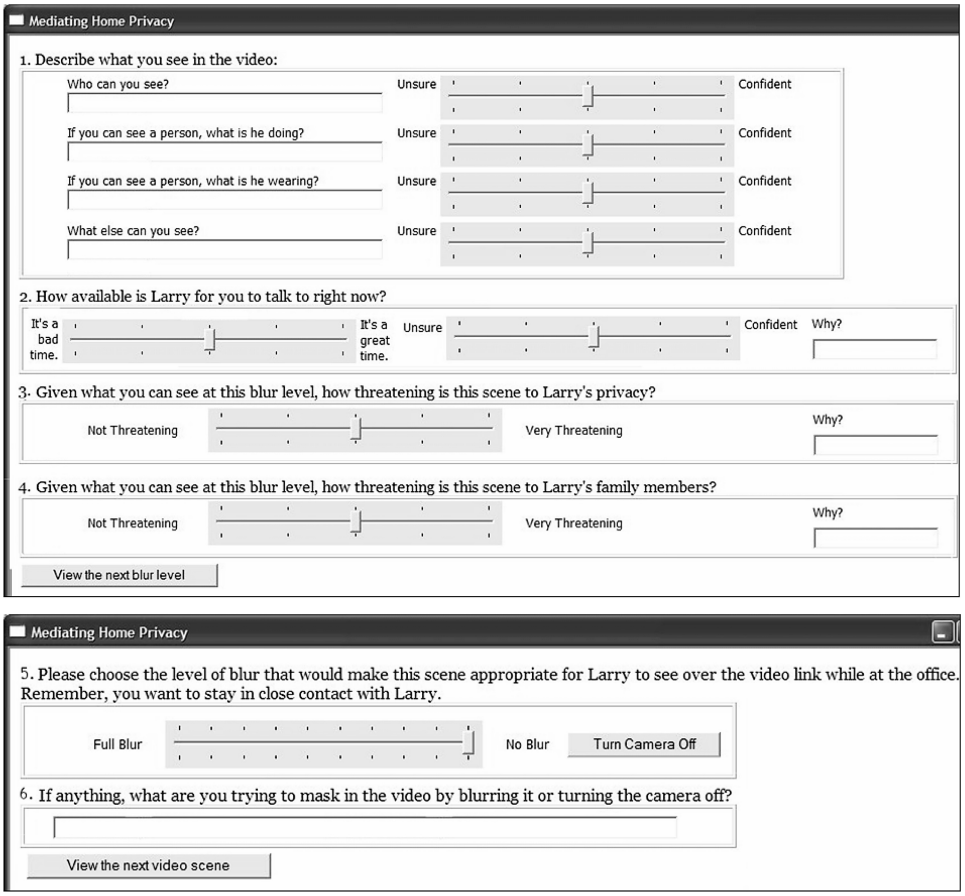

Fig. 3. During-test questionnaire (extracted from the study): privacy/awareness questions for each blur level (top), choosing a blur level (bottom).

2.7 Materials: Questionnaires

Three questionnaires were used during the study to gather data:

1. A **Pre-test Questionnaire** gathered demographics, such as age, gender, occupation, computer experience, and telecommuting experience. In our results, we discuss only gender, as we found no significant differences between other demographically grouped participants.
2. A **During-test Questionnaire** asked participants about each of the blur levels for all video scenes. We used two side by side 17″ CRT displays: the left showed a video scene, while the right presented questions about the video that the person could answer on-screen. As suggested by the questions in Figure 3, we asked people what they could see in each scene at that blur level. Figure 3 (top) shows awareness and privacy related questions asked for each of the blur levels. Similarly, Figure 3 (bottom) shows the set of questions used for each scene after the participant viewed all the blur levels

Fig. 4. A sample forced sort of scenes by privacy risk showing the 300 cm "line of privacy risk."

for it. These questions ask the participant to choose a blur level that would make the scene appropriate for a colleague to view.

3. A **Post-test Questionnaire** was divided into two parts. First, we gathered each participant's opinion of balancing privacy and awareness using blur filtration, and asked participants if they would use an open video link in an office if it was blurred and also at their home if it was blurred. Second, we asked participants to perform a forced sort of five pictures (one for each video scene, printed on standard 21.59 × 27.94 cm pieces of paper) according to how risky they felt each scene was to their privacy if they were the person in the scene. Participants were then asked to place the sorted pictures on a "line of privacy risk" that was 300 cm long (Figure 4): one end represented low risk, the other end represented high risk. Participants were told that they could leave as much space between pictures as they liked, but no two pictures could overlap. The "line of privacy risk" is used in a post hoc assessment/validation of our original rating of each scene's privacy risk.

2.8 Participants

Participants were twenty people—ten females and ten males—holding a range of professional occupations typical of telecommuters, for example, researchers, administrators, consultants, and software developers. We deliberately excluded undergraduate students as we thought their youthful attitudes towards privacy and exposure would tend to be more liberal than those of working professionals. All participants were recruited through email or with a poster advertisement and were paid $25 CND for their participation. Participants ranged in age from 21 to 55 years old, with a mean age of 29 years, and all were regular computer users with experience working in an office environment. Participants were also randomly counterbalanced for telecommuting experience—10 participants (6 male, 4 female) frequently telecommuted either currently or in the past, while the remaining 10 had little or no telecommuting experience.

2.9 Method

The study is a *within subjects* design. Each male participant was shown all five video scenes where the telecommuter was male, while each female participant saw the scenes using the female as the telecommuter; thus, all twenty participants saw each condition (scene type) in the experiment. After completing the pre-test questionnaire, participants were given the telecommuting scenario (Section 2.4). They were then asked to role-play, where they were first told they were at the office and that they would look at each scene in turn in order to determine whether or not Larry/Linda was available for interaction.

1. Participants viewed one of the five video scenes at the first fully blurred level (i.e., blur level 1 in Figure 2).
2. As they viewed each blur level, they answered awareness and privacy related questions (Figure 3, top).
3. They repeated steps 1 and 2 for the same scene at each of the remaining blur levels. This always progressed from fully blurred to completely unfiltered, and answers from previous blur levels remained visible so the participant could simply modify his or her answers.
4. They were then asked to imagine themselves as the telecommuter in each scene and were now themselves being watched by their colleague (Larry/Linda) at the office.
5. They chose a blur level for the scene and gave a reason (Figure 3, bottom). At this point, participants were able to view all blur levels at their discretion.
6. Upon completion of the first video scene, Steps 1–5 were repeated for each of the remaining four video scenes.
7. After completing all five video scenes, they answered the post-test questionnaire, and performed the forced sort of all scenes.

The first scene shown to participants in Step 1 was always our most benign control scene containing the working telecommuter (Figure 1). We used this scene first to offset the chance that a participant may become "ultra-conservative" if they first saw a risky scene and thus rate later less risky scenes as being more threatening than normal. The viewing order for the remaining four scenes was randomized. Participants did not see video scenes where the telecommuter was of the opposite sex because it was felt that imagining yourself as the opposite sex for a portion of the questions may be quite difficult and could confound the results.

When identifying awareness cues for a particular blur level (Steps 1–3), participants were able to use the information they had gained by viewing the scene at previous blur levels. For example, when viewing blur level 4, a participant had already seen the video at blur levels 1–3 and was able to use this information to help infer awareness information about the current blur level. While this is unlike real-life, it does provide us with a *blur threshold*, a lower bound for awareness. This threshold is the first point at which it was possible for participants to accurately deduce awareness cues. In Step 5, participants are asked to choose a blur level *after* seeing the scene in full fidelity (and knowing if it

was embarrassing or not) because we wanted to know how each scene's level of risk affected a participant's selection of blur level.

3. RESULTS

Our results are divided into three sections. First, we validate our original risk assessment of each scene, where we compare it with the results of the forced sort. Second, we address our three research questions by analyzing our results. Finally, we determine people's willingness to actually use blur filtration within a home media space, as captured on the post-test questionnaire.

We should mention that our original data analysis divided our participants into telecommuters and nontelecommuters. However, our analysis showed little difference between these two groups. For simplicity and clarity, we exclude this distinction in the following discussion and figures unless absolutely necessary.

3.1 Perceived Privacy Risk of Scenes

To discover how participants perceived the privacy risk of each scene, we had them perform a post-study forced sort of representative nonblurred pictures of each scene along a "line of privacy," with one end indicating no risk and the other high risk (as in Figure 4).

There was reasonable consistency in participant responses: we saw only six distinct orderings, and even those did not differ much. Figure 5 shows these orderings as indicated in the 'totals' column: 6 of the 20 participants gave the first sequence, 5 the 2nd, 3 the 3rd and 4th, 2 the 5th, and 1 the 6th (these total numbers are further separated into male/female in the Figure).

All participants placed Working as least risky (column 1), while 18 of the 20 had Changing as the most risky scene (column 5). The two male dissenters felt Kissing was more risky than Changing. The major difference between the orderings is the placement of the middle three scenes. For 5 of the 6 orderings, No Shirt, Picking Nose, and Kissing were the middle three scenes, albeit in varying positions. We can ascribe part of this variation to gender differences, i.e., how males rated male actors with No Shirt *vs*. how females rated female actors with No Shirt.

As a whole, we believe that participants' ordering of scenes confirms our original assessment of each scene's risk: Working is low risk, Picking Nose, No Shirt, and Kissing are moderate risk, and Changing is high risk.

Relative order does not indicate the strength of participants' convictions of risk. To capture this, we analyzed how participants position scenes on the line of privacy. Figure 6 graphs our results, with each scene type on the x-axis, and privacy risk on the y-axis (measured by position on the line of privacy: 0 cm—no risk to 300 cm—high risk). We also performed an ANOVA—scene type (5) x gender (2)—that suggested there is no difference in how males *vs*. females determined a scene's risk factor ($p = 0.78$), but a significant difference in the risk associated with a scene type ($p < 0.05$), for example, particular scenes were viewed as being riskier than other scenes.

The figure and an ANOVA test suggest the scenes below can be ranked into four categories of risk. First, almost all judged the Working scene (Figure 6, far

Ordering of Scenes (most frequent to least frequent)	Frequency		
	Male	Female	Total
1 	5	1	6
2 	1	4	5
3 	1	2	3
4 	1	2	3
5 	2	0	2
6 	0	1	1

Fig. 5. The frequency of each ordering of scenes found in the forced sort by males and females. Male participants used the male equivalences of the scenes shown.

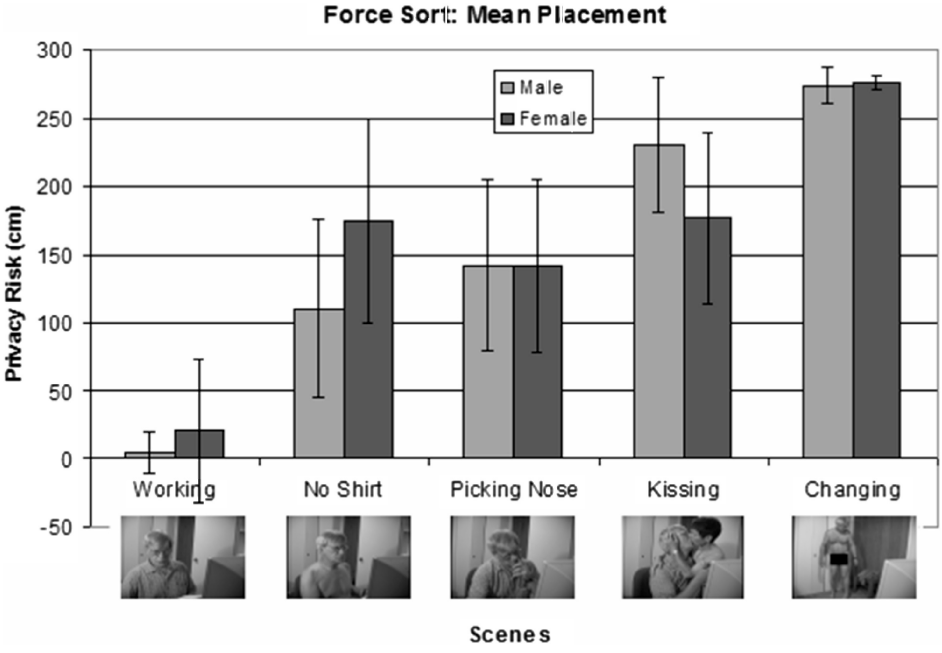

Fig. 6. The mean placement of scenes according to risk, from low risk (0 cm) to high risk (300 cm), during the forced sort.

left) as very low risk: the mean is 4.8 ± 15.5 cm for males and 20.8 ± 52.1 cm for females. A post hoc analysis[1] shows the next category above Working collects No Shirt and Picking Nose into a low-moderate risk rating ($p < 0.01$). Kissing has a somewhat greater moderate-high risk rating ($p < 0.01$). All judged Changing (far right) as very high risk ($p < 0.01$); mean image positions were 274 ± 13.1 cm for males and 276.1 ± 5.1 cm for females.

3.2 Determining Awareness

For each level of blur, participants were asked to write what they could see in the scene, rate how available the person was for conversation, and indicate the confidence they had in their guesses (Figure 3, Questions 1 and 2). We took this information and separated it into four awareness categories:

1. ***activity:*** the main activity found in the scene
2. ***person:*** who was in the scene
3. ***appearance:*** what the person(s) in the scene was wearing
4. ***availability:*** how available the telecommuter is for interaction right now

We initially included 'background' as a fifth awareness category. However, participants typically did not mention background items that added any awareness value, and they also stopped noting this information as the study progressed. This supports our belief that background information becomes unremarkable over time. Consequently, we do not incorporate 'background' into our results. We now pose a series of questions to be answered by our observations.

At what blur levels did people correctly identify awareness cues? Cues from each of the four awareness categories (listed above) were generally identifiable between blur levels 3 and 5 (Figure 2 showed these blur levels). Figure 7 plots the median and range of blur levels at which participants were first able to correctly identify categories of awareness cues for each scene. That is, the figure gives the blur threshold where subjects could just determine these cues. We judged correctness for the activity/person/appearance categories by verifying that the participants' descriptions matched what was actually happening in the scene, regardless of their confidence in their response. Because availability is a subjective measure, we judged an availability response as correct when the participant indicated they were quite confident (3 or greater out of 5) in their answer (Figure 3, Question 2).

For all scenes on average, when determining what activity was occurring in each scene, 75% of participants were able to do so between blur levels 3 and 4 (Table I). The appearance of the actor in each scene was determined by 75% of participants between blur levels 3 and 5 (Table I). Determining who was in each scene was performed between blur levels 3 and 5 by 65% of participants (Table I). Availability was determined by 65% between blur levels 3 and 5 (Table I). The remaining participant breakdown is summarized in Table I. These numbers show that for the majority of participants, their threshold for

[1]All post hoc analyses, unless otherwise stated, were performed with a series of T-Tests using Bonferroni correction.

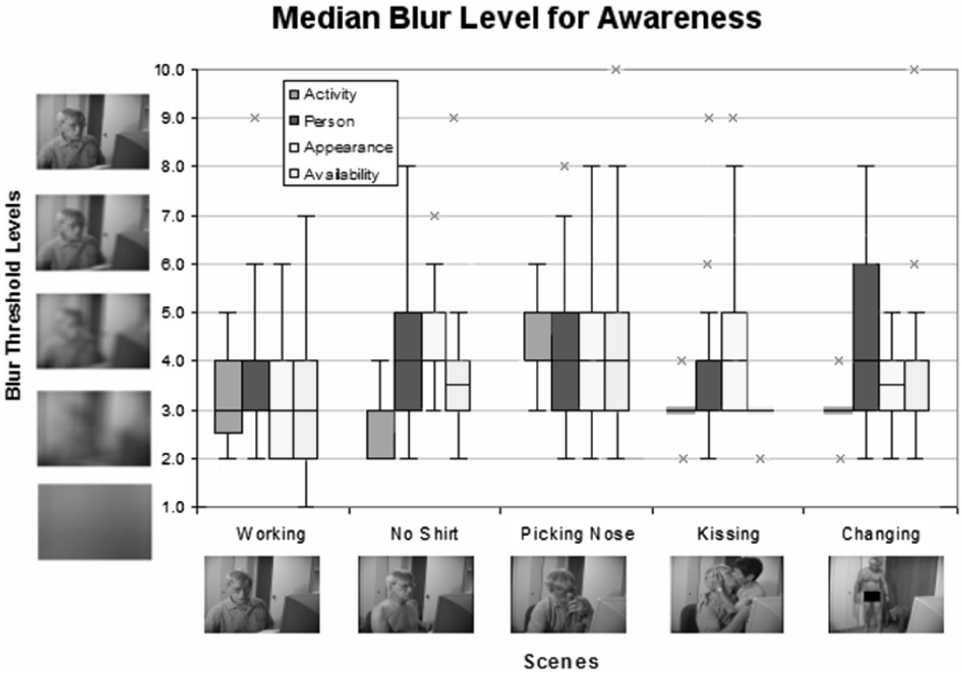

Fig. 7. The median and range of blur levels at which participants were first able to identify awareness cues for each scene.

Table I. Awareness Cues: The Percent of Participants Able to Identify Awareness Cues at Specific Blur Levels

	Blur Levels	Percent of Participants
Activity	<3	25%
	3–4	75%
Appearance	<3	15%
	3–5	75%
	5–6	10%
Person	<3	15%
	3–5	65%
	5–8	20%
Availability	<3	20%
	3–5	65%
	5–6	15%

identifying awareness cues with reasonable confidence is *between blur levels 3 and 5*, although a few participants recognized cues even earlier.

Does scene type affect the blur level at which people begin to correctly extract awareness cues? Yes, the scene type did affect the blur level at which participants begin to correctly extract awareness cues. A series of Friedman ANOVAs (scene type (5) × awareness category) shows that there is a significant difference in the blur levels participants used to first identify awareness information across the five scene types for availability ($\chi^2(4) = 17.62$, $p < 0.05$), activity ($\chi^2(4) = 40.55$, $p < 0.05$), and appearance ($\chi^2(4) = 25.38$, $p < 0.05$).

No difference exists between scenes for determining the person ($\chi^2(4) = 8.592$, $p = 0.072$). A series of Wilcoxon Matched-Pairs Signed-Ranks tests show that the differences are quite varied among scenes for availability and appearance. For activity, the differences lay between the Picking Nose scene and all other scenes (Picking Nose *vs.* Changing, $z = -3.56$, $p < 0.05$, Kissing, $z = -3.70$, $p < 0.05$, No Shirt, $z = -3.90$, $p < 0.05$, Working, $z = -3.13$, $p < 0.05$). In general, participants needed higher video fidelity to identify the picking nose activity than other activities. This is reasonable as this activity involved only a small part of the scene and contained the smallest amount of movement.

Did people's ability to extract availability information from a blurred scene depend on the awareness category? Yes, a series of Friedman ANOVAs (awareness categories (4) x scene type) shows that there is a significant difference in blur levels found for each of the awareness categories for three of the five scenes: Changing ($\chi^2(3) = 28.09$, $p < 0.05$), Kissing ($\chi^2(3) = 18.03$, $p < 0.05$), and No Shirt ($\chi^2(3) = 24.22$, $p < 0.05$). The remaining two scenes, Picking Nose ($\chi^2(3) = 1.02$, $p = 0.80$) and Working ($\chi^2(3) = 3.09$, $p = 0.38$), saw no difference between awareness categories. That is for Changing, Kissing, and No Shirt, only particular categories of awareness information could be determined at particular blur levels. A series of Wilcoxon Matched-Pairs Signed-Ranks tests show, for these three scenes, that participants determined activity at slightly blurrier levels (within one to two blur levels) than what was needed to determine who the person in the scene was and what this person's appearance was like.

How confident were people in their awareness responses? Participants were not very confident in their initial answers (even when they were correct) and in most cases did not become confident until fidelity increased another 2 or 3 more levels. Figure 8 shows the confidence participants had in their ability to determine awareness cues. This mean represents the average confidence that participants had in identifying all awareness components: activity, person, appearance, and availability. We use this to represent the amount of awareness presented by each of the blur levels, as their confidence reflects their belief that they were correctly interpreting the scene. A two-factor ANOVA (scene types (5) $\times$ blur levels (10)) confirms that for all scenes, there is a significant difference ($p < 0.05$) only in the amount of awareness presented by the blur levels.

How did people rate a person's availability in an unblurred scene? Although less important for balancing privacy and awareness, we were curious to know how participants ranked the telecommuter's availability for each of the scenes when they were shown completely clear (Figure 3, Question 2). The Working scene represented *being available* for most participants (mean = 4.3, $s.d. = 1.3$, 1-not available to 5-highly available). Participants believed No Shirt and Picking Nose represented *some availability* (mean = 3.4, $s.d. = 1.5$; mean = 3.2, $s.d. = 1.7$ respectively). Participants rated Kissing and Changing Clothes as *being unavailable*, (mean = 1.2, $s.d. = 0.9$, mean = 1.1, $s.d. = 0.2$, respectively).

In summary, all of these results show that people begin perceiving all categories of awareness cues between blur levels 3 to 5, and that scene type does

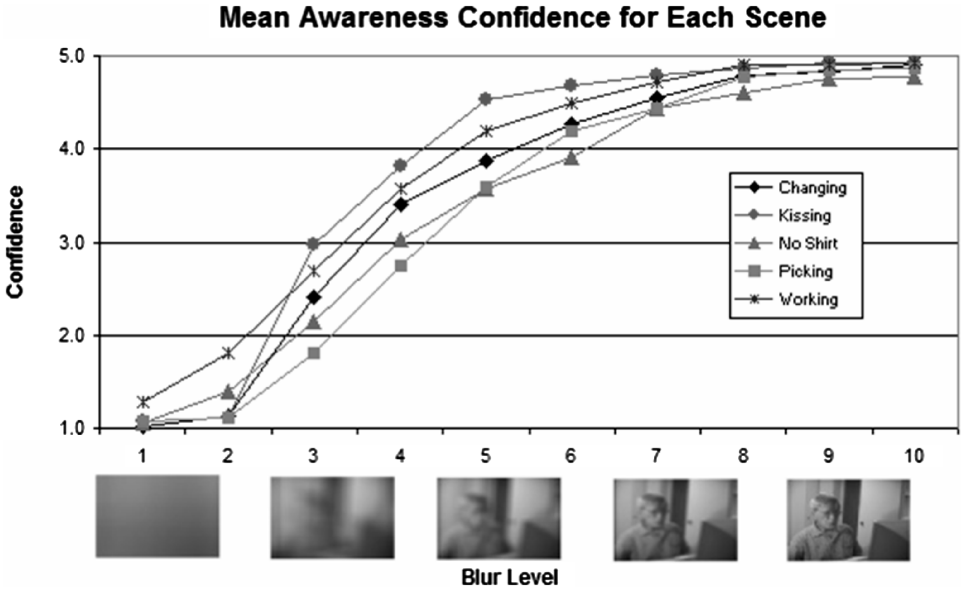

Fig. 8. The mean level of awareness confidence found at each blur level (1-low confidence to 5-high confidence).

make a difference to this result. Furthermore, we've shown that their confidence in their guesses increases with fidelity.

3.3 Privacy Threat

For each level of blur, participants were asked to rate how threatening the scene was for the telecommuter and family members, given what they could currently see (Figure 3: Questions 3 and 4). In this section, we first describe the privacy threat to telecommuters, followed by the threat to family members.

Does the perceived threat to the privacy of the telecommuter differ by blur level? Yes, blur level affects the perceived privacy threat to the telecommuter. The mean privacy threat indicated at each blur level is shown in Figure 9. At levels 1 and 2, participants perceived little to no threat for all scenes. After this, the perceived threat increased with fidelity. This increase occurs dramatically between blur levels 3 and 5 (the region indicated in the figure), and levels off by blur level 7. A two-factor ANOVA (scene type (5) $\times$ blur levels (10)) verifies that the privacy threat between different blur levels does differ significantly ($p < 0.05$). Figure 9 shows that these differences typically start between blur levels 2 and 3, and increase steadily until blur level 5. We see this result even for the Working scene (which remains mostly nonthreatening), suggesting that participants associated added threat with greater image fidelity, even in nonrisky scenes.

Does the privacy threat to the telecommuter differ by scene? The same ANOVA also verifies that there is a significant difference in the threat for telecommuters between scenes ($p < 0.05$). A post hoc analysis of overall mean privacy threat ($p < 0.01$) shows the scenes may be partitioned into three

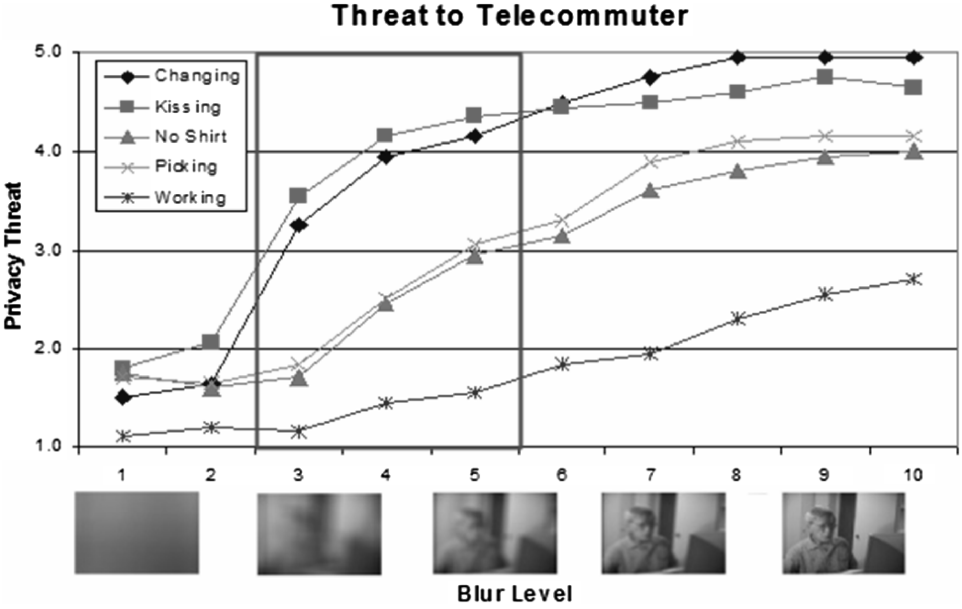

Fig. 9. The level of privacy threat (1-low threat to 5-high threat) to the telecommuter at each blur level. The rectangle highlights blur levels 3 to 5, shown to provide awareness.

categories of threat. The low risk category consists of the Working scene. A moderate risk category includes the No Shirt and Picking Nose scenes. A high risk category holds the Changing and Kissing scenes.

What, if anything, made each scene threatening to the telecommuter? Participants usually associated threat with the person's particular activity or appearance. As fidelity increased, these acts and their details became clearly visible and thus more threatening. Several participants also commented that they felt the scenes would be more threatening if they were viewing a colleague of the opposite sex.

Does blurring affect the privacy threat to family members? Despite the fact that a family member appeared in only one scene, participants still found the scenes to present some level of threat for family members, discussed below. This threat is similar to that posed to the telecommuter: single factor ANOVAs ($p < 0.05$) performed on a scene-by-scene basis reveal no significant differences between the threat to family members and the threat to the telecommuter, except in the Picking Nose scene. There are two obvious distinctions, however: the mean threat at a given blur level is generally lower for family members than it is for the telecommuter, and the Kissing scene posed the highest risk to family members, while the Changing scene posed the highest risk to the telecommuter.

What, if anything, made each scene threatening to family members? Participants' responses were quite similar to those given for the telecommuter: threat was associated with the visibility of the person's risky activity or appearance. Curiously, participants rated the Changing scene as very threatening to

family members, even though no family member is ever present in the scene. The most common reason given by participants for this rating concerns the *potential* for threat: at any time a family member could walk into the room, and the fact that one wasn't there now was almost moot. This reason was given despite the fact that our question (which was accompanied by verbal explanation) specifically asked participants to rate the threat based on what could be seen currently—participants had a tendency to infer what could happen even when instructed not to. A second, less common reason given was that participants felt that the family members may suffer the consequences e.g., embarrassment or ridicule, should the telecommuter's reputation be affected by a privacy violation.

In summary, all these results show that only blur levels 1 and 2 make all scenes nonthreatening. The results also allow us to partition the scenes into three categories of risk: low (Working), moderate (Picking Nose and No Shirt), and high (Kissing and Changing).

3.4 Choosing Blur Levels

Participants were asked to imagine themselves as the telecommuter (Larry or Linda) and then, for each scene, choose a blur level (from 1 to 10) that they felt would make the scene appropriate for their colleague to view (Figure 3: Question 5). They also had the option to 'turn the camera off,' which we codified as a blur level of 0.

What blur levels did participants choose to make a scene appropriate for a colleague to view? The results vary with risk category (found in the previous privacy analysis) but do not differ in statistically significant ways by gender. Figure 10 plots our results, where the y-axis shows the median selected blur levels chosen by participants for each scene. As one would expect, participants chose more revealing blur levels for the low-risk Working scene (median = 6) than for higher risk scenes, e.g., Changing (median = 1). The results from a Friedman ANOVA looking for differences by scene ($\chi^2(4) = 56.26$, $p < 0.05$) and a series of Wilcoxon Matched-Pairs Signed-Ranks tests show that the responses to this question partition the scenes into the same three risk categories we found in previous analysis. We were curious if gender made a difference. A series of Mann-Whitney–Wilcoxon Rank Sum tests found that there is no statistically significant difference between the blur levels chosen for a particular scene by males *vs.* females.[2]

When did people choose to turn off the camera? Nearly half of all participants chose to turn the camera off for the riskiest scene, yet only one turned it off for the least risky scene. That participant was adamantly opposed to using video at home and turned the camera off for every scene. Table II summarizes the proportion of participants who felt no blur levels were adequate for a scene and chose to turn the camera off (i.e., blur level 0) broken down by gender. For every scene, more female participants turned the camera off

[2](Changing, $u = 36.5, w = 91.5, z = -1.03, p = 0.29$, Kissing, $u = 28.0, w = 83.0, z = -1.71, p = 0.09$, No Shirt, $u = 34.5, w = 89.5, z = -1.21, p = 0.22$, Picking, $u = 30.0, w = 85.0, z = -1.53, p = 0.12$, Working, $u = 29.0, w = 84.0, z = -1.68, p = 0.09$).

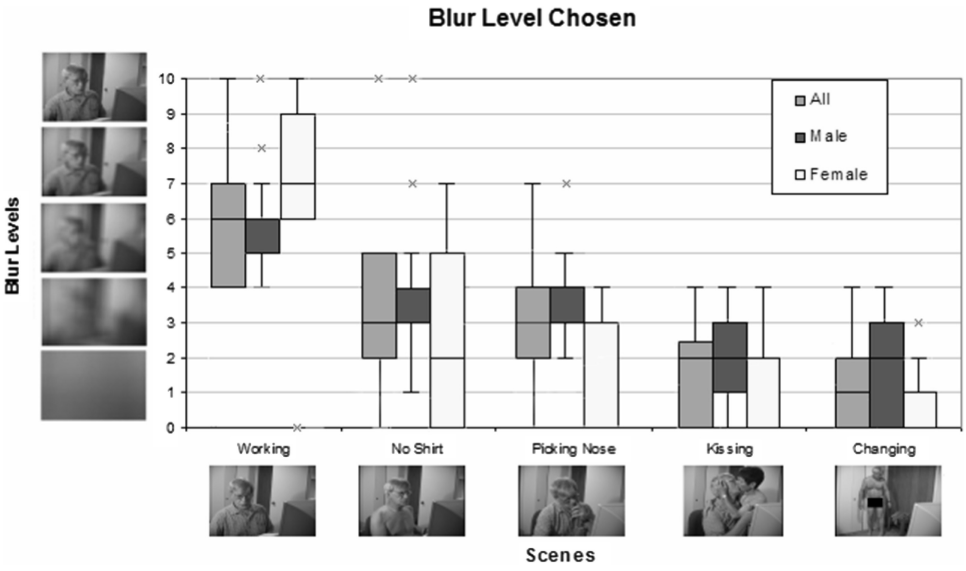

Fig. 10. The median and range of blur levels chosen by participants for each scene. Blur level 0 represents choosing to turn the camera off.

Table II. The Percent of Participants Who Chose to Turn Off the Camera

	Male (n = 10)	Female (n = 10)	All (n = 20)
Working	0%	10%	5%
No Shirt	0%	40%	20%
Picking	0%	30%	15%
Kissing	20%	50%	35%
Changing	30%	60%	45%

than male participants, and a two-factor ANOVA (gender (2) $\times$ scene type (5)) showed the propensity to turn the camera off does in fact differ in a statistically significant way according to gender ($p < 0.05$). Ignoring this gender difference, we can see in the 'All' column of Table I that the five scenes break down into roughly the same three risk categories determined in other analyses.

In summary, these results show that participants choose more distorted blur levels or more participants choose to turn the camera off altogether in order to make a video scene appropriate for a colleague to view, as the risk to privacy posed by a scene increases. Perhaps more importantly, we see that as the risk posed by a scene increases, more people abandon the blur filter in favor of turning the camera off altogether, and that nearly half of participants turn the camera off in order to make high risk video acceptable.

3.5 Willingness to Use Blurred/Unblurred Video

In the post-test questionnaire, we asked participants how willing they would be to use video in their own home to connect to a colleague they work closely

Table III. Participants Who Would/Would Not Use Blurred Video in an Office and at Home

	Participants (n = 20)
Office - yes	13
Office - no	7
Home - yes	9
Home - no	11

with (1-unwilling to 5-willing). The mean willingness for all participants to use unblurred video was 1.9 ($s.d. = 1.0$), while blurred video (that could have its blur level adjusted) was 3.3 ($s.d. = 1.3$). These values are significantly different ($p < 0.05$). We also checked to see if there were significant differences between male and female responses, but none were found.

Participants were also asked what they liked and disliked about using blurred video to balance privacy and awareness. Common likes included: being able to show availability while masking sensitive details, having the ability to control one's privacy, and being able to easily stay in contact with others. Common dislikes included: not being able to easily determine availability from blurred video, not knowing what the other person thinks they are seeing, and having to decide how much to blur and to alter this blur level for various scenes. Several participants said that they felt there was no balance between privacy and awareness: at the point where they could tell what was going on, they didn't feel the person's privacy was adequately being preserved. One participant also indicated a concern that blurred video could be unblurred by the viewer. As mentioned previously, one person was adamantly opposed to using video.

When asked if they would—given the opportunity—actually use blurred video in an office, 65% of participants said they would (Table III). Those who wouldn't use blurred video from an office said they preferred other means of gaining awareness, such as email, instant messaging, phone, or simply just walking over to see a person. They also commented that they felt their personal security would be violated when using blurred video, as the balance between privacy and awareness simply wasn't there.

Participants were then asked if they would—given the opportunity—actually use blurred video from home: 45% said they would (Table III). Most of those who said they would use blurred video at home imposed caveats and restrictions: they wanted to choose the room where the camera was located and they wanted a mirror facility to know what was being captured. They also wanted control over: the blur level, whether or not the camera was on, and—by setting the camera angle—the part of the room that was being captured. Several also commented that they would simply leave the room to do private things that they would not want their colleagues to see. Those who said they would not use blurred video at home explained that they would find it intrusive, that it would violate their personal security, and that they felt blurring did not balance privacy and awareness. They also said that they saw the home as a place where they could go to achieve solitude from their colleagues. They felt that conventional mechanisms—email, instant messaging, or phone—are

adequate means for gaining awareness. One participant said that she would be fine with using blurred video at home, but didn't feel her husband would want it.

4. DISCUSSION

We set out to evaluate blur filtration for its effectiveness in balancing privacy and awareness for video-based telecommuting situations. In particular, we wanted to know whether or not blur levels existed that could provide adequate awareness, while still preserving privacy. The answer to this overarching question can be found by looking at the answers to our research questions.

4.1 Question 1

At what blur levels are participants able to identify who is in the scene, what they are doing, and what they are wearing?

Aggregating the results across all scenes tested, we found that awareness cues were first identifiable between *blur levels 3 and 5.* While these blur levels may seem quite blurry in Figure 2, it is important to remember that participants saw full motion videos and that this motion aids in identification. The levels we found for providing awareness are somewhat more filtered (2 to 3 levels) than those found by Boyle et al. [2000]; thus, in our study, participants were able to garner awareness cues from blurrier scenes. We believe this difference is a result of using videos of a greater fidelity than Boyle et al. [2000].

4.2 Question 2

At what blur levels is privacy adequately preserved for the telecommuter and others in the home?

Blur levels 1 and 2 are the only levels that adequately preserve privacy for all scenes. It is clear that these blur levels do not overlap the awareness levels of 3 to 5; thus, *there are no general-purpose blur levels that can balance privacy and awareness in any scene*. If we analyze this on a scene-by-scene basis, we see that the Working scene, representing a mundane home situation, is the only scene where privacy preserving levels overlap the awareness range. Thus, we can see that blur filtration is only able to balance privacy and awareness for mundane home situations. This is consistent with the Boyle et al. [2000] result, which found overlap for what we consider here to be mundane scenes. The more important result, however, is that blur filtration is *not* able to balance privacy and awareness for the high risk home situations that we are interested in.

4.3 Question 3

What blur levels do participants choose in order to make a given scene appropriate for a colleague to view?

The blur levels that participants chose varied for the three different risk categories that we found: low threat, moderate threat, and high threat. Participants chose more distorted blur levels or more participants chose to turn the camera off altogether as the risk to privacy posed by a scene increases. While this is expected, the contribution and importance lays in the fact that people

begin to abandon the filtration technique with increased risk. Participants said they simply did not trust the blur filtration because, regardless of the detail of information being transmitted, there is still a fear that if the camera is facing them, high fidelity video is being captured and broadcast. Participants asked for other strategies that could give them more control over their privacy. This brings us to the conclusion that blur filtration by itself is not an appropriate technique for balancing privacy and awareness for home situations.

4.4 External Validity

Because our results are based on a laboratory experiment, concerns about external validity and how our findings would generalize to other home/work situations and how they are affected by other factors, are especially important. Perceptions, judgments and choices about privacy, after all, are always situated in a real world context.

In the real world, people's willingness to give up their privacy in order to gain awareness is quite complicated and many factors come into play. One factor is personality. Some people are more flamboyant and/or gregarious. As a result we would expect that they may be much more willing than others to give up more of their privacy in exchange for conversational opportunities. Of course, those who may be shy and quite reserved would have the opposite reaction. Another factor is workplace pressures and how they drive the need to collaborate. If the need is high, people may use the media space in spite of their normal concerns about privacy because of the self-induced pressure to get things done in a timely way. A third factor is workplace power relationships. If people in a position of power insist upon others being in the media space—either explicitly through fiat or implicitly through expectations—then those others may participate in spite of their reluctance or concerns. A fourth factor is similar: if the cultural practice of a group is to use the space, then social pressure and the need to conform may compel otherwise reluctant people to participate. These, and likely other factors as well, are clearly important. They speak about the choice people make to use the media space in spite of the risks inherent in its use. Our concern in this article, however, is about how privacy is protected once people make the choice to use the media space.

To simplify the situation, our experiment tried to minimize the negative aspects of this choice by controlling real-world factors like the ones mentioned, to test a very specific situation. Our scenario placed people in a position where they would want to participate in the media space with a person that they know well, where both parties would benefit from the link. This scenario presents something akin to a "best case" situation for participation. Yet even with many of these real-world factors controlled, we see that blur filtration cannot handle the privacy-awareness complexities that we present. From this, it is clear that given the increased challenge of handling a full facet of real world factors, blur filtration by itself is much too limiting a technique to be used in actual practice. It is a false safeguard. Blur filtration will not make the shy person any more willing to use the media space. Neither will it protect those workers who either have personal incentive to use it or are compelled to do so because of work pressures,

a power relationship, or cultural practice. Nor will people's mistrust of blur filtration disappear. Blur filtration is not a privacy panacea for media spaces.

4.5 Generalizing to Video Obfuscation

From a technical perspective, blur filtration is just one of the many image fidelity reduction techniques that attempt to obscure sensitive details. While each frame of the video is sent through the media space, its appearance is altered to hide information. Other techniques include pixelization and edge-detection filters [Zhao and Stasko 1998]. A fair question to ask is whether our findings and concerns about blur filtration would generalize to these other obfuscation techniques. We believe the answer is yes.

All video obfuscation techniques rely on the same basic premise for protecting privacy: all or part of the video is obscured to a point where sensitive details are hidden. We know that blur filtration failed at providing an obfuscation level that could balance privacy and awareness for several reasons. One is perceptual: if people could extract meaningful awareness information, they could (to some extent) also extract enough information to raise privacy concerns. The second is judgmental: people simply did not trust the technique given the situations presented to them regardless of whether they could actually perceive privacy-invasive information at a particular blur level. The presence of a camera was always seen as a risk, no matter how the image was being obscured. It is this second point that leads us to believe that other video obfuscation techniques will fair no better than blurring, as they all rely on this same basic premise of masking visual information obtained from a camera.

Another related image manipulation technique uses substitution where information in the source frame is used to compose a new artificial frame. One example is an eigen-space filter which replaces socially inappropriate scenes using a predefined set of "socially correct" images [Crowley et al. 2000]. Another example is the shadow-view filter where dark pixels are placed on a static reference image depicting the location of activity in the video [Hudson and Smith 1996]. We know that these techniques would safeguard the risky situations in our study, at least in principle. For example, using an eigen-space filter, our picking nose scene would only show the actor working. More risky scenes would show a similar mundane situation, or an empty room. However, in spite of this, our second point above suggests that this will do no better than video obfuscation techniques since people still have an inherent mistrust of any technique based on image capture. While it may be the case that people learn to trust this over time, their initial exposures and acceptance of the media space will be tempered by this mistrust. This deserves future study.

Consequently, we hypothesize that obfuscation techniques are simply too limiting by themselves for handling the many intricacies of privacy. The next section outlines the implications from these results for the design of privacy-protecting strategies.

5. DESIGN IMPLICATIONS

By taking a step back from the study results, we can see that several important issues arise for developing privacy-protecting strategies to balance privacy and

awareness in a home media space that can potentially capture and transmit risky scenes. The four issues are:

1. A home media space is not for everybody.
2. Privacy control and feedback must be available.
3. Differences exist between individuals within the home.
4. Two distinct cultures are connecting.

Each of these issues is articulated in the responses participants gave in the study's questionnaires and give a better understanding as to why it is difficult to develop privacy-protecting strategies for a home media space. These issues have crucial design implications for a home media space and suggest potential solutions to the problem of balancing privacy with awareness. Table IV summarizes a set of observations from the study (Column 1), the implication drawn from each observation (Column 2), and an example design practice for addressing the implication (Column 3). The design examples are discussed in more detail in Section 6 where we describe our own prototype design of a home media space [Neustaedter and Greenberg 2003a]. The next sections describe each design implication in more detail.

5.1 It's Not for Everybody

The first issue is that a home media space is clearly not for everybody. This study looks at an idealized situation where two intimate collaborators have a need and desire to work closely together—this situation is not always the case. Several participants commented that they felt privacy would be more threatened if their colleague was a person of the opposite sex. In practice, many telecommuting relationships will be less than ideal, for example, two people who have just met, telecommuters are of the opposite sex, or one person is in a power position over another. In these cases, the privacy threat will increase even though awareness will remain constant. On the other hand, situations involving just family and friends may require less privacy and more awareness.

Despite the idealized telecommuting situation presented in the study, it is safe to say that there are certain people for whom a home media space will work for and there are those for whom it will not work. Undoubtedly, the person in our study who turned the camera off for all scenes is not a suitable candidate for a home media space. On the other hand, others in the study who on average were moderately willing to use blurred video from home or those who actually said they would use it, may be more suitable candidates. It may also be the case that concerns drop after a period of usage, similar to the Active Badge system [Want et al. 1992].

For those individuals that a home media space is suitable for, there is always a varying degree of suitability based on individual preferences. It is important to remember that all participants must choose to be a part of a home media space. These media space participants include the telecommuter, work colleague, and other individuals who may be subject to the media space such as family members of the telecommuter. Each must be given an opportunity to decide whether or not they wish to participate in this space and if one participant declines,

Table IV. The Set of Study Observations That Led to Each Design Implication and Example Design Practices to Address Each Implication

Observation	Implication	Example Design Practice
Several participants said they would not use video in a home, or that their family members would object to it.	A home media space is not for everybody.	Use context-aware technology to match individuals and their privacy expectations (e.g., via RFID tags or active badges).
Participants who said they would use video in a home wanted a high degree of control over the media space, along with sufficient feedback of the current privacy state.	Privacy control and feedback must be available.	Provide control with gesture-based input and physical interface controls. Make the camera a source of visual and audio feedback by allowing it to rotate towards and away from the user.
Participants were concerned about the privacy of others in their home; the reputation of others in a home could easily be jeopardized by a media space.	Differences exist between individuals within the home.	Use context-aware technology to detect other home dwellers and use that to control (e.g., to mute) the media space.
Participants struggled when answering questions about the appropriateness of each scene because they didn't know how to judge different contexts, e.g., home vs. office.	Two distinct cultures are connecting.	Make the recording state easily discernable with timely feedback to promote self-appropriation.

then the media space should not be set up or an even greater privacy threat will arise. The implication is that without a user interface or social procedures that can allow people to withdraw or participate, privacy threats will increase.

5.2 Provide Control and Feedback

The second issue is that home media space participants desire a sufficient level of control over their privacy, just as they want an adequate level of feedback informing them of their achieved privacy. Privacy control and feedback must be available for all participants in a home media space if privacy threat is to be reduced. Participants in the study desired to stay in contact with their colleague, yet each had their own individual preferences for how much awareness and privacy they desired. It was also clear that their desires were not static. Thus, participants need individual and continuous control over awareness and privacy. Control in a home media space could mean such things as control over the camera, what it captures, and how it captures. If blur filtration

is used to alter what others see, then control is needed over what blur levels are used for different situations. Feedback in a home media space could mean such things as sound cues, or LED lights to notify people that the camera is capturing.

In everyday situations, we regulate our privacy very effortlessly and typically without thought. It is important when designing a home media space to stay within this paradigm: privacy-protecting techniques should be simple to use, if requiring any effort at all. Unlike blur filtration, which is overly simplistic because it does not recognize the context of its use, context-aware approaches for balancing privacy and awareness can adapt to the current situation and adjust privacy levels as needed. Privacy regulation is a situated, contextualized activity; therefore, it makes sense to use a context-aware user interface for privacy control. While such solutions may appear to take control away from the user, they can be augmented with other simple and lightweight privacy regulation techniques like adjustable physical controls so users can fine-tune a privacy/awareness balance.

5.3 Differences within the Home

The third issue is that a home often has multiple people present with varying expectations of privacy. Privacy expectations will depend on a participant's current situation as well as his motivation to be a participant in the home media space. Telecommuters who gain a benefit from the space may be more motivated to participate in it than family members who gain no benefit. While family members can try to negotiate privacy and/or normalize their behaviors around the media space, this may be difficult in practice and perhaps even impossible. For example, a telecommuter may be working at home and behaving appropriately for the home media space, while a family member may not be if she walks into the room and undresses. If the home is small and the office space is shared with (say) the main bedroom, this introduces unavoidable conflict. The telecommuter wants to use the home media space, but certainly the family member does not. Here, the family member would likely force the telecommuter to turn the camera off and then the telecommuter would no longer be achieving his desired privacy. Balancing privacy and awareness in these situations is difficult and individuals may need to compromise their desired privacy for the desired privacy of another.

5.4 The Clash of Cultures

The fourth issue is that two cultures, a home culture and an office culture, are forced to mix when a home media space is used. Both of these cultures have their own privacy expectations and what is appropriate for home is not always appropriate for the office. Offices are considered to be semi-public areas where individuals are expected to behave in a manner that is suitable for others to see. When using a home media space, office participants expect to see scenes that are appropriate for them. Yet homes are considered to be private areas where individuals have the freedom to relax and gain solitude [Altman and Chemers 1980]. For this reason, those in an office culture may have their privacy violated

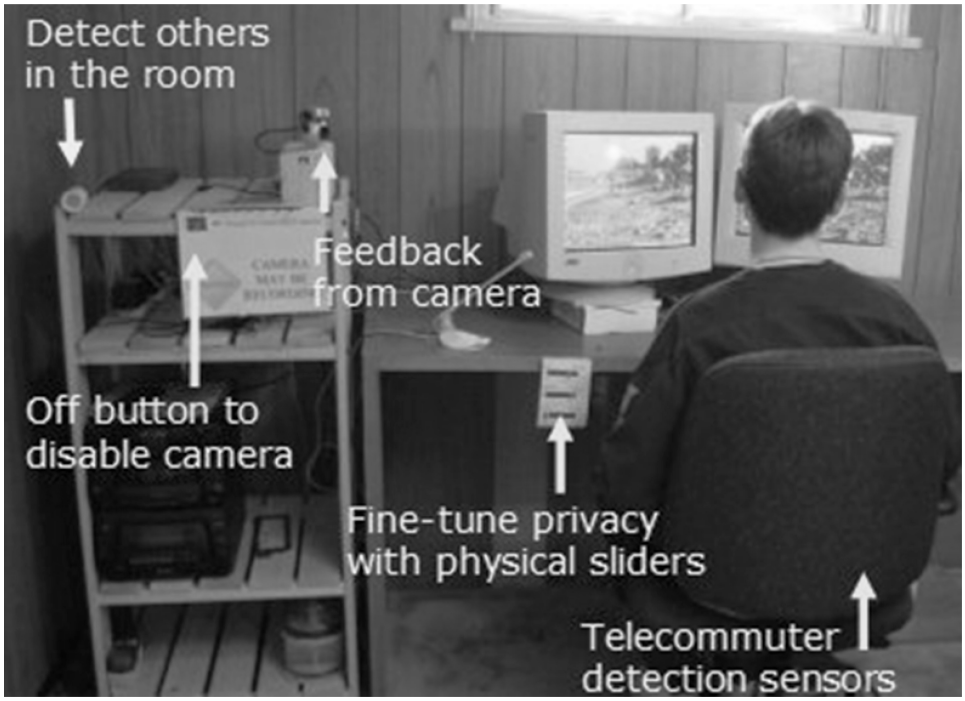

Fig. 11. The layout of our context-aware home media space.

by seeing something over the video link that is inappropriate for an office. Just the same, home participants in the space may have their privacy violated by having something captured over the video link that they don't want others to see. A home media space must attempt to balance the needs and desires of both these cultures.

6. AUGMENTING VIDEO WITH CONTEXT-AWARE PRIVACY PROTECTION

The results of our experiment and its design implications suggest that context-aware computing can be used as a tool for regulating privacy in home-telecommuting situations while presenting little overhead for users. For these reasons, context-aware systems that can detect and control privacy for users of a home media space may be desirable. Based on this, we have developed our first prototype design of a context-aware home media space [Neustaedter and Greenberg 2003a]. The home media space is set up in the spare bedroom/home office of a telecommuter (Figure 11) where sensing technology detects who is around and then infers privacy expectations through a simple set of rules, for example, if only the telecommuter is present, more awareness is provided by the home media space, or if someone other than the telecommuter is present, more privacy is provided by the home media space. While we talk about privacy in a linear manner, we recognize that it is much more complex and have designed our system with this in mind. Additional details about our system can be found in Neustaedter and Greenberg [2003a, 2003b].

In an effort to keep privacy regulation lightweight and interpretable, *implicit* actions such as entering/leaving the room, or sitting down in the office chair, can be used to regulate privacy. We understand, however, that context-aware systems can still make mistakes [Erickson 2003] and it is important that these mistakes do not cause increased privacy violations. For this reason, we first warn users if an implicit action has initiated a privacy decreasing operation; and second, we provide an opportunity for users to override this operation. Continuous *visual* and *audio feedback* of the current state of the system, for example, the sound and motion of a rotating camera [Boyle et al. 2000], or lit/unlit LEDs, allows users to always know how much privacy is being maintained. Timely and appropriate feedback is crucial so that users are aware if their desired privacy is being met by the system. Users can fine-tune the privacy/awareness balance with *explicit* actions, such as blocking the camera with their hand [Boyle et al. 2000], or adjusting a physical slider.

Privacy regulation and feedback are provided with specific elements in our home media space design:

- ***Camera state.*** The camera can be in one of three states: Play (the camera is recording), Pause (the camera is temporarily not recording), and Stop (the camera is permanently not recording and only an explicit action will move it out of this state).
- ***Capturing angle.*** The camera, mounted on a rotating motor, is placed near the door and, given the desired camera angle, can capture any region of the room, except the doorway.
- ***Video fidelity.*** Users can adjust the captured video's fidelity by explicitly adjusting the level of blur filtration used, the camera's frame rate, or the camera's frame size.
- ***Gesture-activated blocking.*** Users can easily turn off the camera by explicitly blocking it with their hand.
- ***Gesture-activated voice.*** Users can easily open an audio channel by explicitly moving their hand over a microphone.
- ***Easy-off button.*** Users can instantly turn off the camera by touching an off button.
- ***Telecommuter detection.*** We know if the telecommuter is present at the computer by detecting the implicit act of someone sitting down in or standing up from the desk chair. A radio frequency identity (RFID) tag in the pocket of the telecommuter identifies if the person sitting is the telecommuter.
- ***Family/friend detection.*** We know if family/friends are present in the room by detecting the implicit act of walking into and out of the room.
- ***Visual feedback.*** We use several visual cues to let the user know how much privacy is currently being maintained, for example, a sign, LEDs, the camera's direction, mirrored video, and the position of physical and graphical controls.
- ***Audio feedback.*** We also use audio cues to let the user know how much privacy is currently being maintained, for example, the sound of a camera clicking and the sound of the camera rotating.

Our home media space design presents one approach for using context-aware technology along with a set of simple rules to provide privacy regulation in home-based media spaces. Our future direction involves formally evaluating our design and its use of context-aware computing for regulating privacy in home media spaces.

7. CONCLUSION

We began with the research goal of determining how well video-blurring safeguards privacy in always-on video links that connect home-based telecommuters with office colleagues. Previous research has shown that blur filtration is able to balance privacy and awareness for mundane office situations [Boyle et al. 2000], yet it was not clear whether this would hold for risky situations present in a home environment. Based on this, our research contributes two main points.

1. An evaluation of blur filtration shows blur filtration and, by implication, other video obfuscation techniques are not able to balance privacy and awareness for home situations; and,
2. An initial set of design implications are proposed to guide researchers in the future design of home-based media spaces that can adequately balance privacy and awareness.

Our first contribution reinterprets the results found by Boyle et al. [2000]. While they found that blur filtration can balance privacy and awareness for mundane work situations, we found that privacy and awareness are *not* balanced by any blur levels for the risky home situations that we are interested in. Furthermore, user feedback from our study allows us to believe that video obfuscation techniques will in general not work. This result is very important, for it seriously questions the premise that video obfuscation techniques by themselves suffice for privacy protection while still providing awareness. Clearly, other privacy-protecting strategies and technologies are required. While this result may not seem surprising in retrospect, we must remember that many researchers (e.g., Hudson and Smith [1996]; Zhao and Stasko [1998]; Crowley et al. [2000]) have or are pursuing research where video obfuscation is used as a technique for balancing privacy and awareness. Our results show that people do not feel comfortable with relying on such techniques and often they mistrust them; people prefer to use techniques offering more direct control over their privacy. Boyle et al. [2000] did not find this because they did not test situations occurring in home environments where threats to privacy increase. While some of the home situations we have presented are extreme cases that would occur infrequently, once they happen there are real and serious consequences such as violated trust.

Our second contribution is more preliminary: we outline a set of design implications that suggest strategies for balancing privacy and awareness in home media spaces. These strategies are intended as avenues for future exploration. Of significant importance is that people want privacy-preservation methods that alter the media space environment (e.g., turning the camera off, rotating

the camera). This in turn gives them control over the space and feedback into the state of the space. For this reason, we suggested in Section 6 that context-aware computing is a promising approach to privacy regulation. A good context-aware environment offers little additional overhead for the user: sensing technology can be used to understand the context of a location and regulate privacy accordingly, and ambient displays can provide the user with appropriate feedback of the state of the media space. As such, we are currently designing and evaluating a context-aware home media space for balancing privacy and awareness. In our home media space, privacy and awareness is balanced with explicit and implicit actions using contextually-aware privacy rules. For example, actions such as blocking the camera, standing up out of the desk chair, or walking into the home office, will adjust the achieved level of privacy. Visual and audio feedback centered on the camera makes this level easily discernable at any time.

Video will proliferate as a technique for providing awareness between those separated by distance, for example, video conferencing for work or between family and friends, live webcams. The question is not whether we should use video in homes; this is already happening and cannot be stopped. The real question is: what are the problems with using video and how might we solve them? We have outlined these problems by looking at one application of video use in a home: home-based telecommuting. However, the ideas we present contribute to home-based video conferencing in general. Many people desire to use technology such as video in their homes and it is important that designers of home-based video conferencing software realize that the privacy risks are real and their designs must address them.

ACKNOWLEDGMENTS

We are grateful to the two actors who performed in our video scenes, and past reviewers for their valuable comments on how to improve an earlier draft of this article.

REFERENCES

ALTMAN, I. AND CHEMERS, M. 1980. *Culture and Environment*. Wadsworth Publishing Company, 1–12, 75–119, 155–214.

BELLOTTI, V. 1998. Design for privacy in multimedia computing and communications environments. In *Technology and Privacy: The New Landscape*, Agre and Rotenberg, Eds. MIT Press, Cambridge, MA, 63–98.

BELLOTTI, V. 1996. What you don't know can hurt you: Privacy in collaborative computing. In *Proceedings of the Conference on HCI (HCI'96)*. Springer, 241–261.

BELLOTTI, V. AND SELLEN, A. 1993. Design for privacy in ubiquitous computing environments. In *Proceedings of the Third European Conference on Computer-Supported Cooperative Work (ECSCW'93)*. Kluwer Academic Publishers, Milan, 77–92.

BLY, S., HARRISON, S., AND IRVIN, S. 1993. Media spaces: Bringing people together in a video, audio, and computing environment. In *Comm. ACM 36*, 1, ACM Press, 28–46.

BOYLE, M., EDWARDS, C., AND GREENBERG, S. 2000. The effects of filtered video on awareness and privacy. In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW '00)*. [CHI Letters 2(3)], ACM Press, 1–10.

BOYLE, M. AND GREENBERG, S. 2005. The Language of Privacy: Learning from video media space analysis and design. In *ACM Trans. Comput. Human Interaction (TOCHI) 12*, 2, June. ACM Press, 328–370.

CROWLEY, J. L., COUTAZ, J., AND BERARD, F. 2000. Things that see. In *Comm. ACM 43*, 3. ACM Press, 54–64.

ERICKSON, T. 2002. Some problems with the notion of context-aware computing. In *Comm. ACM 45*(2), February. ACM Press, 102–104.

FISH, R. S., KRAUT, R. E., RICE, R. E., AND ROOT, R. W. 1993. Video as a technology for informal communication. In *Comm. ACM*, *36*, 1, ACM Press, 48–61.

GREENBERG, S. 1996. Peepholes: Low cost awareness of one's community. In *Companion Proceedings of the Conference on Human Factors in Computing Systems (CHI'96)*. ACM Press, 206–207.

GREENBERG, S. AND KUZUOKA, H. 2000. Using digital but physical surrogates to mediate awareness, communication and privacy in media spaces. *Per. Tech. 4*, 1. Elsevier.

HUDSON, S. E. AND SMITH, I. 1996. Techniques for addressing fundamental privacy and disruption tradeoffs in awareness support systems. In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW'96)*. ACM Press, 248–257.

JANCKE, G., VENOLIA, G. D., GRUDIN, J., CADIZ, J. J., AND GUPTA, A. 2001. Linking public spaces: Technical and social issues. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI 2001)*. ACM Press, 530–537.

KRAUT, R., EGIDO, C., AND GALEGHER, J. 1988. Patterns of contact and communication in scientific observation. In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW '88)*. ACM Press, 1–12.

LEE, A., GIRGENSOHN, A., AND SCHLUETER, K. 1997. NYNEX Portholes: Initial user reactions and redesign implications. In *Proceedings of the International Conference on Supporting Group Work (Group'97)*. ACM Press, 385–394.

MANTEI, M., BAECKER, R., SELLEN, A., BUXTON, W., MILLIGAN, T., AND WELLMAN, B. 1991. Experiences in the use of a media space. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI'91)*. ACM Press, 203–209.

NEUSTAEDTER, C. AND GREENBERG, S. 2003a. The design of a context-aware home media space for balancing privacy and awareness. In *Proceedings of the Fifth International Conference on Ubiquitous Computing (UbiComp 2003)*. Springer-Verlag, 297–314.

NEUSTAEDTER, C. AND GREENBERG, S. 2003b. The design of a context-aware home media space. In *Video Proceedings of the Fifth International Conference on Ubiquitous Computing (UbiComp 2003)*.

PALEN, L. AND DOURISH, P. 2003. Unpacking privacy for a networked world. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI'2003)*. ACM Press, 129–136.

TANG, J. C., ISAACS, E., AND RUA, M. 1994. Supporting distributed groups with a montage of lightweight interactions. In *Proceedings of the Conference on Computer-Supported Cooperative Work (CSCW'94)*. ACM Press, 23–34.

WANT, R., HOPPER, A., FALCÃO, V., AND GIBBONS, J. 1992. The active badge location system. In *Proceedings of Transactions on Information Systems, 10*, 1, Jan. ACM Press, 91–102.

ZHAO, Q. A. AND STASKO, J. T. 1998. Evaluating image filtering based techniques in media space applications. In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW'98)*. ACM Press, 11–18.

Received November 2003; revised March 2006; accepted March 2006 by Paul Dourish