THE UNIVERSITY OF CALGARY

Balancing Privacy and Awareness in Home Media Spaces

by

Carman Gerard Neustaedter

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES

IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE

DEGREE OF MASTER OF SCIENCE

DEPARTMENT OF COMPUTER SCIENCE

CALGARY, ALBERTA

MAY 2003

THE UNIVERSITY OF CALGARY

FACULTY OF GRADUATE STUDIES

The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies for acceptance, a thesis entitled "Balancing Privacy and Awareness in Home Media Spaces" submitted by Carman Gerard Neustaedter in partial fulfillment of the requirements for the degree Master of Science.

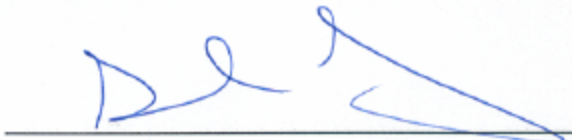

Supervisor, Saul Greenberg
Department of Computer Science

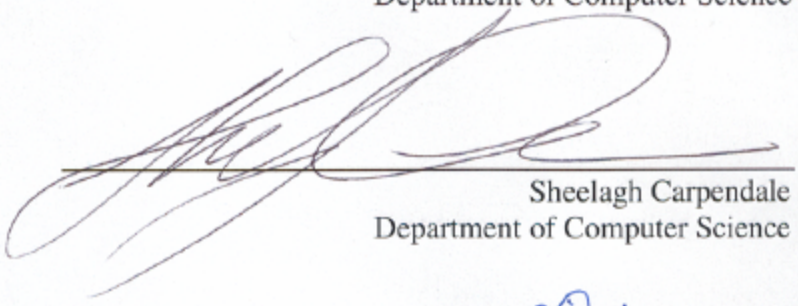

Sheelagh Carpendale
Department of Computer Science

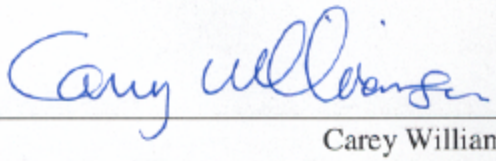

Carey Williamson
Department of Computer Science

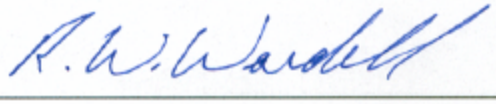

External Examiner, Ron Wardell
University of Calgary

May 29, 2003
Date

Abstract

Always-on video provides rich levels of awareness for collaborators separated by distance, yet it has the potential to threaten privacy as sensitive details may be broadcast to others. This threat increases for telecommuters who work at home and connect to office-based colleagues using video. In this thesis, I address the problem of how one can develop and evaluate privacy-protecting strategies and user interface design techniques for balancing privacy with awareness in a *home media space (HMS)*—defined as an always-on video media space used within a home setting. First, using a controlled experiment, I show blur filtration is not able to balance privacy and awareness for typical home situations involving a telecommuter. Second, I develop a framework for the design of a HMS that identifies a set of appropriate privacy-protecting strategies. Third, I present the rationale and prototype design of a context-aware home media space, designed to balance privacy and awareness for telecommuters and others in the home.

Publications

Materials, ideas, and figures from this thesis have appeared previously in the following publications:

Neustaedter, C., Greenberg, S. and Boyle, M. (2003) **Balancing Privacy and Awareness for Telecommuters Using Blur Filtration.** Report 2003-719-22, *Department of Computer Science*, University of Calgary, January, 2003.

Neustaedter, C., and Greenberg, S. (2003) **The Design of a Context-Aware Home Media Space for Balancing Privacy and Awareness.** Report 2003-722-25, *Department of Computer Science*¸ University of Calgary, May, 2003.

Acknowledgments

This thesis would not have been possible without the help and support of many people.

To Saul, my supervisor and mentor: thank you for inspiring me to pursue graduate school and continually motivating me to be the best that I can be. When I lost sight of my path or did not know which route to take, you were always there to guide me. You have taught me so many things in such a short amount of time Saul.

To my friends from the Interactions Lab: thank you for your ideas, insight, and most of all, friendship. You were there when I needed advice or an opportunity to unwind. You made university fun and an experience I will never forget.

To Michael Boyle: thank you for your knowledge and laying the intellectual foundation for my work including your theory of privacy in video media spaces and previous study of distortion filtration. Your knowledge and understanding of privacy and video media spaces is unsurpassed and your help with my thesis was invaluable.

To Adam: thank you for providing me with your friendship and advice when I needed it most. Despite the distance that has separated us, you have continued to be an invaluable friend. I will always treasure our talks about life and the experiences we shared.

To my parents away from "home," Karen and Rick: you welcomed and loved me with open arms. You have both taught me many lessons in life and I will always cherish your advice. Karen, thank you for your guidance and making sure I never left your place with an empty stomach. Rick, thank you for the "guy talks" and playing stress-relieving video games with me.

To my favourite sister, Marlene: thank you for your support and guidance through life. You were always there when I needed to talk, whether it be about school, life, or love. You have showed me how to be courageous and know that life always has a light at the end of the tunnel.

To my wife, Kirstin: I can not thank you enough. Your love has and always will be my guiding light. Thank you for "putting up with" my privacy-intrusive research, stressful times, and busy schedule. You give me the strength to pursue my dreams.

Dedication

I dedicate this thesis to my parents, Norma and Paul. You watched me grow up, taught me about life along the way, and then supported me as I left home to pursue my dreams. You have always guided me, loved me, and wanted the best for me. I am very proud to have you as my parents and know that I have made you proud with my accomplishments.

Table of Contents

List of Tables

List of Figures

Chapter 1. Introduction

In this thesis, I address the problem of how one can develop and evaluate privacy-protecting strategies and user interface design techniques for balancing privacy with awareness in home-based video media spaces. To set the scene, I begin this chapter with a brief overview of existing research on how technology—particularly video-based media spaces—can provide informal awareness between distance-separated intimate collaborators. Next, I discuss privacy issues inherent in the interface design of video media spaces. Finally, I present the specific problems of privacy issues in home-based media spaces and how I will solve each problem in this thesis. I conclude with an organizational overview of this document.

1.1 Background

Throughout a typical day, co-workers naturally converse and interact amongst each other in what is known as *casual interaction*—the frequent and informal encounters that either occur serendipitously or are initiated by one person (Fish et al., 1993, Hudson and Smith, 1996). Casual interactions foster knowledge and help individuals accomplish both individual and group work (Kraut et al., 1988, Fish et al., 1993). My particular interest is in casual interactions between *intimate collaborators,* defined as those individuals who have a real need or desire for close coordination and communication (Greenberg, 1996). *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 (Kraut et al., 1988, Bellotti and Sellen, 1993, Gutwin et al., 1995). 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 co-workers are distributed.

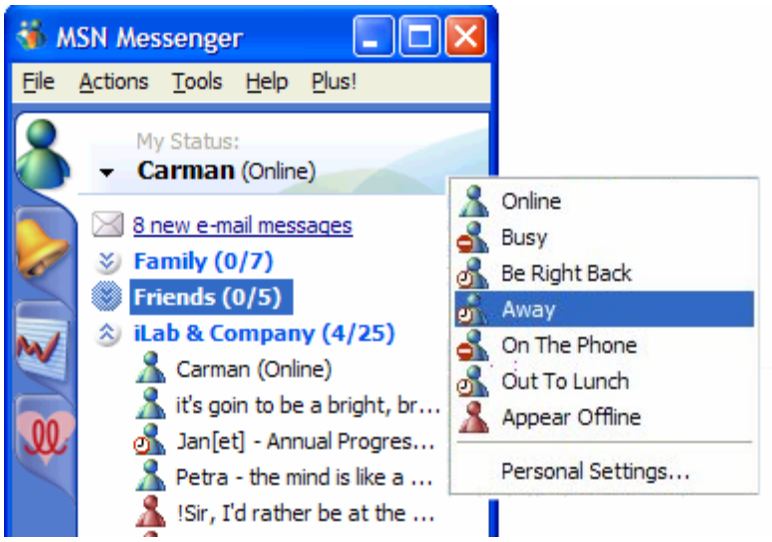

Figure 1.1: Availability states in MSN Messenger.

A variety of existing techniques exist to provide informal awareness for distance-separated collaborators, yet some provide better awareness than others. One popular approach for gaining awareness is the use of availability states, i.e., online, away, busy, in instant messengers, such as MSN Messenger (Figure 1.1) or ICQ. Here, idle indicators change a user's state automatically or users are able to select an availability state, e.g., using the pop-up menu in Figure 1.1. Although very useful, these low fidelity states can provide less than an ideal description of the actual availability of a collaborator because they indicate presence rather than availability, and even this is just an approximation. As a result, the privacy of the collaborator is at risk because co-workers can distract them by interrupting at an inappropriate time.

My particular interest lies in providing informal awareness across distance through the use of a *video media space*. A video media space uses always-on video to capture the scene around a potential collaborator and broadcast it to others in the workgroup (Mantei et al., 1991). Video is capable of providing rich awareness because one can actually see the other person, much like in co-located settings. Yet video comes with many privacy risks, even when used between intimate collaborators in benign settings such as work offices. Rather than seeing someone across the room or in a different office as is normally the case, a video media space can make it appear as though a colleague is sitting close by. Figure 1.2 shows a typical media space where two distance-separated colleagues gain informal awareness using a video channel while they work. Here, both

Figure 1.2: A typical video media space where video is used to provide informal awareness between distance-separated colleagues. PC cameras (circled) capture each collaborator and this information is broadcast to the other colleague (circled on the display of the person on the left).

collaborators have a video camera (circled) that captures and broadcasts their image to the other person; for example, the person on the left sees a closely cropped image of the person on the right (circled).

Video media spaces have been installed and tested in office and research lab settings (e.g., Fish et al., 1990, Mantei et al., 1991, Dourish and Bly, 1992, Bly et al., 1993, Dourish, 1993, Tang et al., 1994, Lee at al., 1997, Coutaz et al., 1998, Greenberg and Kuzuoka, 2000). These spaces typically connect friends and peers who inhabit similar organizational settings and who either are early adopters of technology or have a vested interest in the system. The situation is complicated with *telecommuters*: people who choose to work from home. Many telecommuters still desire close contact with colleagues at the office and use technologies such as email and instant messaging to maintain a limited amount of awareness. As with office-based media spaces, a *home media space*, defined as an always-on video media space used within a home setting, can also provide a rich level of awareness for telecommuters by connecting them with their colleagues at the office. Yet privacy risks increase dramatically. The main problem is that the home is not the office; activities, people, and appearances that are appropriate for the home are often inappropriate when viewed in an office environment by a colleague. For example, it is appropriate to work at home shirtless on a hot summer day, while the same level of dress is inappropriate for most offices. People normally need an emotional release and the privacy of their home allows them to relax and often deviate from social

Figure 1.3: An unfiltered media space view, the view with a blur filter, and the view with a pixelized filter.

customs that they regularly adhere to on a daily basis (Altman, 1975). By introducing video media spaces into homes, the privacy of the telecommuter and others in the home can be greatly threatened. These risks are discussed in detail in Chapter 2.

Methods such as *distortion filters* have been studied to find a reasonable trade-off between providing awareness and preserving privacy in video media spaces (Zhao and Stasko, 1998, Greenberg and Kuzuoka, 2000, Boyle et al., 2000). Distortion filters such as *pixelize* or *blur filters* attempt to preserve privacy by filtering out sensitive information while still providing a level of awareness. Figure 1.3 shows three images of the same collaborator: the left image is unfiltered, the middle image is distorted with a blur filter, and the right image is distorted with a pixelize filter. In using such methods, the amount of awareness decreases because information of a lower fidelity is being broadcast from the video media space. Similarly, as awareness levels increase, privacy decreases as more detailed information from the video media space is broadcast to collaborators. While research has shown distortion filters, such as the pixelize and blur filters, are able to balance privacy and awareness in office situations (Zhao and Stasko, 1998, Boyle et al., 2000), it is not clear if such techniques are suitable for far riskier home situations.

In conjunction with methods to balance privacy and awareness comes a necessity for simple, lightweight user interfaces for video media spaces. These user interfaces typically afford various strategies for presenting privacy *feedback* and *control*. *Feedback* allows users to know whether or not they are attaining their desired level of privacy. Bellotti (1996, 1998) outlines that feedback in a media space involves "informing people when and what information about them is being captured and to whom the information is being made available." Once media space participants know how much privacy is being

attained, they need the ability to adjust their current level of privacy to a desired level. This comes in the form of privacy *control* and as Bellotti (1996, 1998) points out, control involves "empowering people to stipulate what information they project and who can get a hold of it."

When presenting privacy control and feedback, two main problems exist with user interfaces for video media spaces:

a) ***The interface makes it difficult to alter privacy levels.*** If users are not presented with an interface that can easily alter privacy levels, they may resort to doing away with the video link completely. While this gives complete privacy, it comes at the cost of no awareness of their colleagues. Alternatively, they may do nothing and risk having no privacy at all.

b) ***The interface does not provide sufficient feedback of privacy levels attained.*** With insufficient feedback, again, the user may resort to doing away with the video link because of a fear of too much information being broadcast. This again would jeopardize awareness between collaborators. Cues of the level of privacy being maintained may also help users to properly *appropriate* themselves, defined as the act of creating a socially acceptable appearance and/or behavior (Bellotti, 1998).

It is clear that without adequate user control and sufficient feedback of privacy, video media spaces are unable to accomplish the task of providing informal awareness.

1.2 Thesis Problems

This thesis is about privacy in video media spaces used between telecommuters and office workers. Figure 1.4 illustrates the context and scope of my research. In particular, I address the following problems in this thesis:

1. **We do not know if blur filtration is able to balance privacy and awareness in a home media space.** Previous research (Boyle et al., 2000) has shown that distortion filters, such as the blur filter, are able to balance privacy and awareness for benign office situations. Yet we do not know if this balance is achievable for home use of video, as home situations present far riskier situations than office environments.

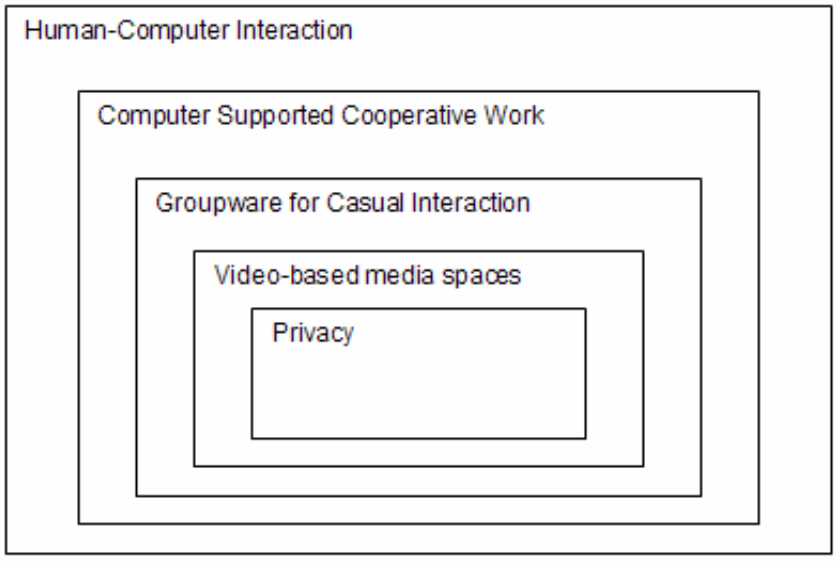

Figure 1.4: The context and scope of my research.

2. **We do not know what other privacy-protecting strategies, if any, are appropriate for balancing privacy and awareness in a home media space.** Research on privacy-protecting strategies for video media spaces has again primarily focused on office settings, rather than homes. It is unclear what other privacy-protecting strategies, aside from distortion filters, may be suitable for balancing privacy and awareness in home settings.

3. **We do not know what user interface techniques are appropriate for presenting users with privacy-protecting strategies in a home media space.** Privacy-protecting strategies for balancing privacy and awareness in a home media space must be presented to users in a simple, lightweight user interface. Research has previously focussed on designing video media spaces for office situations rather than home-settings.

1.3 Thesis Goals

In this thesis, I will address the aforementioned problems with the following goals:

1. **I will evaluate blur filtration for its effectiveness in balancing privacy and awareness in a home media space.** I will define and run a controlled experiment that will evaluate blur filtration's ability to balance privacy and awareness for home

situations that vary in the amount of perceived privacy risk presented, from little or no risk to very high risk *(Problem 1)*.

2. **I will investigate other privacy-protecting strategies for balancing privacy and awareness in a home media space.** I will conduct a literature review on privacy mechanisms within social-psychological theory, looking at mechanisms used in face-to-face situations by various cultures. Based on this literature review, previous research in video-media spaces, and results from the experiment in Goal 1, I will outline a framework for the design of a home media space, which will describe other potential privacy-protecting strategies *(Problem 2).*
3. **I will design a home media space that presents users with privacy-protecting strategies.** Using the framework developed in Goal 2, I will design a home media space, which will present user interface techniques that are appropriate for affording users with privacy-protecting strategies in a home *(Problem 3)*. The home media space design will not be formally evaluated, yet will present one approach for the design of such a space and the use of the framework from Goal 2.

1.4 Organizational Overview

This thesis is divided into seven chapters:

In Chapter 2, I present a literature survey of privacy and video media spaces. I begin with the motivation for this work, which is supporting awareness to promote casual interaction for distance-separated collaborators. Next, I discuss social-psychological theories of privacy and how they relate to video media spaces and home environments. Then I present existing research on privacy preservation techniques for video media spaces.

In Chapter 3, I discuss the methodology for a controlled experiment that evaluates one privacy preservation technique, blur filtration, for its effectiveness in balancing privacy and awareness for home situations containing a telecommuter *(Goal 1)*. The study looks at a series of typical home situations that vary in risk from an expected low risk to an expected high risk. Chapter 3 includes an outline of the study's null hypotheses, variables, materials, and procedure.

In Chapter 4, I discuss the results of the controlled experiment defined in Chapter 3, which includes an analysis of blur filtration levels that provide users with awareness cues, along with blur filtration levels that adequately preserve privacy *(Goal 1)*. I also look at what blur filtration levels people choose to use for home situations, as well as how willing they are to use a home media space. I conclude the chapter with a set of design implications for privacy-protecting strategies to be used in the design of a home media space. These implications articulate the difficulties in designing strategies for balancing privacy and awareness in home media spaces.

In Chapter 5, I take a step back and summarize research on privacy mechanisms used by various cultures to regulate and control privacy in face-to-face situations. Next, I use this research and the experiment results discussed in Chapter 4 to develop a framework for designing a home media space *(Goal 2)*. This framework contains a set of privacy-protecting strategies that can be used within a home media space to afford the user with control and feedback of privacy.

In Chapter 6, I discuss the design of a home media space that uses the design framework presented in Chapter 5. This involves discussing user interface design principles for providing users with a plethora of privacy-protecting strategies *(Goal 3)*. The home media space design is not formally evaluated, but presents one approach for the design of such a space and the use of the framework presented in Chapter 5.

In Chapter 7, I conclude this thesis by summarizing how I achieved each of my research goals. I also list my research contributions and suggest areas for future work in home media spaces.

Chapter 2. Video Media Spaces and Privacy

In this chapter, I set the scene with a literature review of video media spaces and privacy. My goal is to provide the reader with sufficient background knowledge of how video media spaces support casual interaction and awareness, and also the privacy issues that arise when using such spaces. First, I discuss the importance of casual interaction and its relationship with informal awareness. Second, I summarize social-psychological definitions of privacy and their relationship with home situations involving telecommuters who use a video media space. I conclude by reviewing existing research on design techniques that have been used to balance privacy and awareness in video media spaces. This chapter will act as a motivation and foundation for future chapters, where the techniques and theories presented here will have a direct impact on my work involving home media spaces.

2.1 Supporting Awareness for Collaborators

In this section, I discuss the motivation behind media spaces, which is supporting casual interaction and informal awareness between distance-separated collaborators. First, I describe in more detail the importance of casual interaction for co-workers and how informal awareness supports it. Second, I show how informal awareness can easily be lost when collaborators become separated over distance. Finally, I discuss how video media spaces can provide rich awareness over distance by utilizing a visual channel.

2.1.1 Casual Interaction

Casual interactions are unstructured meetings and exchanges of information between co-workers. Casual interaction is also referred to as informal communication, distinguishing it from traditional formal communication (e.g., scheduled meetings) through official

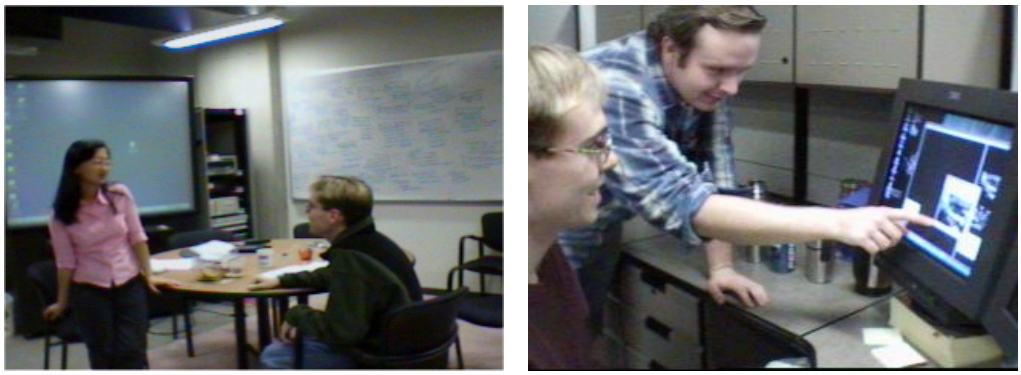

Figure 2.1: Casual interactions: two co-workers converse as one comes into the office (left); and, two co-workers discuss computer software when one asks the other a question.

structures within a company (Fish et al., 1990). Serendipitous or one-person initiated encounters that occur between workers throughout a typical day can easily result in casual conversation. Casual interactions therefore are spontaneous, frequent, and usually initiated by one person (Fish et al., 1992, Hudson and Smith, 1996). Depending on a person's job type, research has shown that 25 to 70% of a person's time is spent in face-to-face interaction, and that informal encounters comprise almost a third of all office activity (Whittaker et al., 1994). Often individuals will meet while passing through a hall, strike up a conversation in an elevator, or even discuss new ideas on a coffee break. The number of chance encounters and the potential they provide is endless. Figure 2.1 (left) shows casual interaction between two co-workers who briefly met as one came into the office. Figure 2.1 (right) shows two co-workers discussing some computer software after a casual question led to conversation. While seemingly unimportant, recurring casual interaction proves crucial, for this is how people maintain loyalty between themselves, build relationships, coordinate activities, accomplish work, and foster knowledge (Kraut et al., 1988, Fish et al., 1992, 1993).

2.1.2 Informal Awareness

Casual interaction is made possible through *informal awareness* (Kraut et al., 1988, Bellotti and Sellen, 1993, Gutwin et al., 1995). Informal awareness is a naturally gained understanding of who is around, what tasks they are performing, and whether or not they are available for conversation or collaboration. Informal awareness is the glue that holds casual interaction together. It provides cues that influence how people interact with

Figure 2.2: Gaining awareness from the amount of work piled on a desk, the absence of a worker from his or her desk, and a look of intense concentration.

others, for example, whether or not one decides to initiate a conversation with a colleague (Bellotti and Sellen, 1993). Figure 2.2 shows that cues of informal awareness can come from various sources: the amount of work piled on someone's desk (far left), the absence of another at his or her desk (middle), or a look of intense concentration from a co-worker (far right). One may see that the co-worker one needs to speak with is currently on the phone, or that a collaborator just arrived at the office. People garner these cues very easily as a result of working in a social environment, e.g., by looking around their shared office and by simply walking down a hallway and glancing into open office doors.

2.1.3 The Problem of Distance

Informal awareness is easy to obtain for those in the same physical environment (Greenberg, 1996). However, when people are separated by distance, awareness disappears unless technology is there to support it. Distance can mean several things: people can be separated between floors of a building, located in different buildings, or, even worse, situated in different cities. A study by Kraut et al. (1988) found that researchers from both academia and industry whose offices are close together are six times more likely to collaborate than those separated by distance. The distances do not have to be large; Kraut showed an exponential drop-off of interactions as distances increased. Separation of colleagues and workers becomes even more prevalent with an increasing number of home workers or telecommuters (Whittaker et al., 1994). Without informal awareness, people are unable to easily move into and out of tightly coupled casual interaction.

2.1.4 The Importance of Video

Although sound is often necessary for communication, people's visual channel predominates in gathering many visual awareness cues. After analyzing episodes of casual interaction, Fish et al. (1990) found that the visual channel was used to initiate all of the observed interactions. They also note that the visual channel was crucial for locating and identifying a colleague. According to Whittaker and O'Conaill (1997), visible behaviors are monitored with very little conscious effort by people and provide cues to help people move into and maintain conversations.

There are now many technologies that help people interact casually over distance, e.g., instant messaging (Nardi et al., 2000). Of these, video provides the richest awareness to collaborators because it displays many visual awareness cues in a natural and easily understandable manner. As well, video seamlessly supports both awareness and conversation; thus, it's easy for people to act on their awareness information simply by moving into conversation over the video link.

Mantei et al., (1991) defined a *video media space* as a system with an always-on video connection, which provides the rich visual channel needed for informal awareness. This awareness can then be used to smoothly move people into and out of casual interaction while separated by distance (Mantei et al., 1991). In a sense, a video media space provides a surrogate of face-to-face physical presence over distance, allowing collaborators to use the same visual cues for gaining awareness as in co-located settings. The problem is that while video media spaces are able to provide rich awareness, they also have the potential to broadcast sensitive details and violate the privacy of collaborators.

2.2 Social-Psychological Views of Privacy

When video media spaces are used to provide informal awareness, privacy issues undoubtedly arise from the transmission of information that may be considered privacy sensitive. In a home setting, this information can be extremely sensitive as it is coming from a location that is largely regarded by its occupants as being private. People working at an office expect that certain actions will be observed by others, yet this is not always

the case for homes. In this section, I outline social-psychological definitions of privacy and their relationship with home situations involving video. First, I define privacy and several important ways it can be violated. Second, I present several case studies of typical home situations involving a home media space, used to illustrate the privacy definitions, and articulate the issues facing telecommuters and others in the home when video is used.

2.2.1 Definitions of Privacy

The concept of privacy is often quite vague. Yet most people know when their privacy has been compromised and these privacy violations can even raise deep emotional responses. The problem is that privacy is many things and in order to define it, its meaning must be deconstructed. Privacy, as defined by Altman (1975) in general terms, is an "interpersonal boundary-control process, which paces and regulates interaction with others." At certain times people seek more privacy, while at others they seek interaction. Altman notes that privacy is like a cell membrane that continuously alters its permeability. This shifting of permeability either creates more privacy for an individual by closing off external contact or allows more contact from other individuals. Thus, interaction is tightly coupled with privacy (Altman, 1975). To create an optimal state of privacy, the individual's desired privacy must be equal to his or her actual achieved privacy (Altman, 1975).

Boyle (2003) has deconstructed Altman's definition of privacy to create a theory of privacy in video media spaces. Boyle's theory states that privacy can be violated in several fundamental ways by a video media space: an invasion of solitude, a breech of confidentiality, a loss of autonomy, or a combination of these. As Boyle explains,

- ***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. Solitude is normally only threatened when interaction occurs.
- ***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.

While these privacy violations may occur because of a media space's design, participants in a media space may also fail to *appropriate* themselves correctly for their current situation (Bellotti, 1998). *Appropriation* is the act of creating a socially acceptable appearance and/or behavior for a given situation (Bellotti, 1998, Boyle, 2003). Often individuals are unaware of how to appropriate themselves because they are not properly informed of the given situation, or because they may be in a mixed context (to be discussed shortly).

The next section discusses five case studies, which will illustrate these privacy violations, as well as outline privacy issues for telecommuters and others in the home. Each of the case studies is very realistic and is derived from actual events reported to me by telecommuters who habitually use always-on video connected to work colleagues. The point of these case studies is both to set the scene and to show how privacy violations can be explained using our definition.

2.2.2 Case Study 1: Working with No Shirt

The first case study involves a telecommuter living in a dual role as worker and as home occupant.

It is a hot day and Linda is shirtless wearing only a bra. Forgetting her attire (because this is not a problem in the home context), she enters her home office to quickly check her email. The always-on video captures Linda shirtless and she (too late) notices her colleagues aghast at this public display.

First, we can see that Linda's **confidentiality** is being breeched; *her colleague has learned more about her than is desired*. The colleague has now seen Linda wearing just her bra. Second, the media space is threatening Linda's **autonomy** because she wishes to work without being fully clothed, yet this has unwanted consequences; *she is no longer able to choose how she participates in the space*. Third, Linda is appropriately dressed for home, but definitely not for an office and it is very unlikely that such an appearance would even be seen at an office; *she has failed to properly **appropriate** herself*.

2.2.3 Case Study 2: Picking One's Nose

The second case study results from a telecommuter's unconscious acts, the ease of forgetting that a distant colleague is (virtually) sitting right across from him, and from the lack of feedback that he is actually in a public setting. While unconscious acts can be seen at an office, they are more likely to be seen at home because people typically relax from social customs when at home.

Larry is working on his home computer when he suddenly sneezes. Naturally, he proceeds to blow his nose, forgetting that a camera on top of his monitor captures this at a *very* close range. Larry now begins to pick his nose at great length. His colleague views the scene over the video link and is disgusted at how inconsiderate Larry is being, while really Larry had not intended for him to see this.

First, one can expect that all people pick their nose at some point or another, yet now Larry's colleague is more certain that Larry actually does pick his nose. Larry's **confidentiality** has been breeched because *his colleague knows more details about him then he would like.* Second, Larry would like to pick his nose without others seeing, but the media space has threatened his **autonomy** because he is unable to do so; *he has lost control over how he is being recorded.* Third, while Larry is clothed appropriately for the office, his behavior is not appropriate for others to see in detail. Many would tolerate seeing this accidentally at a distance, yet now they are seeing a close-up, somewhat disgusting view of it; *Larry has failed to properly* ***appropriate*** *himself*.

2.2.4 Case Study 3: Kissing a Partner

A privacy risk also arises for other family members and friends in the home who may gain no benefit from the video link yet still incur its privacy threat. A *threat/benefit disparity* occurs when the benefit that an individual receives from a media space is unequal to the threat that comes with it (Boyle, 2003). The third case study is representative of typical activities between individuals within a home that one may not want others to see, e.g., disciplining children, arguing with a partner, or showing affection.

Linda is working in her home office in the early morning when her lover (who has just woken up) enters the room wearing only pajamas and gives Linda a big wet kiss. All this is captured on camera, and Linda's colleague has seen this much to her dismay. Linda's lover realizes this and becomes infuriated; he tells her never to use the camera again.

First, the **confidentiality** of Linda and her lover is being breeched because they do not want others to know details of their personal life, yet now *her colleague knows too many details about their personal lives*. Second, Linda and her lover wish to kiss, but *they do not want others to see it*; their **autonomy** is being threatened. The threats to autonomy and confidentiality are worse for Linda's lover because he gains no benefit from the media space, yet still incurs the privacy threat. One can imagine that this could have serious consequences if this were a surreptitious relationship. Third, while Linda is appropriate in appearance, her lover's appearance is not appropriate for an office; *he has failed to create an **appropriate** appearance* for the situation. Furthermore, *both Linda's and her lover's behavior is not **appropriate*** for office environments and would not normally be seen by others (albeit such affairs may occur behind closed office doors).

2.2.5 Case Study 4: Shown Naked

The fourth case study is representative of situations that one clearly does not want colleagues to see and results from the dual purposes typical of most home offices.

Larry's home office is also his spare bedroom. One hot day Larry takes a shower in the bathroom next door. He towels off, and then goes into the spare bedroom to retrieve a bathrobe in the closet. A few moments after entering the room, Larry realizes that the camera is directed at him. He drops to the floor, crawls to the camera, and knocks it off the computer.[1]

First, *Larry is showing his colleague more than he wants to* and his **confidentiality** is breeched as a result. Second, while Larry wants to dress, *he does not want others to view this* act nor his naked body—clearly his **autonomy** is violated. Third, by being

[1] Reported by Saul Greenberg during personal communication.

naked there is no way Larry is appropriated correctly for an office environment; *he has failed to create an **appropriate** appearance* for this situation.

2.2.6 Case Study 5: Interrupted During a Break

The fifth case study involves a telecommuter working at home and being interrupted at an inappropriate time.

Linda sits down at her computer desk and forgets to enable the always-on video link. While she has spent the day working on a report, Linda takes a break when her young daughter comes into the room to discuss a school project. Linda's colleague is gaining no awareness information from the media space and does not realize that Linda is busy with her daughter. The phone rings and it's Linda's colleague calling with a question. Linda quickly answers the question with impatience and annoyance at being interrupted during her break.

This case study does not threaten **confidentiality** and **autonomy** as it is a fairly mundane situation where Linda has created an **appropriate** appearance and behavior for a colleague to view at an office. If the media space were on, Linda's daughter may face privacy risks. In this case, Linda's **solitude** is being violated because *she is being interrupted at an inappropriate time.* Linda is busy meeting with her daughter, yet because the colleague has lost all awareness information he cannot easily decide if now is a good time to interact with Linda.

2.3 Privacy Preservation Techniques

While video media spaces can provide awareness, they also introduce privacy risks; in home settings, privacy risks arise for both telecommuters and other individuals in the home and these can be extremely threatening. The problem is that an increase in the amount of awareness information transferred provides an increasing level of awareness, yet the more information transferred, the greater the privacy risk (Hudson and Smith, 1996). Too little an amount of awareness can also cause invasions of another's solitude by those wishing to engage in casual interaction with no knowledge of the activities of another (Bellotti, 1998).

In an effort to help mitigate privacy concerns over video links, researchers have studied techniques to provide a balance between awareness and privacy. Boyle (2003) outlines four thematic approaches to preserving privacy in video media spaces: pull the plug, mirror, reciprocity, and fidelity reduction. All offer potential techniques that may be viable solutions for preserving privacy in a video media space within a home setting. This section discusses each of Boyle's privacy preservation techniques in turn and outlines existing media space designs that have incorporated the techniques, as well as some that have failed to do so. Existing media space research will be used in subsequent chapters to guide the design of home media spaces.

2.3.1 Pull the Plug

The easiest and most certain way for most people to preserve privacy is to "pull the plug" on the camera; this type of action was seen in the fourth case study (Section 2.2.5). As Boyle (2003) discusses, a "pull the plug" mechanism is simply an easy method for disabling the capturing device or software. This includes both physically unplugging the camera, and lightweight techniques to block or turn off the camera. For example, Figure 2.3 shows a camera turned to face a wall (left), a plastic visor used to block a camera's lens (middle), and a software control for turning the camera off (right). Of these, the last is least satisfactory for the person who has no certainty (or trust) that the software has stopped transmitting the image. Boyle (2003) notes that users require opportunities for complete privacy; thus, "pull the plug" mechanisms are mandatory for all media space designs.

Figure 2.3: Three example "pull the plug" mechanisms for controlling privacy: turning the camera to face a wall (left), flipping down a plastic visor (middle), and toggling a software control (right).

Several problems exist with "pull the plug" mechanisms for balancing awareness with privacy. First, "pulling the plug" completely eliminates any awareness. A side effect is that it can slightly increase privacy violations through untimely interaction. This was found in the fifth case study when Linda's video link was not enabled and her colleague interrupted while Linda was talking with her daughter. Here, Linda's colleague had no awareness information to help decide if and when to move into interaction. Second, many media space designs cause additional awareness problems by not making it easy to reverse a "pull the plug" operation. For example, one may be inclined to simply start unplugging cables at the back of her computer if her privacy becomes greatly at risk. When awareness is next desired, the user may have to painstakingly find out what cord is attached to the camera and where to plug the cord into the back of her computer, with the chance that she may simply not bother. Third, some "pull the plug" techniques provide little feedback of whether or not the capturing device has actually stopped recording. Figure 2.3 shows a current Logitech™ PC camera (middle), which comes with a plastic visor that can be flipped down to achieve privacy. While the visor is semi-translucent, it is close enough to the camera's lens that only an image of the visor is transmitted. The fact that the camera's lens is blocked indicates to the user that his image will not be seen clearly by others, yet there is still uncertainty about what is being transmitted.

The following three media space designs illustrate either a failure to provide a "pull the plug" mechanism, or design problems with an existing "pull the plug" mechanism.

XEROX PARC's Cruiser system (Fish et al., 1993), designed to support teleconferencing and casual interactions between employees, failed to provide an easy mechanism for completely pulling the plug on the media space. Informal awareness was supported by two main mechanisms in the Cruiser System: users could *Glance* into a colleague's office where they would see a video snapshot of the office, or users could use an *Autocruise* where the system would show a video snapshot from a random office (to mimic casual encounters, such as passing in the hall). *Glances* and *Autocruises* could be initiated at any time. Colleagues on the receiving end (who were captured by the video) had no easy means to block this if they desired privacy. Fish et al. (1993) found that as a

result, management would often physically disable the system by disconnecting wires when sensitive issues were being discussed.

Apple's first prototype of the Virtual Café (Bellotti, 1998) was designed to broadcast images from a local café to a web site so Apple employees could see how busy the café was throughout the day. The camera captured employees of the café, as well as café customers. While all employees had agreed to be captured on video, many customers did not realize a camera was capturing them (a sign warning them was missed by most). Moreover, the system failed to offer café customers a "pull the plug" mechanism that could disable the capturing device. There was nothing customers could do about being recorded (besides leaving the café), if they even realized they were being recorded in the first place!

Microsoft's Virtual Kitchen (Jancke et al., 2001) connected three office kitchens with audio and video channels in an effort to promote social interaction between employees (Figure 2.4). The Virtual Kitchen would record the activities in each kitchen and broadcast this information to the other kitchens. Devices for recording the audio and video were placed throughout the kitchen and an OFF button outside the kitchen could be used to interrupt transmission for those not wanting to participate (Figure 2.4). For example, if you were about to enter the kitchen, but didn't want to be captured you could simply press the OFF button. The situation is different, however, if someone is already present in the kitchen. In the case that you may interrupt a (virtual) conversation in

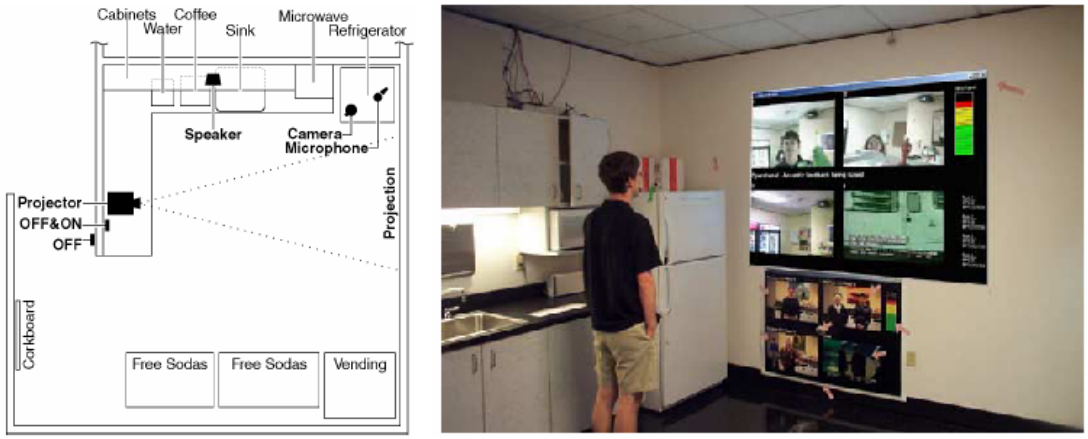

Figure 2.4: Microsoft's Virtual Kitchen: placement of interface components (left) and the projected display (right)—images copied from Jancke et al., 2001.

progress by pushing the OFF button, heat and motion sensors were used inside the kitchen to detect when someone was present. If someone was already present in the kitchen, the OFF button would be disabled and you would no longer have the "pull the plug" mechanism. Although the OFF button was present in this media space design, disabling it, in effect, removed the mandatory "pull the plug" privacy preservation mechanism.

Boyle et al.'s (2000) Nanana media space was an experimental video media space designed to support awareness between two distance-separated collaborators. Nanana uses gesture as input for a "pull the plug" mechanism: users could cover the camera with their hand to stop the transmission of video. When the camera was "blocked," it rotated to face a wall using a servo motor. The same blocking gesture returned the camera to an unblocked state.

2.3.2 Mirror

A second approach to preserving privacy in video media spaces involves providing a mirrored image of what is being captured by the media space (Boyle, 2003) (Figure 2.5). This type of feedback enables users to properly appropriate themselves for the given situation, by letting them see exactly what they will appear like to other media space

Figure 2.5: A sample mirror facility on the right monitor shows the collaborator's own image.

users, e.g., users can position the camera at an appropriate angle, or decide if they are properly groomed. A mirror facility, if reliable and visible, also provides the user with the knowledge of the camera's current state: on or off. Boyle (2003) adds that mirror facilities are natural and lightweight because people are used to seeing themselves in a mirror and adjusting appearances as a result. Figure 2.5 shows a sample mirror facility on the right display where the user can see what is being captured and broadcast to his collaborators.

While mirror facilities do show you what is being captured, they typically fail in showing who is actually viewing the captured scene (Boyle, 2003). It is also important to remember that other applications or screen savers may end up hiding the mirror facility, rendering it useless.

The University of Toronto's CAVECAT system (Mantei et al., 1991), designed to allow group members to meet while distributed, provides an example of a media space designed with a mirror facility. CAVECAT displays video for four media space participants, arranged in a 2 x 2 grid. Users are able to see their own video in one of the grid's cells. CAVECAT participants were able to use the mirror facility to properly appropriate and frame themselves in the camera's view. The mirror facility coincidentally served another purpose; if media space participants happened to be at a co-worker's office that was also running CAVECAT, they could monitor activities happening in their office while they were gone. However, this in turn could violate the privacy of others in that space who were not aware that they were being recorded.

2.3.3 Reciprocity

Media spaces that implement reciprocity ensure that "if you can see someone else, they can see you and that if you can hear someone else, they can hear you." (Fish et al., 1990) Reciprocity can allow you to know who is accessing your captured video (Boyle, 2003).

Reciprocity need not be all or nothing, as is the case in face-to-face conversations: as you move closer to other people you see them more clearly, just as they see you more clearly (Greenberg and Kuzuoka, 2000). This feature of face-to-face communication can be replicated in a media space's design: as you move closer to the capturing device you

see and are seen more clearly (Greenberg and Kuzuoka, 2000). Reciprocity makes deciding the level of detail a cooperative process between media space participants because now both you and your co-workers must decide what is captured by the media space (Boyle, 2003).

Despite the fact that reciprocity in media spaces replicates an attribute of face-to-face communication, various problems exist when attempting to provide reciprocity. First, as Boyle points out, reciprocity can cause a loss of autonomy because you may no longer be able to fully decide *how* you wish to be a part of the space; your decision will now depend on others. This can further be complicated in situations involving a telecommuter and colleagues at the office where varying privacy expectations may exist, e.g., the telecommuter may expect more privacy because she is at home while an office colleague may desire little privacy. As Hudson and Smith (1996) discuss, reciprocity forces both spaces to be public. Second, to properly provide reciprocity, it is necessary to orient and position cameras in appropriate locations. Improper positioning of cameras can cause viewers to be able to stand outside the captured region, yet still able to view others over the video link (Fish et al., 1990). Moreover, camera placement can be awkward, e.g., cables may restrict movement of the camera. Third, reciprocity can create unnatural effects, such as inaccurate sizes and distances. For example, small display sizes can cause others to appear much smaller over a video link than they do in real life. Just the same, a video link can make colleagues appear much closer or further away than they actually are.

A full reciprocity rule is imposed in XEROX PARC's Cruiser system, where users are guaranteed that if they can hear and see someone else, that person can see and hear them (Fish et al, 1992). In the case of NYNEX Portholes (Lee et al., 1997), only a partial degree of reciprocity is made available. Portholes posts video snapshots of workgroup members on a site accessible only by other group members. Here, users are capable of learning who has accessed their own images in the last five minutes. Because the feedback presented here is not immediate, reciprocity is limited.

In Bellcore's VideoWindow (Fish et al., 1990), potential conversations failed to materialize because reciprocity was not provided properly. Users would often either

Figure 2.6: An unfiltered media space view, the view with a blur filter, and the view with a pixelized filter.

stand too close to the capturing device, leaving an unrecognizable image; or, stand out of the camera's capturing area, yet still able to view the video of others, i.e., the viewing angle and capturing angle were not the same (Fish et al., 1990). This situation can occur easily and unknowingly if a mirror facility is not provided.

2.3.4 Fidelity Reduction

Fidelity reduction reduces the captured video's fidelity in order to preserve privacy (Figure 2.6). One type of fidelity reduction is the use of *distortion filters:* an algorithmic reduction of image fidelity that hides sensitive details in a video image (Zhao and Stasko, 1998, Greenberg and Kuzuoka, 2000, Boyle et al., 2000). Two filters include the *pixelize filter* that produces a mosaic of solid rectangles and the *blur filter* that naturally blends regions of an image to produce a blurred effect. Figure 2.6 shows three images of the same collaborator: the left image is unfiltered, the middle image is distorted with a blur filter, and the right image is distorted with a pixelize filter. Here, the amount of awareness decreases because less detailed information is actually being broadcast from the video media space. Similarly, as awareness levels increase, privacy decreases as more detailed information from the video media space is broadcast to collaborators.

A second type of fidelity reduction is found in scene reconstruction. Scene reconstruction uses techniques such as background subtraction and eigen-space filters to remove unnecessary information from images (Crowley et al., 2000). It can then place alternative backgrounds that have little or nothing to do with the actual location the user is situated in. Figure 2.7 shows two crude images of a user being broadcast from his office. To preserve his privacy, scene reconstruction is used to subtract out the background information from the room. The image on the left shows the office

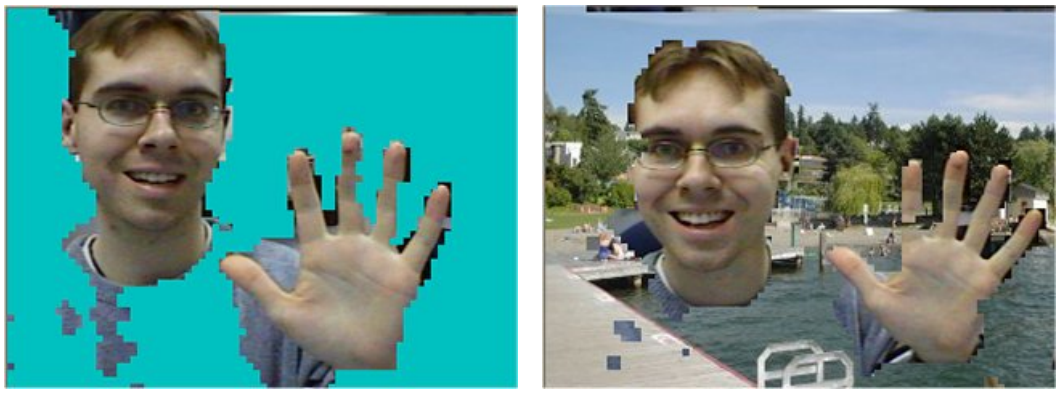

Figure 2.7: Two samples of background subtraction/reconstruction: a plain background replaces an office background (left) and a beach scene from Kirkland, WA, U.S.A. replaces an office background (right).

background replaced by a plain background, while the image on the right shows a similar office background replaced with a beach scene from Kirkland, WA, U.S.A. Boyle (2003) points out that this technique has the disadvantage of providing false information to others, which could in effect hinder any hope at providing awareness. Scene reconstruction also assumes that the information users wish to hide is the background details of their location. This may not always be the case however, e.g., none of the privacy violations in our case studies concern background, and scene reconstruction fails to mask the person's appearance and activity.

To evaluate filtration, Boyle et al. (2000) studied two distortion filters—the pixelize and blur filters—and how they balance privacy and awareness in mundane and benign office situations, e.g., people working or reading, people chatting, people eating lunch. They found that both filters offered a filtration level that adequately preserved privacy and still provided awareness for these situations. The blur filter, however, was found to balance privacy and awareness over a larger range of filtration levels than the pixelize filter. This study is discussed in more detail in Section 3.1.

Greenberg and Kuzuoka (2000) use the idea of providing video of lower fidelity to balance privacy and awareness in their Active Hydra surrogate (Figure 2.8, right). The Active Hydra surrogate contains a video and audio connection, along with a physical proximity sensor. The proximity sensor is used to mimic face-to-face situations where conversations usually arise when people are located close together (Greenberg and Kuzuoka, 2000). When both users are physically close to their Active Hydra surrogate,

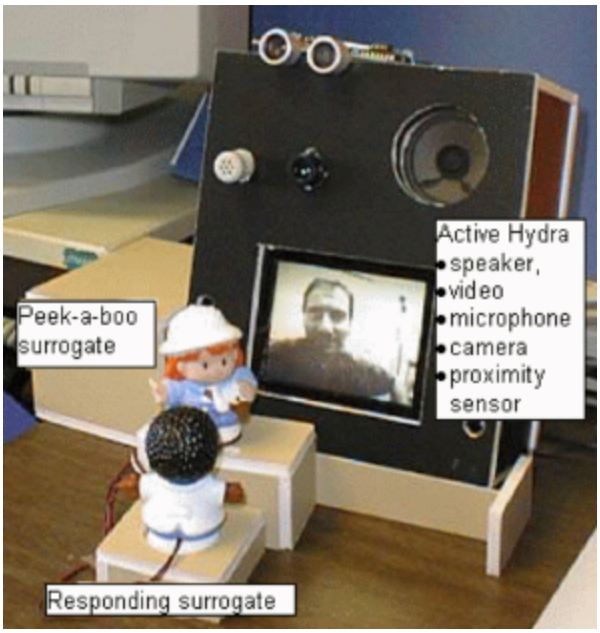

Figure 2.8: Greenberg and Kuzuoka's Active Hydra surrogate for providing informal awareness—image copied from Greenberg and Kuzuoka, 2000.

they see and hear each other in full fidelity. As users move away from the Active Hydra surrogate, audio is first disabled. In their digital video version of the system, moving further back causes the video to degrade using either a pixelize filter or a blur filter (examples shown in Figure 2.6). Finally, even further movement causes the video channel to only provide flashing still images. Their first analog video version of the system only degraded video fidelity with flashing still images. The Active Hydra surrogate provides reciprocity by making the fidelity of the audio and video transmission dependent on the proximity of both users to their Active Hydra unit.

The Nanana media space (Boyle et al., 2000) builds on Greenberg and Kuzuoka's work and also uses fidelity reduction to preserve privacy while providing awareness. Nanana provides only a video connection and uses a pixelize distortion filter, similar to the Active Hydra unit, to degrade video fidelity. In this media space, the fidelity present in the transmitted video is again dependent on how close the two end users are to their own computer and capturing device. When both users are physically close to the capturing device, as detected by a physical proximity sensor, each sees the other in full (high) fidelity, i.e., a low filtration level is applied with a high frame rate. Similarly, when both users are far away from the capturing device each sees the other in low

fidelity, i.e., a higher filtration level is applied with a low frame rate. If one user is close to the device and the other is far, a moderate filtration level is applied with a high frame rate. Nanana uses a Microsoft Windows CE™ to mirror the video captured and also provides reciprocity in the sense that both parties have the same filtration level and frame rate applied to their video.

2.4 Summary

In this chapter, I have briefly summarized how video media spaces support casual interaction and awareness, and the privacy issues that arise when using such spaces. First, I discussed the importance of *casual interaction*: the spontaneous, frequent, and unstructured encounters between co-workers throughout a typical day (Fish et al., 1992, Hudson and Smith, 1996). Casual interactions have been shown to foster knowledge, help accomplish work, and build relationships between co-workers (Kraut et al., 1988, Fish et al., 1990, 1992, 1993). Second, I showed how *informal awareness*—a naturally gained understanding of who is around and whether they are available—holds casual interaction together by allowing people to decide if and when to move into interaction (Kraut et al., 1988, Bellotti and Sellen, 1993, Gutwin et al., 1995). While informal awareness is easily gained when people are located close together, the problem is that awareness breaks down as people become separated by distance (Greenberg, 1996).

To support rich awareness over distance, utilizing the visual channel is crucial as Fish et al. (1990) found that the visual channel was used to initiate all of their observed interactions. Thus, a *video media space* is designed to provide rich awareness by supporting the visual channel through the use of an always-on video link (Mantei et al., 1991). The problem, however, is that video media spaces may broadcast sensitive information that may threaten a user's privacy. Social-psychological theories by Altman (1975) and Boyle's (2003) theory of privacy in video media spaces showed that privacy can be threatened by a video media space in three fundamental ways:

- ***An invasion of solitude:*** being distracted or interrupted
- ***A breech of confidentiality:*** losing control of what others know about you

- ***A loss of autonomy:*** losing control over how and when you participate in a media space

These privacy violations, as discussed by Bellotti (1998), may also arise from a failure to *appropriate*—create a socially acceptable appearance and/or behavior—oneself for the given situation.

A series of case studies illustrated how both telecommuters and others in a home face privacy threats from a video media space:

- ***The telecommuter lives a dual role as worker and home occupant:*** appearances and behaviors that are appropriate for the home may not be appropriate to be viewed at the office, e.g., the telecommuter works shirtless.
- ***The telecommuter is now (virtually) sitting right across from a colleague:*** unconscious acts are viewed at a very close range, e.g., the telecommuter is caught picking his nose.
- ***Threat/benefit disparity:*** individuals in the home who may gain no benefit from the video media space still incur its privacy threat, e.g, the telecommuter's lover is captured kissing her.
- ***The dual purposes typical of most home offices:*** the home office is also a spare bedroom, e.g., the telecommuter is shown naked after towelling off.
- ***Colleagues lose awareness of the telecommuter****:* colleagues are unable to accurately determine if a telecommuter is available for interaction, e.g., the telecommuter is interrupted by a co-worker when busy taking a break.

Privacy is clearly an issue when using video media spaces. To help provide a balance between privacy and awareness, it is thus necessary to empower users with privacy preservation techniques. Boyle (2003) has organized the privacy preservation techniques used in various video media spaces into four thematic approaches:

- ***Pull the Plug:*** an easy method for disabling the capturing device or software, e.g., blocking a camera's lens, turning it to face the wall, utilizing a software control.

- ***Mirror:*** providing the user with a mirror image of what is being recorded to support self-appropriation.
- ***Reciprocity:*** replicating an attribute of face-to-face situations where "if you can see someone else, they can see you." (Fish et al., 1990)
- ***Fidelity Reduction:*** reducing the video's fidelity through techniques such as distortion filters or scene reconstruction.

While these methods appear useful, they come with several fallibilities. "Pull the Plug" mechanisms typically eliminate awareness completely and must be easily reversible. Mirror facilities are quite reliable if they are visible, but screensavers or other applications may block the mirror. In addition, many mirror facilities do not show who is actually viewing the video. Reciprocity has the problem that it forces privacy to be a cooperative process, yet the privacy expectations of collaborators may not always be the same. As well, differences between the viewing and capturing angle can sometimes allow people to stand outside of the capturing region, yet still able to view the video link. Fidelity reduction, filtration in particular, offers potential, but it is not clear if filtration levels exist that balance privacy and awareness for risky home situations.

Previous research on privacy preservation techniques has focused mostly on office settings and as a result, it is not clear what privacy-protecting strategies are appropriate for home usage of a video media space. In the remaining chapters of this thesis, I will draw on the body of literature presented in this chapter to investigate privacy-protecting strategies for balancing privacy and awareness in a home media space. Chapter 3 evaluates one privacy-protecting strategy by presenting a controlled experiment aimed at determining if blur filtration is able to balance privacy and awareness for home situations involving a telecommuter. Chapter 4 presents the results of this study.

Chapter 3. A Methodology for Evaluating Blur Filtration

The preceding chapter presented several existing techniques for preserving privacy in video media spaces, yet it is unclear which techniques are able to balance privacy and awareness for home situations that may arise as a result of a home media space. Previous research (Boyle et al., 2000) has shown that distortion filters, such as the blur filter, are able to balance privacy and awareness for office situations using a video media space; however, we do not know if this is true for the riskier home situations in which we are interested.

In this chapter[2], I discuss a controlled experiment designed to evaluate blur filtration for its effectiveness in balancing privacy and awareness for typical home situations involving a telecommuter. First, I outline previous research of distortion filters for balancing privacy and awareness. Second, I discuss the experiment's methodology where I outline the research questions the study addresses, the null hypotheses, and, the independent and dependent variables. Third, I outline the materials used in the study: five video scenes typifying home telecommuting situations, the blurred video scenes and blurring algorithm, and three questionnaires (pre-test, during-test, post-test). Each of the video scenes used in the study is also analyzed for its perceived privacy risk. I conclude the chapter with an outline of the experiment's procedure. Chapter 4 gives the results of the experiment.

---

[2] A version of Chapters 3 and 4 is published as:

Neustaedter, C., Greenberg, S. and Boyle, M. (2003) **Balancing Privacy and Awareness for Telecommuters Using Blur Filtration.** Report 2003-719-22, *Department of Computer Science,* University of Calgary, January.

3.1 Background

In an effort to help mitigate privacy concerns over video links, other researchers have studied *distortion filters:* algorithmic reduction of image fidelity that hides sensitive details in a video image while still revealing awareness information (Hudson and Smith, 1996, Zhao and Stasko, 1998, Greenberg and Kuzuoka, 2000, Boyle et al., 2000). Specifically, Boyle et al. (2000) studied two distortion filters and how they balance privacy and awareness in mundane and benign office situations, e.g., people working or reading, people chatting, people eating lunch. The two filters were the *pixelize filter* that produces a mosaic of solid rectangles, and the *blur filter* that naturally blends regions of

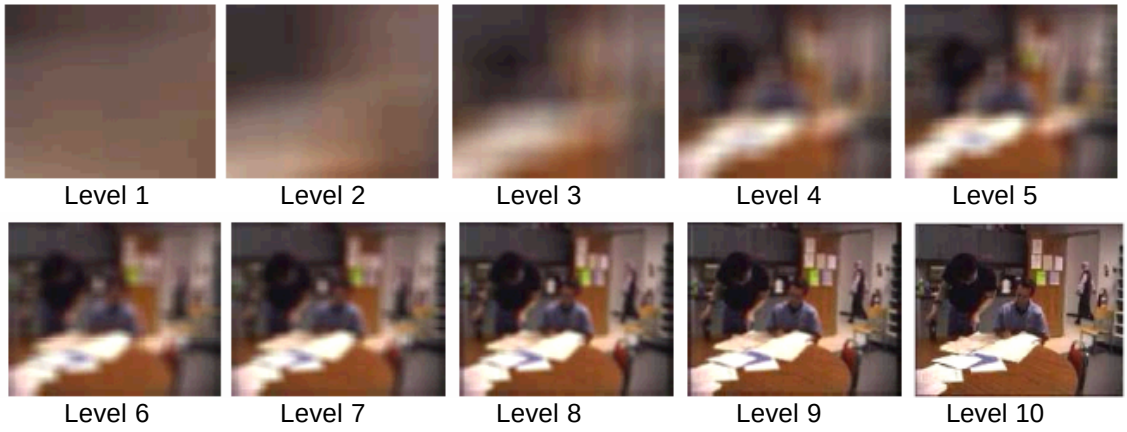

Figure 3.1: The filtration levels evaluated by Boyle et al. (2000) for the blur filter, showing one of their five scenes. Level 10 is the unfiltered scene. Copied from Boyle et al. (2000).

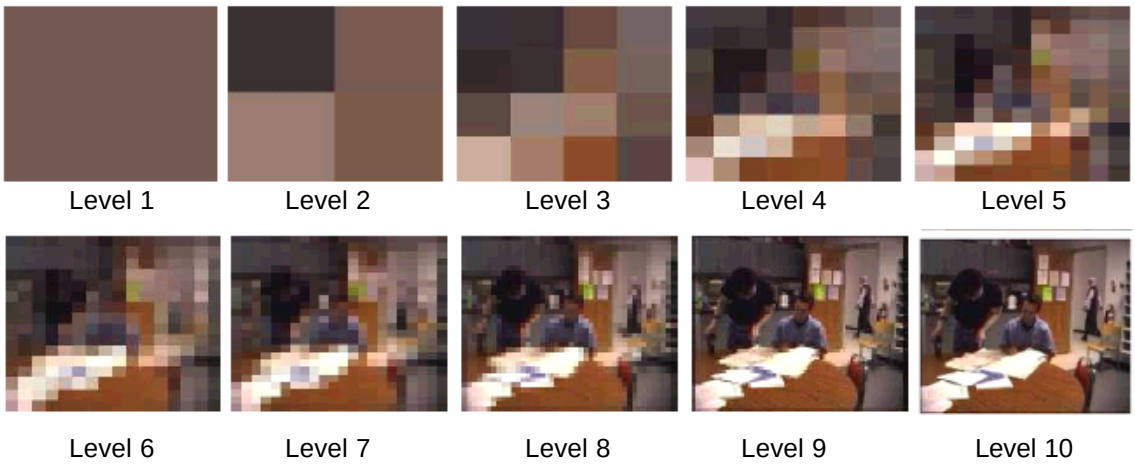

Figure 3.2: The filtration levels evaluated by Boyle et al. (2000) for the pixelize filter, showing one of their five scenes. Level 10 is the unfiltered scene. Copied from Boyle et al. (2000).

an image to produce a blurred effect. Figure 3.1 shows the levels of filtration that Boyle et al. (2000) evaluated for the blur filter, while Figure 3.2 shows the levels of filtration evaluated for the pixelize filter.

Boyle et al. (2000) wanted to know if there existed filtration levels that could adequately provide both privacy and awareness for office situations. Figure 3.3 illustrates this point by showing a privacy/awareness spectrum. The left end of the spectrum represents privacy and the right end represents awareness. Heavily filtered scenes (e.g., Boyle et al's Level 1 in Figures 3.1 and 3.2) would reside at the far left end of the spectrum; one can gain complete privacy, yet no awareness is provided for colleagues. Lightly filtered scenes (e.g., Boyle et al's Level 9 in Figures 3.1 and 3.2) would reside at the far right end of the spectrum; one can provide complete awareness, yet gain little or no privacy. Filtration levels that provide both privacy and awareness would reside somewhere in the centre of the spectrum where the privacy and awareness bars overlap; thus, these filtration levels would provide an overlap or balance between privacy and awareness. However, if no filtration levels are able to provide an adequate level of both privacy and awareness, then the two bars, in fact, would not overlap as is shown in Figure 3.3.

Boyle et al. (2000) found that both the blur filter and pixelize filter offered filtration levels that preserved privacy while still providing awareness for office situations. Specifically, filtration levels 3 to 5 for the blur filter (Figure 3.1) and levels 5 to 6 for the pixelize filter (Figure 3.2) provided an adequate level of both privacy and awareness. For these results, the overlap between the privacy and awareness bars in Figure 3.3 would

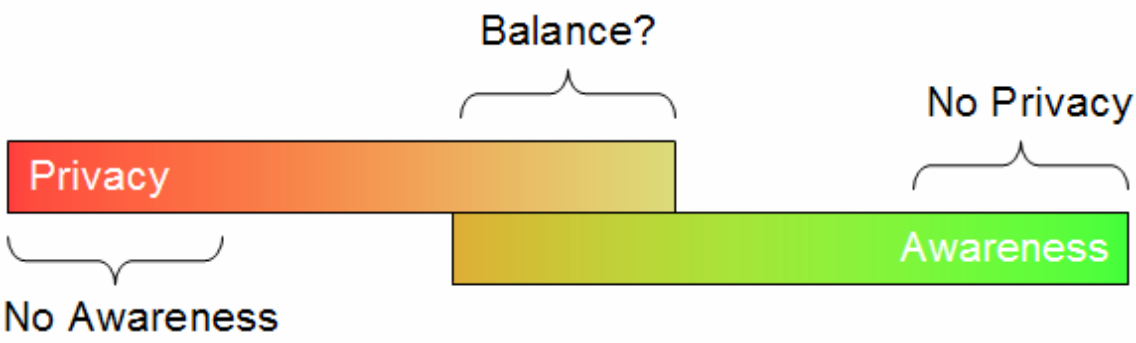

Figure 3.3: The privacy/awareness spectrum: one end represents complete privacy and the other end represents complete awareness (described more in the text).

contain blur levels 3 to 5 and pixelize levels 5 to 6.

However, Boyle et al. (2000) did not study the effects of their distortion filters on situations that may be extremely sensitive to privacy violations, such as those typified in the case studies from Chapter 2. We expect to find that equivalent or higher filtration levels are needed as risk increases. Yet it is not clear if there will still be filtration levels that can provide both privacy protection and awareness. Consequently, we set ourselves the research goal of determining how well video-blurring safeguards privacy in always-on video links that connect the home-based telecommuter with their office colleagues. To achieve this goal, we constructed an experimental study to test blur filtration with a set of scenes typifying home telecommuters that range greatly in their privacy risk. Scenes include mundane situations, such as working at a computer, to moderately risky situations such as the telecommuter kissing her partner, and to extremely threatening situations, such as being shown completely naked. The major difference between our study and Boyle et al.'s (2000) is that we are testing video usage in a home telecommuting scenario rather than an office, where we explore scenes that are much more threatening to one's privacy than everyday mundane office situations.

The next section of this chapter outlines the study's methodology and includes the specific research questions the study answers. The results of the study are outlined in Chapter 4.

3.2 Methodology

In our study, participants imagine a scenario where they are a 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. For each blur level, participants answer privacy and awareness questions, described in more detail later. Three main research questions are addressed by the study:

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 it appropriate for a colleague to see the scene and when is privacy adequately preserved?

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 a particular scene, if the blur levels which preserve privacy (Question 2) fall in the range of blur levels which 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 actually 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.

3.3 Hypotheses

Based on the previous research questions, three null hypotheses are tested in the study. The first null hypothesis analyses the viewer's ability to extract awareness cues from the video scenes at each of the ten levels of blur (Question 1). The second null hypothesis analyzes how the ten levels of blur filtration affect the perceived privacy threat within each of the video scenes (Question 2). The third null hypothesis analyses each scene's effect on the viewer's selection of blur level (Question 3).

Hypothesis 1: There is no difference in a viewer's ability to determine particular *awareness cues* from ten different levels of blur (from fully blurred to completely clear) applied to five different videos containing scenes within a home, where scenes vary in risk level ranging from no risk to high risk.

Hypothesis 2: There is no difference in the amount of *perceived privacy threat* to the telecommuter and family members presented by each of the ten levels of blur filtration (from fully blurred to completely clear) applied to five different videos containing scenes within a home, where scenes vary in risk level ranging from no risk to high risk.

Hypothesis 3: There is no difference in the viewer's *selection of blur level* as they try to make a particular scene appropriate for viewing by a distant colleague for five

different videos containing scenes within a home, where scenes vary in risk level ranging from no risk to high risk.

3.4 Independent and Dependent Variables

The independent variables for the study are scene type (5) x blur filter levels (10). The dependent variables recorded in a during test-questionnaire are: a participant's abilities to correctly identify awareness cues, a participant's confidence in identifying awareness cues, a participant's perception of the videos' level of privacy threat, the appropriateness of the scene and each corresponding blur level, and the chosen blur level for safeguarding each video.

3.5 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. Each scene shows a telecommuter performing a different activity or with a different appearance, where all scenes are recorded from the same point of view, i.e., behind the computer monitor at the same angle (Figure 3.4). The scenes are sorted by expected perceived privacy risk, from low risk to high risk (discussed later):

1. ***Working at a computer:*** The telecommuter is working at a computer while wearing clothes appropriate for both home and the office.
2. ***Picking one's nose:*** 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.
3. ***Working with no shirt on:*** The telecommuter is working shirtless at a computer.
4. ***Kissing a partner:*** The telecommuter is working when his/her partner enters the room, kisses the telecommuter intimately, and leads him/her out of the room.
5. ***Changing clothes / Naked:*** The telecommuter enters the room in a robe, is shown completely naked, and then puts on underwear.

Figure 3.4: The five (unfiltered) video scenes typifying home situations facing a telecommuter. Participants did not see the black bars for the Changing scene.

Our scene selections were based on the results of a pilot study: participants were shown a set of ten different home scenes and asked privacy questions about them (discussed in more detail in Appendix B). When asked what they would like to hide in each scene, participants most frequently reported the person's activity and appearance, while the location was reported less frequently. We hypothesize this was because when using a video link, the camera typically remains stationary and records the same background information. The background scene changes very little and soon becomes unremarkable to viewers; after all, it is simply just a room. As a result, each scene in the current study occurs in a home office/spare bedroom where both the person in the room and his/her activity determines risk.

Each scene was recorded twice—once containing a paid male as the main actor and once containing a paid female as the main actress. The actors were deliberately chosen to be middle-aged individuals with the appearance of a working professional, as opposed to university students who may be viewed as being more liberal. Recording was done with a high-quality Canon XL-1 digital video camera. No special recording lighting was used, as this would make the footage of better quality when compared to most homes, e.g., when using a PC camera at home, most individuals do not use specialized lighting required for professional video production. Each recorded sequence was approximately one minute in length. After recording, all ten video scenes were edited into approximately 10-30 second AVI video clips without degrading video quality, i.e., final videos were 720 x 480 pixels and 30 frames per second (fps). This is compared to those used in Boyle et al's (2000) study, which were Intel Indeo™ compressed at 176×144 pixels and 24 fps. While current PC cameras are not capable of capturing video at DV quality, DV format was used during our study in the anticipation that in the near future a similar quality format would be available for video conferencing.

3.6 Materials: Scenarios Provided to Participants

Privacy violations will vary depending on the subjective context of the video media space. For example, a particular person may be more willing to give out personal information than another and desire less privacy, or vice versa. To normalize this, we gave participants a telecommuting scenario, shown 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 setup 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..."*

We chose this scenario to reflect how people in our pilot study described their desired use of a home media space.

3.7 Materials: Assessing the Risk of Each Scene

The level of risk presented in each scene was assessed prior to the study using Boyle's (2003) theory of privacy in video media spaces as applied to our telecommuting scenario. The five selected scenes are similar to the five case studies (Section 2.2) presented in Chapter 2 and as a result, most present the same privacy risks. Boyle's theory states that privacy can be violated in three fundamental ways: a breech of confidentiality, invasion of solitude, loss of autonomy, or a combination of these. 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), e.g.,

create a socially acceptable behavior or appearance. We now use these potential violations to describe each scene and assess its privacy risk.

In scene 1, Linda (or Larry) is working at a computer while wearing clothes appropriate for both home and the office, e.g., clean and casual clothes (Figure 3.4). This scene is representative of a mundane activity that you would typically perform when working at home, e.g., reading the newspaper, checking email. Linda has intentionally appropriated herself correctly and therefore is not experiencing any privacy violations. This situation constitutes 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 3.4). This scene is representative of an unconscious act, e.g., scratching yourself, blowing your nose, or other grooming activities. Larry would like to pick his nose without others seeing, but the media space has threatened his autonomy because he is unable to do so; 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. We feel this is a mild breech of confidentiality. While the telecommuter has created an appropriate appearance for the office, like scene 1, his behavior is not as acceptable for others to see in detail. Because social cues of the other person's presence are lost, Larry doesn't realize his colleague can see this. Many would tolerate seeing this accidentally at a distance, yet now they are seeing a close-up, somewhat disgusting view of it.

In scene 3, the telecommuter is working at a computer wearing no shirt (Figure 3.4)—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 waiting for a shirt to dry in the laundry. The media space is threatening Linda's autonomy because she wishes to work without being fully clothed, yet this has unwanted consequences. Linda's confidentiality is being violated, 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). Linda is appropriately dressed for home, but definitely not for an office and it is very unlikely that such an appearance would even be seen at an office. This makes the scene 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 3.4). This scene is representative of typical activities between inhabitants of a home, e.g., 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. Larry's autonomy is again being threatened. His confidentiality is also being breeched because the telecommuter wants others to know little details of his personal life, yet now they know that he is about to partake in sexual activities with Linda. Linda is also experiencing the similar threats to her 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), making this scene a moderate risk as well.

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 3.4). 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. Like scene 4, Linda is showing others more than she wishes to show them and her confidentiality is breeched as a result. 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 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, e.g., 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.

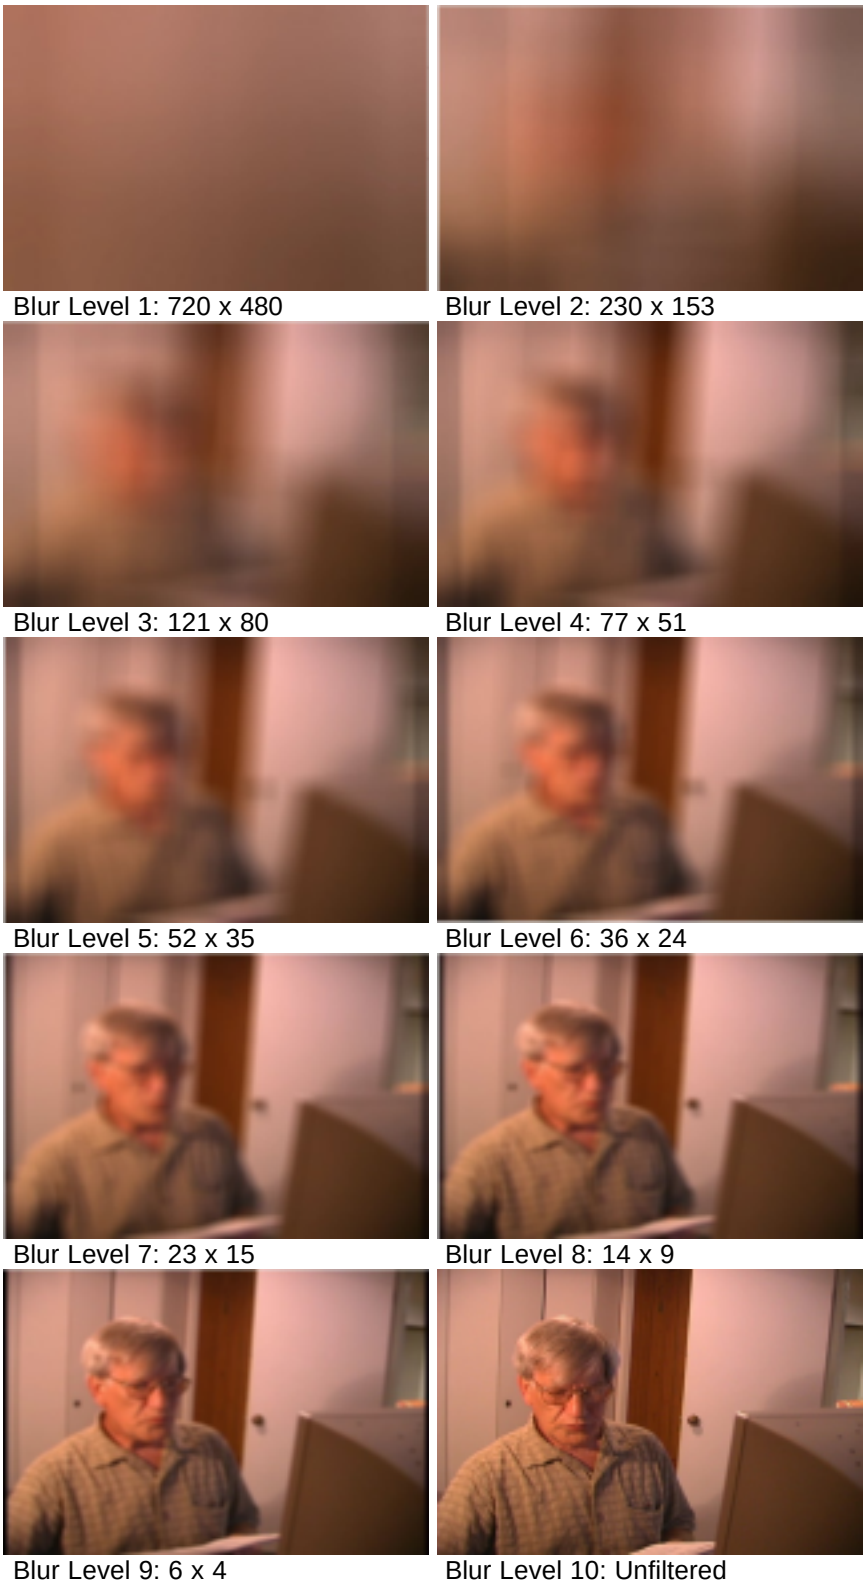

Figure 3.5: The ten blur levels evaluated in the study (currently showing the Working scene with the male actor) and the corresponding size of the pixel neighbourhood used for blurring. Reproduction quality and a lack of motion may cause these images to appear blurrier than the videos used in the study.

3.8 Materials: Blurred Video Scenes

The ten video scenes (five male, five female) were also pre-processed to create a set of videos at each of the ten different blur levels to be evaluated (Figure 3.5). A distortion algorithm from the University of Calgary's GroupLab Collabrary was used for blurring videos (Boyle and Greenberg, 2002). Depending on the level of clarity required for a video sequence, the distortion algorithm averages pixels within a neighborhood to create a video sequence, where the video smoothly changes between regions of pixels in each frame of video. Large neighborhoods create greater distortions. For example, blur level 2 in Figure 3.5 uses a neighborhood of 230 x 153 pixels for blurring. This means that for a video blurred at level 2, each pixel is averaged with the neighbouring 230 x 153 pixels. This is a typical method for smoothing (i.e., blurring) an image. We used a logarithmic function to determine the neighbourhood size for each blur level as in practice it seemed to provide the most natural progression between blur levels.

The blurring algorithm used in our study is the same algorithm used by Boyle et al (2000), yet our blur levels differ somewhat. A comparison of the proportion of pixels blurred in each of our blur levels to that of Boyle et al. (2000) shows that our blur levels are approximately 0.25 to 0.50 blur levels clearer. For example, our blur level 3 is about a quarter of a blur level clearer than Boyle et al.'s (2000) blur level 3. It is important to remember however, that the focus of the study is not to test specific blurring algorithms or neighborhood sizes, but rather to test blur filtration in general.

3.9 Materials: Questionnaires

Three questionnaires were designed for the study: pre-test questionnaire, during-test questionnaire, and post-test questionnaire.

3.9.1 Pre-Test Questionnaire

We gathered data about each participant, such as age, gender, occupation, computer experience, and telecommuting experience. The questionnaire also asked participants about their experience in using video-conferencing software and their personality (shy vs. outgoing, level of self-consciousness). The actual questionnaire is in Appendix C.3.

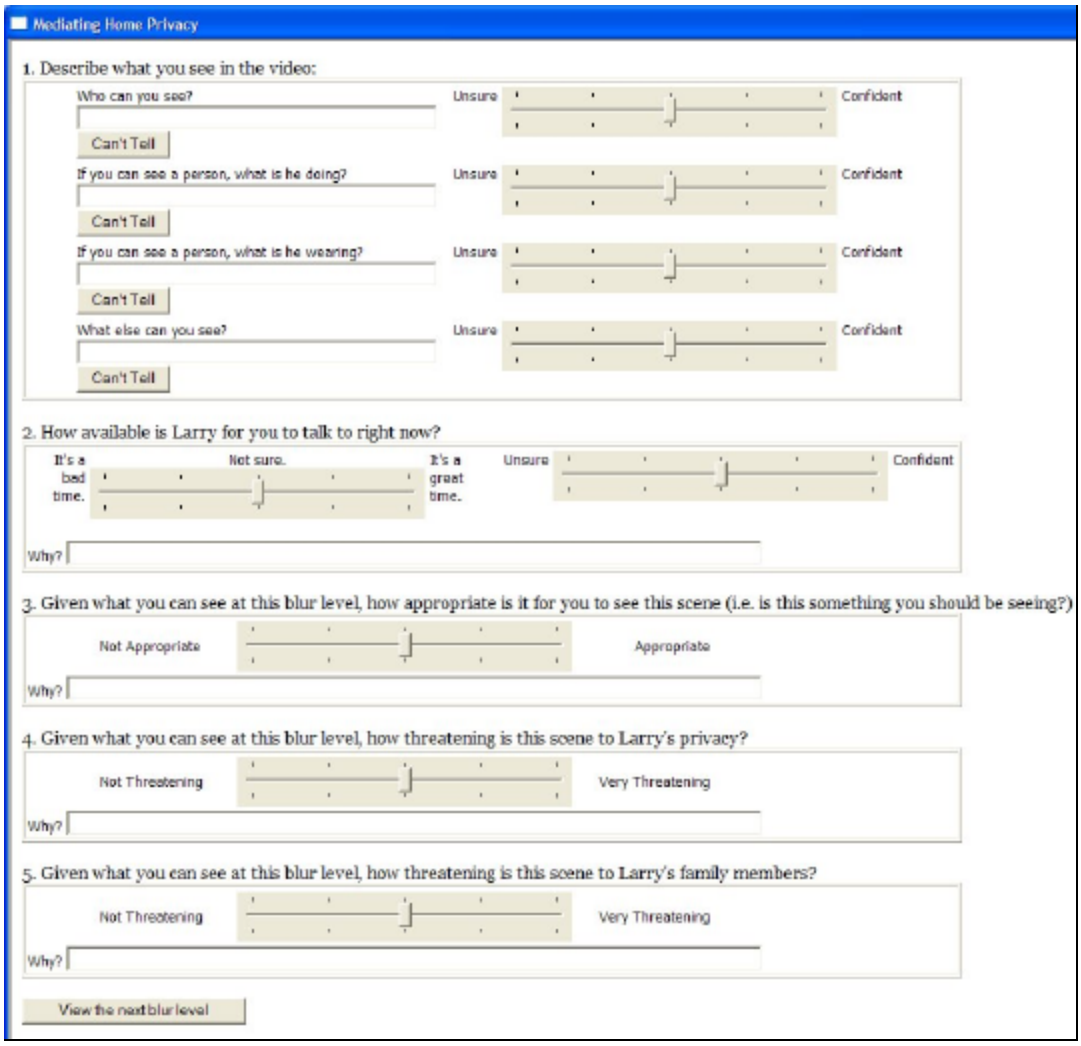

Figure 3.6: During-test questionnaire: questions answered by participants for all the blur levels for each scene.

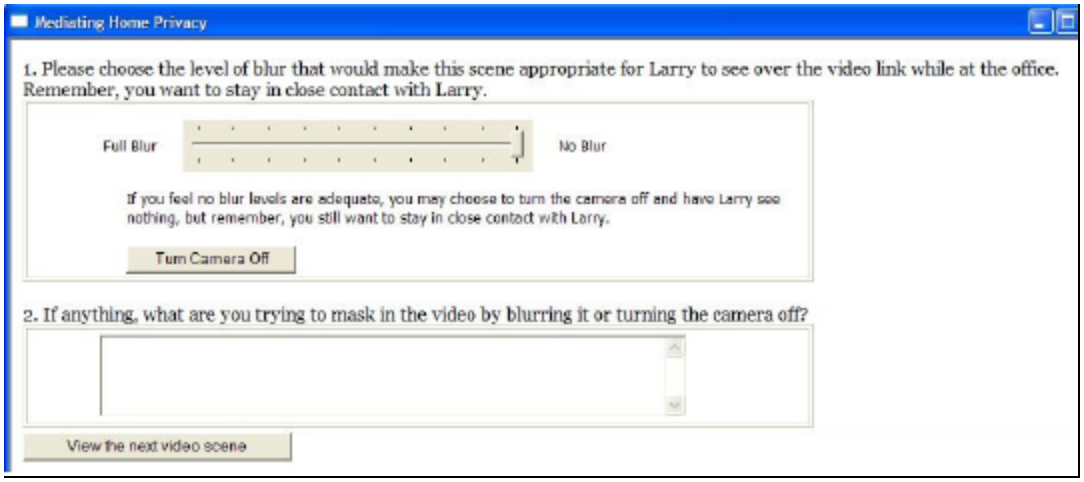

Figure 3.7: During-test questionnaire: questions answered by participants for each scene.

3.9.2 During-Test Questionnaire

We asked participants about each of the blur levels for all video scenes. The questionnaire was web-based and used two 17" CRT displays: the left display showed videos, the right display showed questions. The awareness and privacy related questions asked for each of the blur levels are shown in Figure 3.6 (Research Questions 1 and 2). Similarly, Figure 3.7 shows the set of questions used for each scene after the participant viewed all the blur levels for it. These questions ask the participant to choose a blur level that would make the scene appropriate for a colleague to view (Research Question 3).

3.9.3 Post-Test Questionnaire

In the post-test questionnaire, 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. The full questionnaire is shown in Appendix C.4. A final question asked participants to perform a forced sort of five pictures (one for each video scene, printed on standard 21.59 x 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

Figure 3.8: A sample forced sort of scenes by privacy risk showing the 300 cm "line of privacy risk" and a magnified portion of it.

sorted pictures on a "line of privacy risk" that was 300 cm long—one end represented low risk, the other end represented high risk (Figure 3.8). 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 to assess our original rating of each scene's privacy risk.

3.10 Participants

Participants were twenty people—ten females and ten males—holding a range of professional occupations, e.g., 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 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 balanced 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.

3.11 Method

The study is a *within subjects* design. Each male participant was shown all five video scenes where the telecommuter was male, while female participants looked at video scenes where the telecommuter was female; 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 3.6). 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 (e.g., blur level 1 in Figure 3.5).

2. As they viewed each blur level, they answered awareness and privacy related questions (Figure 3.6).
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 where Larry/Linda was now the viewing colleague at the office.
5. They chose a blur level for the scene and gave a reasoning (Figure 3.7). 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 from no privacy risk to high privacy risk (Figure 3.8).

The first scene shown to participants in Step 1 was always our most benign control scene containing the working telecommuter (Figure 3.4). 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 later rate the control scene 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 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 did provide us with a lower bound for awareness: we were able to know the first point at which it was possible for

participants to accurately deduce awareness cues. In Step 5, participants were 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.

3.12 Summary

In this chapter, I presented a methodology for evaluating blur filtration for its effectiveness in balancing privacy and awareness for home telecommuting situations. First, I described previous research in using distortion filters to balance privacy and awareness. In particular, Boyle et al. (2000) evaluated the blur filter and pixelize filter for their effectiveness in balancing privacy and awareness in typical office situations. Both were found to offer distortion levels where privacy and awareness were adequately provided for office situations. However, Boyle et al. (2000) did not look at the more risky situations facing telecommuters working from home. Chapter 2 showed that privacy risks are higher for home-based telecommuters than those at an office; it is not clear if distortion filtration is able to balance privacy and awareness for these situations. As a result, the experiment described in this chapter has the research goal of determining how well video-blurring safeguards privacy in always-on video links that connect the home-based telecommuter with their office colleagues. The chapter breaks this goal down into three research questions, summarized as:

1. What blur levels, if any, can provide awareness?
2. What blur levels, if any, can adequately preserve privacy?
3. What blur levels, if any, do people actually choose to use?

Each research question is then associated with a null hypothesis for the experiment. The answers to the first two questions will show whether or not blur levels exist that can provide both awareness and privacy; thus, a balance between privacy and awareness would exist. The answer to the third question will show how users actually feel about using blur filtration for home-based telecommuting situations.

For the study, five video scenes were recorded that we feel typify home telecommuting situations, e.g., each scene represents a typical situation that has been reported to us by telecommuters. The five scenes also range in their perceived privacy risk level, as rated by us using Boyle's (2003) theory of privacy in video media spaces:

1. ***Working***: low risk
2. ***Picking one's nose***: moderate risk
3. ***Working with no shirt***: moderate risk
4. ***Kissing***: moderate risk
5. ***Changing/Naked***: high risk

Each of these scenes was recorded twice, once with a male as the telecommuter and once with a female as the telecommuter, then blurred at ten levels using a standard blurring algorithm.

During the study, participants first answered a pre-test questionnaire, which asks about background information such as demographics, telecommuting experience, and personality. Participants then imagined themselves as a close-working colleague of a telecommuter. Next, participants viewed each video scene in turn at all ten levels of blur, answering privacy and awareness questions about the blurred scene, e.g., how available is the telecommuter?; how well is the telecommuter's privacy being preserved? After viewing all blur levels for a scene, participants were given the opportunity to pick a blur level or turn the camera off in order to make the scene appropriate for a colleague to view, if they were the person in the scene. Participants concluded the experiment by answering a post-test questionnaire and performing a forced sort of the scenes according to their perceived privacy risk.

The next chapter outlines the results of this controlled experiment, answering each of the three research questions in turn. The results are then used to present a set of design implications for privacy-protecting strategies for home media spaces.

Chapter 4. Can Blur Filtration Balance Privacy and Awareness?

In this chapter, I discuss the results of an experiment designed to evaluate blur filtration for its effectiveness in balancing privacy and awareness for typical home situations involving a telecommuter (the experiment's methodology is found in Chapter 3). The results presented in this chapter are divided into four main sections. First, I briefly discuss participant demographics, as captured on our pre-test questionnaire. Second, I validate our original risk assessment of each scene, where we compare it with the participants' forced sort of scenes by privacy risk. Third, I answer our three primary research questions by analyzing the study results. Finally, I determine people's willingness to actually use blur filtration within a home media space, as captured on the post-test questionnaire. I conclude the chapter with a set of design implications for privacy-protecting strategies for a home media space. These implications are based on the results of the study and articulate the difficulties in designing strategies for balancing privacy and awareness in home media spaces.

The original data analysis divided our participants into telecommuters/non-telecommuters. However, our analysis showed little difference between these two groups. For simplicity and clarity, I exclude this distinction in the following discussion and figures unless absolutely necessary.

4.1 Participant Demographics

The pre-test questionnaire confirmed that we had a broad range of participants in our study, i.e., they differed in their personalities and experience with video conferencing. Most participants described their own personality as being neither outgoing nor shy—the mean response was 3.4 (s.*d*.=0.9; median=3; 1-very shy to 5-very outgoing). Most participants were somewhat concerned about how their co-workers perceived them—the

mean response was 3.6 (*s.d.*=0.9; median=4; 1-not concerned to 5-very concerned). When asked how frequently they telecommute, participants' mean response was 2.3 (*s.d.*=1.9; median=2; 0-never telecommuted to 5-frequently telecommuted). We labelled those with a frequency of 3/5 or greater as telecommuters and confirmed their telecommuting experience by talking with each participant. When asked how frequently they used video conferencing software, participants' mean response was 1.8 (*s.d.*=1.2; median=1; 1-never to 5-frequently).

4.2 Assessing the Risk of Scenes

To discover how people perceived the privacy risk of each scene, we had them perform a forced sort of representative non-blurred pictures of each scene onto a "line of privacy," with one end indicating no risk and the other high risk (Figure 3.8).

There was reasonable consistency in participant responses: we saw only six different orderings, and even those did not differ by that much. Figure 4.1 shows these orderings: 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 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 the male actor with No Shirt *vs.* how females rated the female actress with No Shirt. For example, 8 males *vs.* 3 females selected sequences 1, 3 and 5 that place No Shirt as having the least risk (column two) of the three middle scenes. To further illustrate, 4 of the 10 females chose sequence 2, which places a female with no shirt in the 4th (second riskiest) column.

As a whole, we feel that participants' ordering of scenes confirms our original assessment of each scene's risk (Section 3.7): Working is low risk, Picking Nose, No Shirt, and Kissing are moderate risk, and Changing is high risk.

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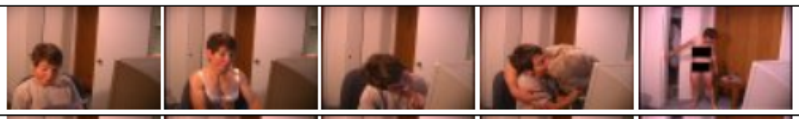	5	1	6
2 	1	4	5
3 	1	2	3
4 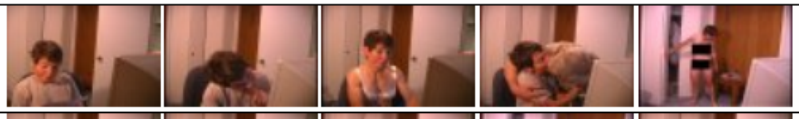	1	2	3
5 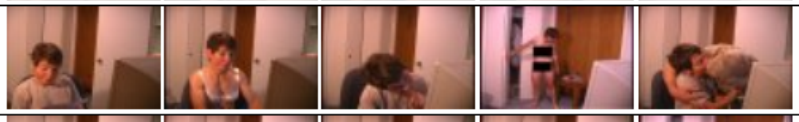	2	0	2
6 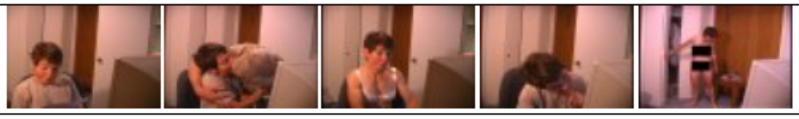	0	1	1

Figure 4.1: The frequency of each ordering of scenes found in the forced sort by males and females. Male participants used the male equivalences of these scenes.

Relative order does not indicate the strength of people's convictions of risk. To capture this, we analyzed how people position scenes on the line of privacy. Figure 4.2 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 ($p < 0.05$). While our previous results (previous page) showed some gender differences in the ordering of scenes, the strength of people's convictions of risk are quite similar between genders.

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 4.2, far left) as very low risk: it is close to the 25 cm rating, and the standard deviation is relatively small. A

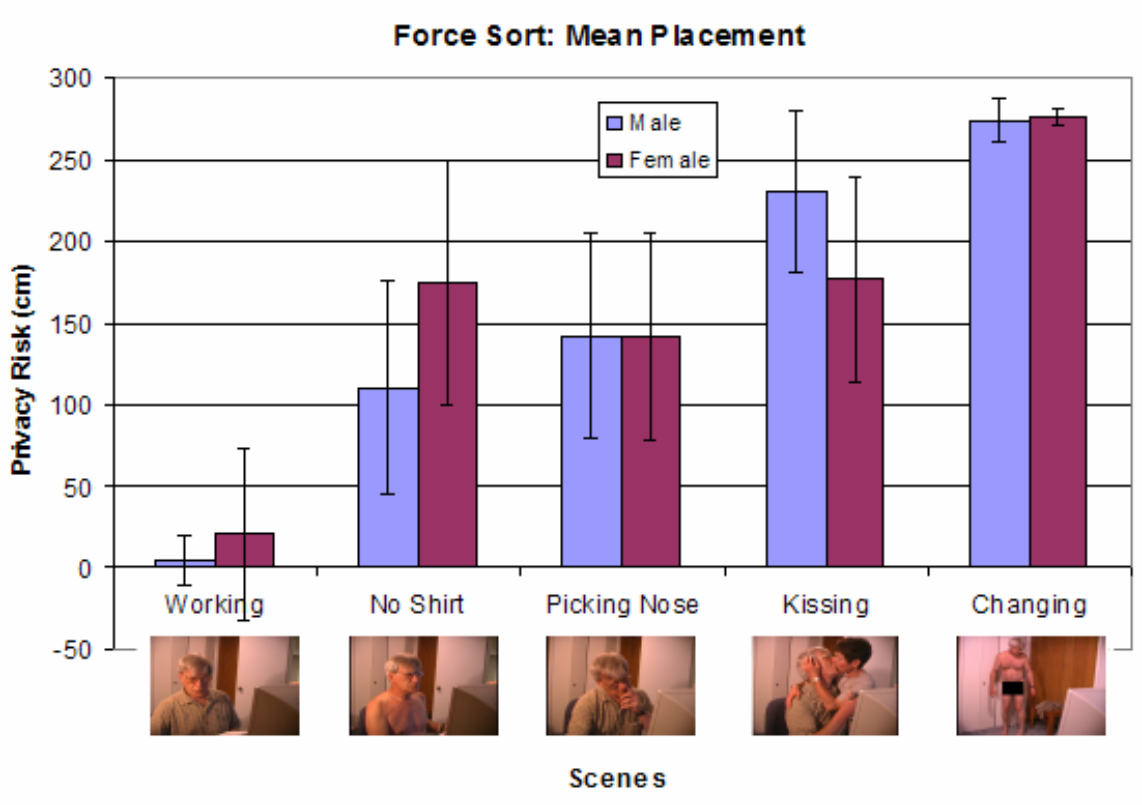

Figure 4.2: The mean placement of scenes according to risk, from low risk (0 cm) to high risk (300 cm), during the forced sort.

post-hoc analysis[3] 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 ($p < 0.01$); images were positioned around 275 cm, and its standard deviation is small.

4.3 Determining Awareness

For each level of blur, participants were asked to write what they could see in the scene, how available the person was for conversation, as well as the confidence they had in their guesses (Figure 3.6: Questions 1 and 2). We took this information and separated it into four awareness categories:

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

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 was for interaction

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. While each video's background is quite mundane, this result was also found in our pilot study where room backgrounds varied (Appendix B). This confirms our belief that background information becomes unremarkable over time. Consequently, we do not incorporate 'background' into our results. In this section, we will now show how our results answer particular awareness questions.

At what blur levels did people correctly identify awareness cues? Cues from each

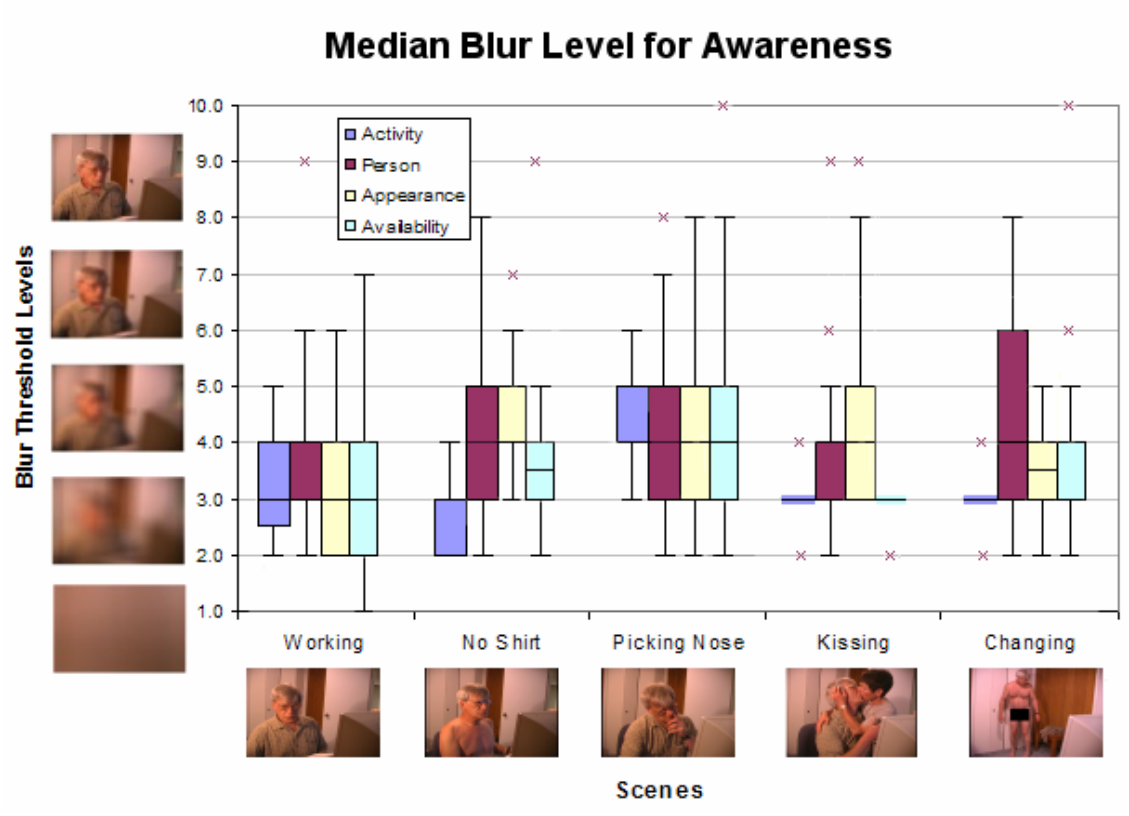

Figure 4.3: The median and range of blur levels at which participants were first able to identify awareness cues for each scene.

of the four awareness categories (listed above) were generally identifiable *between blur levels 3 and 5*. Figure 3.5 shows these blur levels, however, reproduction quality and a lack of motion may cause these images to appear blurrier than the videos used in the study. Figure 4.3 plots the median and range of blur levels at which participants were first able to correctly identify categories of awareness cues for each scene. 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) in their answer (Figure 3.6: Question 2).

	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 %

Table 4.1: Awareness Cues: the percent of participants able to identify awareness cues at specific blur levels.

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 4.1). The appearance of the actor in each scene was determined by 75% of participants between blur levels 3 and 5 (Table 4.1). Determining who was in each scene was performed between blur levels 3 and 5 by 65% of participants (Table 4.1). Availability was determined by 65% between blur levels 3 and 5 (Table 4.1). The remaining participant breakdown is summarized in Table 4.1. These numbers show that for the majority of people, their threshold for identifying awareness cues with reasonable confidence is *between blur levels 3 and 5*, although a few participants notice these cues 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 people 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 people used to first identify awareness information across the five scene types for availability ($?^2(4) = 17.62$, $p < 0.05$), activity ($?^2(4) = 40.55$, $p < 0.05$), and appearance ($?^2(4) = 25.38$, $p < 0.05$). No difference exists between scenes for determining the person ($?^2(4) = 8.592$, $p = 0.072$). A series of Wilcoxon Matched-Pairs Signed-Ranks tests show that the differences are quite varied amongst 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 the smallest amount of movement.

Did people's ability to correctly 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

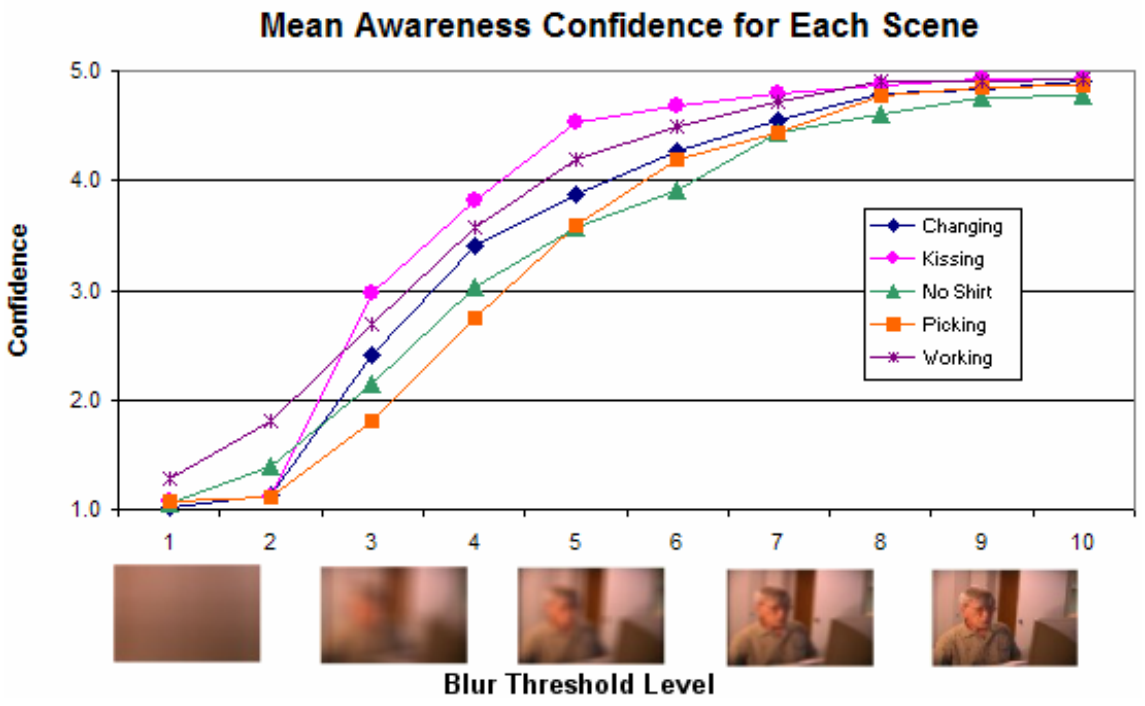

Figure 4.4: The mean level of awareness confidence found at each blur threshold level (1-low confidence to 5-high confidence).

levels found for each of the awareness categories for three of the five scenes: Changing ($?^2(3) = 28.09$, $p < 0.05$), Kissing ($?^2(3) = 18.03$, $p < 0.05$), and No Shirt ($?^2(3) = 24.22$, $p < 0.05$). The remaining two scenes, Picking Nose ($?^2(3) = 1.02$, $p = 0.80$) and Working ($?^2(3) = 3.09$, $p = 0.38$), 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 that people determined activity at slightly blurrier scenes (within one to two blur levels) before they could determine either the person or their appearance for these three scenes.

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 4.4 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) × 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 in knowing how people ranked the telecommuter's availability for each of the scenes when they were shown completely clear (Figure 3.6, Question 2). Participants were asked how available the telecommuter was for conversation at that moment. The Working scene clearly represented *being available* for most participants (mean=4.3, *s.d.*=1.3, 1-not available to 5-highly available). People believed No Shirt and Picking Nose represented *some availability* (mean=3.4, *s.d.*=1.5; mean=3.2, *s.d.*=1.7 respectively). People rated Kissing and Changing Clothes as *being unavailable*, (mean=1.2, *s.d.*=0.9, mean=1.1, *s.d.*=0.2, respectively). Participants typically used the telecommuter's activity and appearance to

determine availability. Some participants had difficulties determining availability even when the scene was shown completely clear.

In summary, all of our awareness results suggest there is a difference in a viewer's ability to determine particular awareness cues over different blur levels across different scene types. In particular, our results show that people begin perceiving all categories of awareness cues between blur levels 3 and 5, and that scene type does make a difference to this result. Furthermore, we've shown that their confidence in their guesses increases with fidelity. With these results we reject our first null hypothesis.

4.4 Perceived Privacy Threat

4.4.1 Appropriateness of the Scenes

Participants were first asked to rate the 'appropriateness' of all the scenes for each of the ten blur levels, i.e., given the telecommuting scenario, was it appropriate for the telecommuter to see the scene contents (Figure 36, Question 3)? As discussed in Chapters 2 and 3, failing to create a socially acceptable appearance and/or behavior can contribute to a privacy threat.

Does the appropriateness of each scene differ by blur level? All scenes are appropriate for *blur levels 1 and 2*, yet as video fidelity increases, only the Working scene remained largely appropriate for *all blur levels*. All other scenes decline in appropriateness to a point where participants felt they were inappropriate to see. Figure 4.5 plots the mean appropriateness values (1-not appropriate to 5-appropriate) for all participants for different blur levels. With all fully blurred scenes where almost nothing is visible, all participants gave the scenes a high appropriateness rating, i.e., approximately 4 out of 5. In general, people alter their appropriateness judgment the most between levels 3 to 5.

A two-factor ANOVA (scene type (5) x blur levels (10)) shows that there is a significant difference between the appropriateness presented between blur levels ($p < 0.05$) and between scenes ($p < 0.05$). Figure 4.5 clearly shows that differences for blur levels lays between levels 3 and 5. A post-hoc analysis shows that the only significant difference for scenes is between the Working scene and all other scenes ($p < 0.01$). That

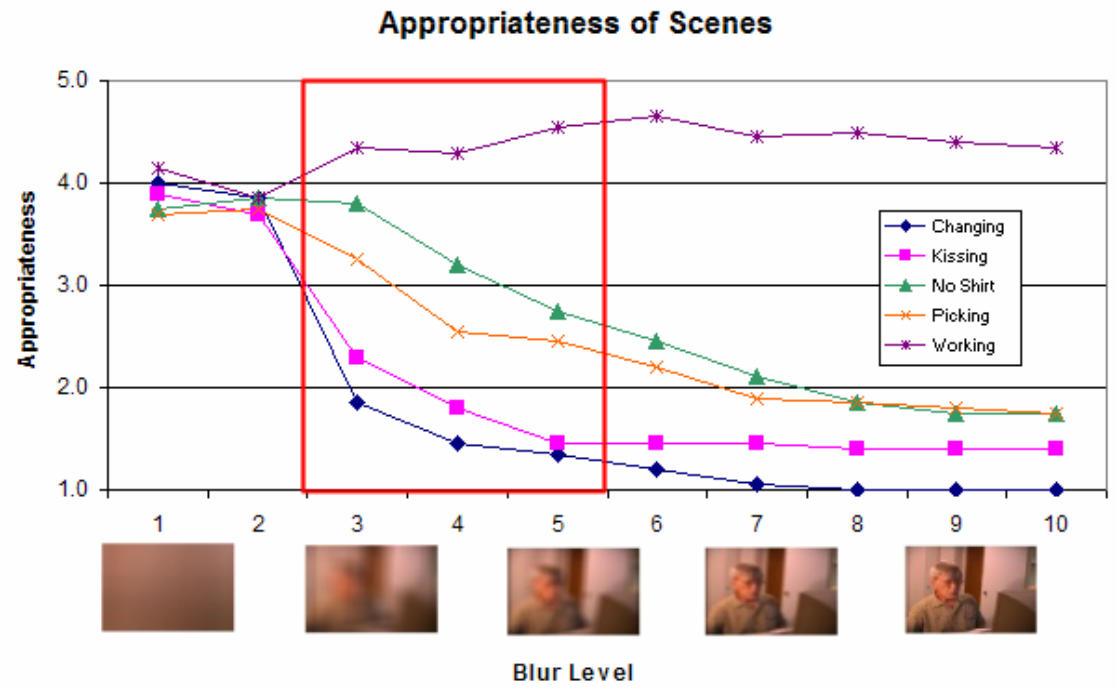

Figure 4.5: The level of appropriateness (1-not appropriate to 5-appropriate) found at each blur level. The rectangle highlights blur levels 3 to 5, shown to provide awareness.

is, there is no significant difference in the appropriateness values for each of the remaining four scenes (No Shirt to Picking Nose, $p = 0.19$, Picking Nose to Kissing, $p = 0.07$, Kissing to Changing, $p = 0.13$).

What, if anything, made scenes inappropriate? Participants mostly said it was either what the person was doing in the scene or what they were wearing (or not wearing). Scenes normally started out appropriate because participants could tell very little about what was going on. As participants discovered what was happening and grew more confident in their assumptions, appropriateness decreased. A strong inverse correlation ($r < -0.91$) exists between the amount of appropriateness and awareness found at each blur level for all but the Working scene, e.g., as awareness went up, appropriateness went down. A few participants said that not being able to tell what was happening in the scene, as was the case for the first two blur levels for most participants, made it inappropriate to view (this is why the mean level was 4 rather than 5 for fully blurred scenes). They felt that more awareness information needed to be provided at these points. Several participants also commented that they felt the scenes would be even less appropriate if they were viewing a colleague of the opposite sex.

4.4.2 Threat to the Telecommuter and Family Members

For each level of blur, participants were also asked to rate how threatening the scene was for the telecommuter and family members, given what they could currently see (Figure 3.6: Questions 4 and 5). 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 4.6. At *blur 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) × blur levels (10)) verifies that the privacy threat between different blur levels does differ significantly ($p < 0.05$). Figure 4.6 clearly 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 non-threatening),

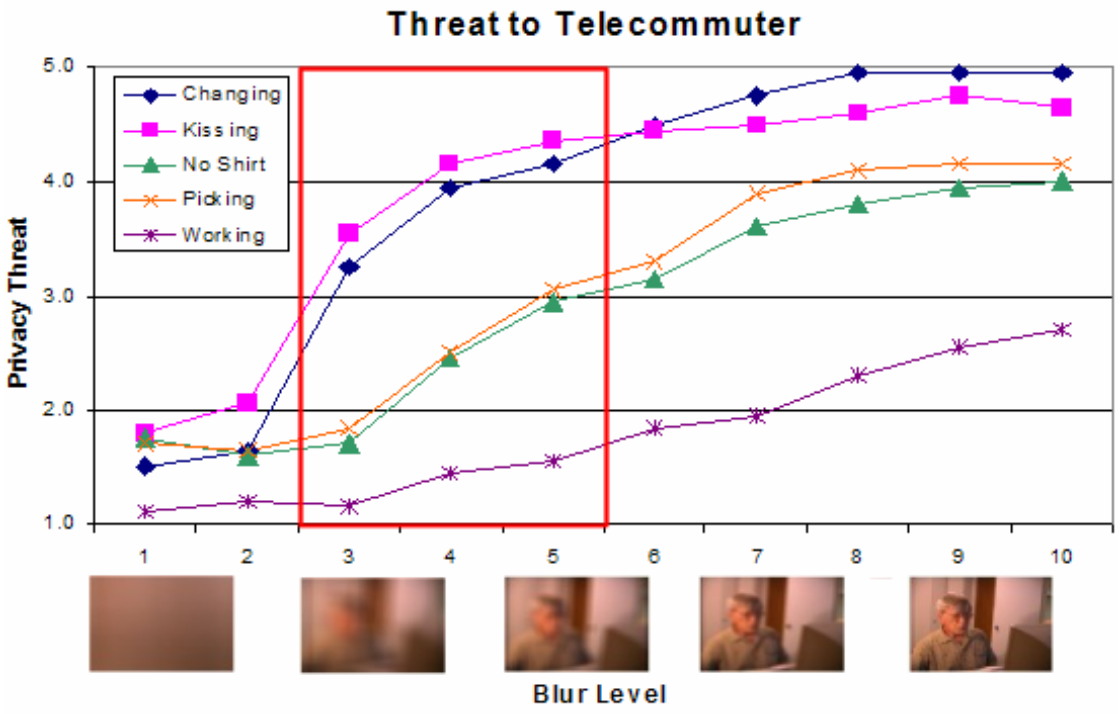

Figure 4.6: The level of privacy threat (1-low threat to 5-high threat) to the telecommuter at each blur level. The rectangle highlights blur levels 3 to 5, shown to provide awareness.

suggesting that people associated added threat with greater image fidelity, even in non-risky 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 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 (Figure 4.7). Our reasoning is 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, i.e., threat was associated with the visibility of the person's risky activity or appearance. Curiously, people 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

Figure 4.7: The level of privacy threat (1-low threat to 5-high threat) to family members at each blur level. The rectangle highlights blur levels 3 to 5, shown to provide awareness.

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, i.e., people had a tendency to infer what could happen even when instructed not to do so. 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 suggest that the perceived threat to the privacy of the telecommuter and family members differs between scenes and increases with fidelity. We can reject the second null hypothesis. In particular, our results show that only blur levels 1 and 2 made all scenes non-threatening. The results also allow us to partition the scenes into three categories of privacy risk: low (Working), moderate (Picking Nose and No Shirt), and high (Kissing and Changing).

4.5 Choosing Blur Levels

Participants were asked to imagine themselves as the telecommuter (i.e., 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 3.7, Question 1). 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 4.8 plots our results, where the *y*-axis shows the median selected blur levels chosen by participants for each scene. As one would expect, people 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 shows 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

Figure 4.8: The median and range of blur levels chosen by participants for each scene. Blur level 0 represents choosing to turn the camera off.

difference between the blur levels chosen for a particular scene by males *vs.* females.[4]

	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 %

Table 4.2: The percent of participants who chose to turn off the camera.

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 4.2 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 than male participants, and a two-factor ANOVA (gender (2) × 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 4.2 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. That is, we can reject the third null hypothesis. 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 the participants turn the camera off when high risk video is presented.

4.6 Willingness to Use Blurred/Unblurred Video

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

[4] (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, $u = 29.0$, $w = 84.0$, $z = -1.68$, $p = 0.09$).

whom they work closely (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 was 3.3 (*s.d.*=1.3). These values are significantly different ($p < 0.05$). We also checked to see if there was a significant difference between male and female responses, but none was 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. Only one person was adamantly opposed to using video.

When asked if given the opportunity, would they actually use blurred video in an *office*, 65% of participants said they would (Table 4.3). Those who would use blurred video from an office gave the following reasons: it would be helpful for people separated across floors or buildings; it could provide availability information while still providing privacy; they could turn the camera off if needed; and they didn't expect to do inappropriate things in the office setting. 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.

	Male (n=10)	Female (n=10)	All (n=20)
Office - yes	6	7	13
Office - no	4	3	7
Home - yes	5	4	9
Home - no	5	6	11

Table 4.3: Number of participants that are willing/not willing to use an always-on video link at home or an office.

Participants were then asked if given the opportunity, would they actually use blurred video from *home*, 45% of participants said they would (Table 4.3). Most of those who said they would use blurred video at home had several restrictions: they wanted to choose the room where the camera was located, they wanted a mirror facility to know what was being captured, and they wanted control over the blur level and whether or not the camera was on. Several also commented that they would simply leave the room to do private things that they would not want the colleague to see. One participant insisted that the viewing colleague would need to know how to use the video link properly and also respect the privacy of the telecommuter. 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 blur levels did not balance privacy and awareness. They also said that they saw the home as a place where they would 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.

4.7 Discussion

This chapter 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 first two 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?

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 3.5, it is important to remember that participants saw full motion videos where motion aids in identification. These values do not take into consideration the confidence of participants; rather they reflect when participants first correctly perceive

what is happening in each scene. As mentioned, at this point most participants were not confident in their responses. The levels we found for providing awareness are somewhat more filtered (2 to 3 levels) than those found by Boyle et al. (2000). We believe this difference is a result of using videos of a greater fidelity than Boyle et al. (2000).

Question 2: At what blur levels is it appropriate for a colleague to see the scene and when is privacy adequately preserved?

Only at *blur levels 1 and 2* is it appropriate for a colleague to view all the scenes. *Blur levels 1 and 2* are also 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. For this scene privacy threat is low for *blur levels 1 to 7* and the awareness levels remain as *blur levels 3 to 5*. Thus, we can see that blur filtration is only able to balance privacy and awareness for mundane home situations. However, designs should allow users to override these default values, perhaps through user preferences or direct controls. This is discussed in more detail in Chapter 6.

This confirms the results found in Boyle et al's (2000) study; for benign situations like working, there are blur levels that are able to balance privacy and awareness. More importantly, this shows that blur filtration is not able to balance privacy and awareness for the high risk home situations in which we are interested. The significance of this is that blur filtration by itself does not suffice for privacy protection in video-based telecommuting situations; other privacy-protecting strategies are required. People simply do not trust techniques where a camera continuously faces them. It matters little how the image is being filtered, whether the camera is capturing or not, or even if the camera is turned on! By implication, this means that other image processing techniques will do no better than blur filtration.

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 choose varied for the three different risk categories that we found: low threat, moderate threat, and high threat. Participants choose more distorted blur levels or more participants choose to turn the camera off altogether as the risk to privacy posed by a scene increases. While this is expected, the major importance lays in the fact that people begin to abandon the filtration technique with increased risk; thus, emphasizing the point that other privacy-protecting strategies are necessary for balancing privacy and awareness in home situations. The next section outlines the implications from these results for the design of privacy-protecting strategies.

4.8 Design Implications

By taking a step back from the study results, we can see that several issues arise for developing privacy-protecting strategies to balance privacy and awareness in a home media space when scenes have the potential to be more risky. 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 offer potential solutions to the problem of balancing privacy with awareness.

4.8.1 It's Not for Everybody

The first issue is that clearly a home media space is not something 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. It is expected that in situations less than ideal, the

privacy threat will increase, while awareness will remain mostly the same. Contrarily, situations involving just family and friends may require less privacy and more awareness, but we do not know this for sure.

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 also those for whom it will not. Undoubtedly, the person 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, Harper, 1996). Here, a badge worn by workers was used to track each person's location throughout the workplace. Incoming telephone calls could then be routed to the phone nearest to the recipient (Want et al., 1992).

For individuals for whom a home media space is suitable, 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 a participant of a particular space declines, then this media space should not be set up or an even greater privacy threat will arise.

4.8.2 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 the 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. For these reasons, context-aware systems that can detect and control privacy levels for users of a home media space may be desirable. 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.

4.8.3 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. One can try to balance privacy for both individuals at the same time, but this is largely difficult and perhaps 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. The telecommuter wants to use the home media space, but certainly the family member does not. The scene formerly presented little privacy threat, but now is quite threatening and very inappropriate! 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 level of interaction with his colleague. Balancing privacy and awareness in these situations is difficult and individuals may need to compromise their desired privacy for the desired privacy of another.

4.8.4 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 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.

4.9 Summary

We began this chapter 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. The results of our study confirm the results found by Boyle et al. (2000); blur filtration can balance privacy and awareness for mundane situations like working. More importantly, we found that privacy and awareness are not balanced by any blur levels for the risky home situations that we are interested in. These results are significant because we have shown that for home-based video links blur filtration by itself does not suffice for privacy protection; other privacy-protecting strategies and technologies are required.

In retrospect, our results may not seem surprising. Yet countless researchers (e.g., Zhao and Stasko, 1998, Crowley et al. 2000) are pursuing avenues where filtration 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. By implication, this means other image processing techniques will not suffice for balancing privacy and awareness in home-based telecommuting situations. This was not found by Boyle et al. (2000) 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.

The results of the study have important design implications for privacy-protecting strategies. First and foremost, a home media space is not suitable for everybody; media space participants must be willing, with a real desire to be a part of the space. For those who choose to participate, a high degree of control over privacy, along with feedback of the degree of privacy being maintained is strongly desired. When designing privacy-protecting strategies, it is also important to consider that in most homes multiple people exist and privacy expectations may vary between them. Home media space designs must address the privacy concerns of both the telecommuter and others in the home. A final consideration is the effect of combining two separate cultures: an office culture and a home culture. Both cultures have varying privacy expectations and a successful design must leverage this.

The next chapter takes a step back in order to identify other privacy-protecting strategies, which may be appropriate for balancing privacy and awareness in a home media space. First, I look at previous social-psychological research on privacy mechanisms used by various cultures for privacy regulation. Next, I use this research to construct a framework of privacy-protecting strategies for the design of a home media space.

Chapter 5. A Framework for the Design of a Home Media Space

In this chapter, I discuss a framework for the design of a home media space. This framework involves the identification of privacy-protecting strategies for balancing privacy and awareness in a home media space. First, using social psychology theory, I define culture and outline four categories of *privacy mechanisms*, which are used by various cultures to regulate privacy within their society: verbal behavior, non-verbal behavior, environmental mechanisms, and cultural mechanisms. Second, I discuss the creation of a *home media space culture*, which is the foundation for presenting privacy mechanisms to home media space users. Third, for each category of privacy-mechanisms, I discuss privacy-protecting strategies for balancing privacy and awareness in a home media space, their feasibility, and their usefulness. These take into consideration the design implications discussed in Chapter 4 and include strategies for providing both privacy control and feedback. The framework presented in this chapter is used in the design of a home media space, which is presented in Chapter 6.

5.1 Defining Culture

Cultures are the foundation for regulating privacy as each culture develops and employs its own privacy-protecting strategies. According to Altman and Chemers (1980), over 150 different definitions of culture exist, depending on the field of study and usage. For the purpose of my research, I have chosen to use the definition presented by Altman and Chemers, which is common to most social psychology research. Thus, for a culture to exist, "people must agree, with or without verbalizing their agreement, that there are common ways to view the world and to behave." (Altman and Chemers, 1980) Participants in a culture need not agree on everything, what is more important is that a common consensus is present on key issues relating to the culture. The creation of a

culture is an evolutionary process whereby a culture is slowly moulded and shaped over time. With this evolution comes the creation and evolution of a set of privacy mechanisms that can be used by the culture to regulate their privacy (Altman and Chemers, 1980). Cultural practices, such as privacy mechanisms, are passed on by educating others like offspring or new members (Altman and Chemers, 1980). Using this definition of culture, we will now look at privacy mechanisms used by cultures to regulate privacy and interaction.

5.2 Privacy Mechanisms of Humans

Privacy mechanisms are the behaviors and actions that humans employ to control privacy. Each and every culture that has lived has used privacy mechanisms to regulate interaction with others (Altman and Chemers, 1980). When individuals require more privacy, they use these mechanisms to let others know they desire less interaction. Just the same, when individuals require more interaction, they use these mechanisms to let others know they desire less privacy. The privacy mechanisms that are used by humans can be classified into four categories (Altman, 1975):

1. ***Verbal behaviors****:* the use of the content and structure of what is being said;
2. ***Non-verbal behaviors****:* the use of body language;
3. ***Environmental mechanisms****:* the use of physical artefacts and features of an environment; and,
4. ***Cultural mechanisms****:* the use of cultural practices and social customs.

These four categories form a coherent system where mechanisms from one category can be substituted for mechanisms from another category depending on the culture and given situation. In fact, research has shown that different cultures employ different privacy mechanisms (Altman and Chemers, 1980). For example, North Americans typically live in houses or apartments where doors and walls are prevalent. As a result, North Americans often rely on environmental mechanisms presented in the architectures of their homes to regulate privacy. If a person seeks privacy to read a book, she can simply go to another room and if necessary, shut a door. In contrast, other

cultures, such as the Iban of Borneo, rely very little on environmental mechanisms (Patterson and Chiswick, 1981). The Iban live in communal longhouses where 22 different families normally live. Families share portions of their living space and social contact is prevalent. The physical environment presents little for privacy mechanisms, so instead, the Iban rely on social practices. One such instance occurs in the evenings, when the shared portions of the longhouses are accepted as being private regions and intrusion is limited (Patterson and Chiswick, 1981).

In the next four sections, I will outline each category of privacy mechanisms in turn. For simplicity, I will use the terms 'privacy mechanisms' and 'privacy-protecting strategies' interchangeably for the remainder of this chapter.

5.2.1 Verbal Behaviors

Verbal behavior consists of the use of content and structure of what is said to control privacy (Altman, 1975). *Content* refers to the words that are spoken and their meanings, while *structure* refers to linguistic features of what is spoken, such as tone, pitch, dynamics, and pronunciation (Altman, 1975). The content of speech normally provides explicit instructions as to how much privacy a person desires. For example, if a family member enters the room where a telecommuter is currently working, the telecommuter may say, "I'd like to be left alone" if he would like to have more privacy or alternatively, "please come in" if he desires interaction. Here, the actual content of what the telecommuter says creates an instruction for the family member. The structure of speech can be much more difficult to comprehend, however, because it has the power to either emphasize the content, or alter it. For example, a pleasant and welcoming tone could be used by the telecommuter when he says, "please come in." The content of what is said gives the family member the impression that the telecommuter truly does want interaction and his solitude is not being intruded upon. On the other hand, the telecommuter may use an impatient and annoyed tone when inviting the family member into the room, suggesting to the family member that in fact the telecommuter desires more privacy, but is compromising this desire. Here, the structure of the telecommuter's speech alters the meaning of the content.

Verbal behavior does have several problems, which can limit its effectiveness in regulating privacy. First, it can be quite easy for messages to get misinterpreted if both the speaker and the listener do not understand a common language, or moreover, the listener does not understand the content or structure of what is being said. Second, the environment can play a role in creating ambiguous verbal behaviors. Walls, doors, or even loud music may cause a telecommuter to alter her speech dynamics and shout a response to a family member. Just the same, individuals in a library may be required to whisper or speak softly to one another. These changes in structure can unintentionally present a mixed or incorrect message. Third, verbal behavior can also be easily affected by culture. What people choose to say, or choose not to say will depend on the speaker and her relationship with the listener. For reasons such as these, other privacy mechanisms may be substituted for verbal behaviors, or combined with them.

5.2.2 Non-verbal Behaviors

Non-verbal behavior consists of the use of body language, such as gestures and posture, to control privacy (Altman, 1975). A person's posture as he relaxes in a chair may indicate that more privacy is desired during a nap, or a gesture, such as a wave, from across the room may signify a desire for conversation. Research has shown that people have a natural tendency to understand most non-verbal behaviors and their privacy expectations (Altman, 1975). When people are located close together, non-verbal behaviors increase (Altman, 1975). For example, in an exam situation, people may try to block or cover their test paper, indicating their desire for privacy. Non-verbal behavior is often tightly coupled with verbal behavior and has the power to again, either emphasize or alter the content of verbal behavior. A person may show her enthusiasm in a conversation by using large sweeping hand gestures as she speaks, or conversely could show dissatisfaction by crossing her arms while speaking and listening.

Some problems exist with non-verbal behaviors. First, as mentioned, non-verbal behaviors increase when people are located close together; thus, a problem for non-verbal behaviors occurs then when people are separated by distance. This can simply mean people are in different rooms, or even worse, different buildings. If technology is not used for communication in such situations, it may be impossible to use non-verbal

behaviors. For example, when using a telephone one can not see the gestures or posture of the people he talks with and must rely solely on verbal behaviors. Second, although people have a natural tendency for understanding non-verbal behaviors, unlike most spoken languages, there does not exist a common language of gestures and postures. Non-verbal behaviors can change between individuals and there always exists a threat of misunderstanding a non-verbal behavior. To alleviate these problems, non-verbal behaviors can be combined with other privacy mechanisms, or substituted.

5.2.3 Environmental Mechanisms

Environmental mechanisms consist of the use of physical artefacts and features of an environment to control privacy (Altman, 1975). The first thing people usually think of when discussing environmental mechanisms are physical objects, such as walls, fences, doors, windows, blinds, and curtains (Altman, 1975). Such objects are used quite frequently in North American culture to control access to one's home. For example, to limit neighbours from viewing one's backyard, a fence may be built or a large row of trees could be planted. Research has shown that environmental mechanisms are more prevalent in cultures containing a wide variety of homes, such as North American culture (Altman and Chemers, 1980).

Other, less obvious, environmental artefacts are also used to control privacy, such as the clothing one wears (Altman, 1975). People can signal their approachability with the type of clothing they wear and research shows that dressing according to common standards conveys an acceptance to the current situation (Altman, 1975). Personal space is also heavily used to control one's privacy and interaction. Spatial zones around a person are used for different types of interaction and different social relationships (Altman, 1975):

- ***public zone***: a distance greater than 12 feet is used for strangers, and those one wishes to have little or no interaction with
- ***social distance:*** a distance of 4 to 12 feet is used for friendly interaction, often with acquaintances

- ***personal distance:*** a distance of 1.5 to 4 feet is used for interacting and conversing with friends
- ***intimate distance:*** a distance of less than 18 inches is used for intimate activities with close friends or family members

Timing is another environmental mechanism commonly used by people to control privacy (Altman, 1975). People may choose to do certain activities when others are not around, thus gaining more privacy. For example, a telecommuter may desire more privacy when reading and may choose to read when family members have left the house.

Environmental mechanisms do have their limitations. As was discussed previously, cultures such as the Iban of Borneo cannot rely on environmental mechanisms like walls and doors because of their living situation (Patterson and Chiswick, 1981). The Iban typically live in communal longhouses that contain approximately 22 different families. The physical environment presents little means for utilizing privacy. Some work cultures may also face similar situations, e.g., groups of cubicles, shared offices, and shared labs. In these cases, cultures may need to rely on other privacy mechanisms when environmental ones are not available.

5.2.4 Cultural Mechanisms

Cultural mechanisms consist of the use of cultural practices and social customs to control privacy (Altman, 1975). Although it may often go unnoticed, each culture contains a set of learned social practices and customs that have evolved and developed over time (Altman, 1975). These practices can be used to gain more privacy when desired and those who do not adhere to such customs are often labelled as outcasts or social deviants. Cultural practices evolve over time and are passed on by educating others, such as offspring or new members (Altman and Chemers, 1980).

Cultural mechanisms for controlling privacy are more prevalent in areas of communal living, such as the longhouses of the Iban of Borneo. In such situations, environmental mechanisms are less available and therefore inhabitants must rely on social protocol for privacy regulation (Patterson and Chiswick, 1981). For example, while the shared space of the longhouses provides little privacy during the day, in the

evenings the Iban rely on social protocols to gain more privacy for their families. Areas that are public within the longhouse during the day become private areas in the evening and most interaction occurs outside the longhouse on a gallery (Patterson and Chiswick, 1981).

In North American culture, the sanctity of closed doors is prevalent as a cultural mechanism for controlling privacy (Altman, 1975). It is common for someone seeking privacy to close a door to a room, such as a person desiring privacy while using a bathroom. While closing a door may seem like an environmental means for controlling privacy, it is the social practice engaged by those seeing the door closed that outlines the cultural mechanism at play; those who see the door closed will normally knock first before entering. Formal status within North American culture also plays a large role in regulating privacy with social customs and practices (Altman, 1975). For example, high ranked officials often have more doors to go through to get to their offices (Altman, 1975).

5.3 A Home Media Space Culture

Cultures use a repertoire of privacy mechanisms to regulate their privacy and these mechanisms vary depending on the culture. We will now turn our attention to those individuals using a home media space. In a home media space, we are dealing with two distinct cultures: an office culture, containing colleagues of a telecommuter; and, a home culture containing a telecommuter and family members. Each of these cultures has its own privacy expectations and as a result, when we link them in a home media space privacy threats arise. In this section, I discuss how these two cultures can create a new culture for users of a home media space. A *home media space culture*—the resulting culture when a home culture is linked to an office culture—is important because through it, just like other cultures, privacy mechanisms are made available for media space users. I define media space participants or users to be telecommuters, others within the home who may be subject to the media space, and colleagues at the office using the space. I include others in the home for the simple reason that they too need to be able to control their privacy if a media space is to be present in their environment. The next section

discusses the creation of a home media space culture and is the foundation for developing privacy-protecting strategies for a home media space.

5.3.1 Defining a Home Media Space Culture

When telecommuters and colleagues at the office use a home media space they are taking part in the evolutionary process of defining a home media space culture (Dourish, 1993). This culture draws on customs from both a home culture, as well as an office culture to create a culture similar to both, yet unique to home media spaces. A home media space culture is important because it helps set the foundation for accepted privacy mechanisms. As a culture, media space participants will rely on specific privacy mechanisms presented to them in the design of the media space. While people have been linking homes and offices for years with technology, such as telephones, instant messaging, and email, none of these technologies has the ability to connect two cultures more than video because of the rich level of awareness presented. As Altman and Chemers (1980) discuss, cultures rely on the existence and acceptance of core beliefs; thus, participants in a home media space culture need to share a common consensus for key issues involving the use of the video media space.

The creation of a home media space culture faces several interesting challenges. As mentioned, the creation of a culture is an evolutionary process and for this reason key issues may not be universally accepted in a home media space at the onset. Privacy threats will be heightened until a point when the culture has become well defined and a set of core beliefs are present. For example, it may be the case that a telecommuter may feel it is fine to work at home while not wearing a shirt, but would prefer if her colleague at the office did not watch her work without a shirt. If a common understanding that it is inappropriate to view the telecommuter working shirtless has yet to be established, then colleagues at the office risk threatening the privacy of the telecommuter.

A further impediment to establishing a common consensus is the openness of participants in the culture. Participants must be willing to accept and understand the desires of others within the space or a consensus will never develop and the culture will fracture. This includes all media space participants: the telecommuter, others present in the home and subject to the media space, and colleagues at the office using the video link.

It is not always the case that each of these parties will be accepting of the culture. However, it may be the case that as the culture evolves, their opinions will also evolve.

The next section outlines privacy mechanisms for a home media space culture and separates them into the four categories of mechanisms, discussed previously in this chapter: verbal behavior, non-verbal behavior, environmental mechanisms, and cultural mechanisms.

5.4 Privacy Mechanisms for a Home Media Space

For each category of privacy mechanisms previously discussed in this chapter, I will now outline the feasibility of using mechanisms from that particular category within a home media space. I also present design techniques for providing feedback and control of privacy for each category. *Feedback* is an essential component of privacy mechanisms because it allows media space participants to know whether or not they are attaining their desired level of privacy. Bellotti outlines that in a media space, feedback involves "informing people when and what information about them is being captured and to whom the information is being made available." (Bellotti, 1998) Feedback must meet three criteria for it to be sufficient for media space participants, according to Bellotti (1998):

1. ***timely:*** there should be little or no delay in presenting the feedback so participants have the opportunity to make timely modifications to what is being captured;
2. ***appropriate:*** feedback must be suitable for the given situation in order for the participant to react accordingly to it; and,
3. ***distinctive:*** feedback must be understandable and unambiguous so participants can react accordingly to it.

Once media space participants know how much privacy is being attained, they need the ability to adjust their current level of privacy being maintained. This comes in the form of privacy *control* and as Bellotti (1998) points out, control involves "empowering people to stipulate what information they project and who can get a hold of it." Control must be presented using a lightweight and understandable technique so it is easy for participants to easily alter their attained privacy level. Moreover, this control is tightly

coupled with privacy feedback. If feedback is not sufficient then measures for controlling privacy cannot be taken, leaving media space participants at risk of having their privacy threatened.

5.4.1 Verbal Behavior: Sound and Voice

Verbal behavior consists of the use of content and structure of what is said to control privacy (Altman, 1975). Using this definition and extending it to include technology that may be present in a home media space, verbal behavior in a home media space can consist of:

1. verbal instructions between media space participants;
2. verbal instructions to devices within the media space; and,
3. verbal instructions or sound cues from devices in the media space to media space participants.

Each of these techniques will now be discussed in more detail.

The first approach, using verbal instructions between media space participants, could mean instructions between either co-located participants or distance separated participants. The simplest approach for using verbal behavior to regulate privacy would be the first choice: between co-located home media space participants, such as individuals with the home. For example, if a telecommuter was working in her home office and using the video link when a family member entered the room, the telecommuter could warn a family member, "The camera is on, please wait until I turn it off." Such verbal behavior could be used to control the privacy of the family member with little difficulty.

For distance-separated participants, communication must be done through the media space. Two choices exist to support this: the media space could use an open audio link with continuous audio being captured; or, an optional audio link could be used where users could, say, press a button to transmit audio. Both choices allow participants to exchange dialog, yet the first introduces a new privacy threat from always capturing and broadcasting audio. One technique proposed by Hudson and Smith (1996) for preserving

privacy for an open audio link alters the audio to a point where it is difficult to understand what is being said. This would compromise the effectiveness of using the verbal behavior to regulate privacy however. The second choice offers a more heavyweight solution because users must explicitly turn on and off the audio link, yet it implicitly preserves privacy because users now choose when they are heard by others. As a result, an optional audio link presents a more favourable option. Greenberg and Kuzuoka (2000)'s Active Hydra uses proximity to determine whether or not an audio link is open; when both users are close to their unit, the audio link is enabled.

When designing such audio links it is important to remember the effect of verbal structure on people's interaction with others. For example, Huang et al. (2002) shows that the volume of one's voice in video-mediated communication can affect his role in group decision making tasks. Artificially loud voices are seen as assertive, while artificially soft voices are seen as submissive by video conferencing participants (Huang et al., 2002).

The second approach, using verbal instructions from participants to devices in the home media space, can provide a high level of control over privacy and could be used to control almost any media space device. For example, a camera could stop recording if a user utters the command "Camera Off!" Several difficulties exist in allowing such verbal control however. Currently the major challenge comes from technology: computers are unable to recognize complex and varying strings of commands. Humans commonly use many variations of speech and currently computers cannot understand changes in the structure of speech, which may easily affect verbal content. If computers were capable of understanding such commands, designers and users would need to be very careful to ensure that the correct devices receive the proper commands (Bellotti et al., 2002). For example, if a telecommuter walks into the room wearing only a bra and issues the command, "Turn on," expecting a second monitor to turn on, it is imperative that a *camera* does not turn on instead and begin capturing video!

The third approach, using verbal instructions from devices within the media space to media space participants, offers a crucial component of privacy feedback. Feedback of the level of privacy being attained is most easily presented through visuals or with audio.

In the case that visuals go unnoticed, audio becomes vital. One approach is to have audio cues as verbal commands played by a device. For example, if a camera is currently capturing the media space when a telecommuter enters her home office, a device could play a verbal command of "Camera is recording!" A second approach is to instead play a recognizable sound (Gaver, 1988). Using the same example, instead of having the computer play a spoken command, the computer could simply play the sound of, say, a camera clicking. Sounds from other devices in the media space can also offer feedback, e.g., the sound of a camera rotating to capture a wall when the user does not want to be recorded (Boyle et al., 2000). This second approach presents a more feasible alternative than the first approach because sounds are more universal than utterances from any one language. Moreover, speech is considered to demand higher cognitive processing than simple sounds. For sounds to be useful, they must of course be recognizable and carry a specific meaning for *all* media space participants.

5.4.2 Non-verbal Behaviors: Presenting and Using Gestures

Non-verbal behavior consists of the use of body language, such as gestures and posture, to control privacy (Altman, 1975). For a home media space, non-verbal behaviors can be used in two ways:

1. gesture-based input for devices within the media space; and,
2. non-verbal instructions between media space participants.

Each of these techniques is now discussed in more detail.

The first approach, gesture-based input for devices within a media space, is certainly feasible and can compliment verbal behaviors much like in face-to-face situations. Here, the user can give the media space explicit instructions using recognized hand or body gestures. For example, using computer vision techniques it is possible to detect if a user tries to block the view of the camera by placing his hand in front of it (Boyle et al., 2000). This gesture can then be used to turn off the camera. When developing gesture-based input for a PDA media player, Pirhonen et al. (2002) found that in general, gestures need to be both easily reproducible and easy to learn by users. They

also discovered that explicit and immediate feedback is necessary for users to be sure they performed gestures correctly.

Two main problems exist with using gesture-based input. First, it can be difficult to detect and interpret gestures as unique inputs. The device that the gesture is intended for must recognize the specific gesture, and furthermore, only that device must respond to it (Bellotti et al., 2002). For example, it would be important that a hand trying to block a camera does not get misinterpreted as a wave and video continues to be recorded as a result. Just the same, other devices should not mistakenly interpret the gesture as being input for them, leaving undesirable consequences. A second problem is that currently there are no methods for rapidly developing systems containing gesture-based input. Most research in this area is being performed in specific application settings, such as recent work by Pirhonen et al. (2002) where a set of specialized gestures was developed for playing music on a PDA.

The second approach for using non-verbal behaviors is simply a replication of that which is done in face-to-face situations where people use body language to control privacy. Co-located participants should have little trouble with this, yet participants separated by distance must rely on the video channel for presenting their non-verbal behaviors. Video must be shown at a level of fidelity high enough for other participants to easily interpret gestures and postures. For example, if a telecommuter desires solitude and moves his hand in front of the camera for a few seconds to signal this desire to his co-worker, it is important that this gesture is not mistaken for a waving hand, which is a common signal for "Hello" in North American culture, and could falsely cause the co-worker to move into interaction with the telecommuter.

5.4.3 Environmental Mechanisms: Virtual Fences, Blinds, and Doors

Environmental mechanisms consist of the use of physical artefacts and features of an environment to control privacy (Altman, 1975). Just as individuals can control their own environment in the physical world, they should be able to control their environment in a home media space. Environmental mechanisms for a home media space can be grouped into four categories:

1. lightweight mechanisms for altering the media space's physical environment;
2. self-appropriation for controlling physical appearance and behavior;
3. adjustable personal space; and,
4. use of timing.

North American culture relies heavily on environmental mechanisms for regulating privacy; thus, a plethora of environmental mechanisms are vital for the success of a home media space within North American culture. Each category of approaches for environmental mechanisms in a home media space will now be explored.

The first approach involves providing users with lightweight mechanisms for altering the media space environment. Mechanisms could include either *explicit* or *implicit* control over privacy. Explicit control could come in the form of either software (e.g., graphical slider) or hardware controls (e.g., physical sliders) whereby the user can alter various environmental attributes, described shortly. Implicit control could involve sensing-based technology detecting the level of privacy desired by media space participants and then automatically altering the environment as needed. Several environmental attributes, specific to media spaces, merit lightweight adjustment either explicitly or implicitly:

- ***camera state:*** the camera can either be recording or not recording.
- ***video fidelity:*** the video could be filtered to hide sensitive details with techniques such as blur filtration. The camera's frame rate could also be adjusted.
- ***camera angle:*** the camera's direction and what is being captured can be adjusted.
- ***camera placement:*** the camera can be placed in any number of locations (discussed below as a method for adjusting personal space).

There remains, of course, other physical attributes that could be controlled, such as the amount that a door is closed, however while important they are less specific to media spaces. When designing a home media space it is important to place as much control in the power of the user as the potential for sensitive information to be broadcast is quite high. This is why using features such as reciprocity for automated control of video

fidelity (discussed in Section 2.3.3) may be inappropriate for a home media space. Reciprocity enforces the notion that participants at both ends of the media space desire the same level of privacy, which may easily not be the case for a home media space.

Just as important as providing lightweight control over environmental attributes, is providing the user with feedback of the current level of privacy being attained. The environment can then be modified automatically to show this state using physical devices such as LEDs, signs, or a mirrored image of what is being broadcast.

Two toolkits, which have been developed in GroupLab at the University of Calgary, make it easy to rapidly prototype media spaces containing lightweight mechanisms for control and feedback of privacy. The first, GroupLab Collabrary, makes it easy to create software with video and audio links and alter attributes such as video fidelity (Boyle and Greenberg, 2002). The second, Phidgets, which contains pre-packaged physical devices and a corresponding software Application Programming Interface (API), makes it easy to rapidly prototype physical interfaces and sensing environments (Greenberg and Fitchett, 2001).

The second approach for controlling privacy in a home media space, self-appropriation, lays in the hands of media space participants. Self-appropriation involves creating an appearance and behavior suitable for the current situation (Bellotti, 1998). Given enough feedback of the level of privacy currently being attained, users have the power to control their own privacy by simply appropriating themselves correctly. This can be difficult in a home media space however. Participants at the home location may be forced to appropriate themselves for the office, which itself can be an infringement on their autonomy. To help alleviate this problem, users can rely on lightweight controls (previously discussed) to help appropriate themselves correctly for both home and the office.

The third environmental mechanism is adjustable personal space. Just like in face-to-face situations, home media space users have the opportunity to utilize personal space for controlling privacy. First, the media space can be set up in any location within the home. The most suitable location for a media space is interestingly enough in a semi-private room. As was found in the results from Chapter 4, the room should be semi-

private in the sense that it is normally dedicated for work use, such as a home office, but could also be used occasionally for other purposes like as a spare bedroom. This type of room offers users a large amount of control over their privacy because it is not commonly used by many people within the home, as opposed to other rooms such as a living room. Second, within the media space the camera can be positioned in any number of locations; camera placement determines what background information is captured. This typically becomes unremarkable over time, as seen in Section 4.3, but care can be taken so that background information does not include areas such as an open doorway where others may be able to see through the doorway into other rooms. Currently, users are restricted to placing cameras within a certain range from their computers because of fixed wire lengths between the camera and computer.

When choosing the camera's location, participants must also be aware of the effect of camera placement on communication. First, the placement of a camera can affect eye contact between two media space participants. If eye contact is desired, a camera should be placed somewhere in the user's direct line of sight, e.g., close to a computer display. Sometimes such placement is difficult, however people are normally less sensitive to eye contact if ones eyes are looking above or below another's eyes (Chen, 2002). This means that if a camera is to be placed outside a person's direct line of sight, users will have fewer difficulties perceiving there is eye contact if the camera is placed either above or below a user's line of sight, as opposed to the left or right of it (Chen, 2002). For media space design, this suggests placing the camera above the display to facilitate eye contact (Chen, 2002). Research by Huang et al. (2002) shows that camera placement can also affect the influence of individuals in group decision making tasks. Artificially tall people, those with a camera looking up at them, are seen as more assertive, while artificially short people, those with a camera looking down at them, are seen as being submissive (Huang et al., 2002). Clearly this type of camera placement can affect relationships between participants of a home media space. If awareness is the primary purpose of the space, as opposed to direct communication such as meetings, these types of placement problems may be less important though.

A media space user can also place the camera at a certain distance away from her in order to use a specific spatial zone (discussed in Section 5.2.3). For example, if a telecommuter desires a colleague to be within her personal spatial zone, she could place her camera roughly two feet away from her. This two foot distance, in addition to the distance between her colleague and his monitor, would place the colleague within the personal spatial zone of 1.5 to 4 feet. Camera distance can also be adjusted with varying lens types and zooming options available in current cameras.

With the fourth environmental mechanism, using timing, users can control their privacy by simply varying when they work. In order to preserve the privacy of others within the home, telecommuters can choose to work and use the home media space when others are not in the home. This mechanism may be less desirable because in order to maintain awareness and close contact, a colleague of the telecommuter may also need to match the same working schedule, which in turn would affect the autonomy of the colleague.

5.4.4 Cultural Mechanisms: Social Solutions

Cultural mechanisms consist of the use of cultural practices and social customs to control privacy (Altman, 1975). Media space participants must always be able to rely on social protocol and cultural mechanisms in the case that technology is not able to balance privacy and awareness adequately. Given that a home media space culture is able to develop, a consensus must be made about several key cultural issues involving the media space:

1. the purpose of the media space;
2. who is allowed to view what is captured; and,
3. what content is appropriate to be seen.

Each of these issues embodies the home media space culture and will now be discussed in more detail.

The first cultural issue involves identifying the purpose of the home media space; what should the media space be used for and, perhaps more importantly, what should the

media space *not* be used for. This thesis is built on the idea that the major purpose of a home media space is to provide awareness between intimate collaborators who are separated by distance. In this situation, the collaborators must have the need and desire to work closely together and maintain a close working relationship. This or another purpose must be clearly defined or there is the risk that the media space will be used for devious purposes such as spying on the collaborator or others captured in the space.

The second cultural issue involves defining who is allowed to view the captured video in the media space. To provide awareness between two collaborators, both should, of course, be allowed to view the captured video in the media space. Collaborators, however, may not feel comfortable with individuals other than a particular work colleague viewing them. For example, if a female telecommuter is normally only viewed by another female colleague in the media space, the telecommuter may feel quite uncomfortable to have a male colleague viewing the space if this is not her desire. It is possible for viewers to stand outside the zone captured by the camera so media space participants may not always know who is watching them. For this reason, a mutual understanding of who should and should not be viewing the media space is necessary.

The third cultural issue identifies what content is appropriate for participants to see in the media space; participants must know who is appropriate to be seen as well as what is appropriate to be seen. "Who" could include or exclude other people in a telecommuter's home, like family members and possibly children. In an office environment, "who" could include colleagues not wishing to partake in the media space. "What" could include certain activities, behaviors, or confidential background information. If the media space begins to capture *something* a participant should not be seeing, or *someone* he should not be seeing, the participant should adhere to the social protocol and not look, or block the region of the screen showing the video being transmitted. As part of this social protocol, participants should attempt to present themselves appropriately for other media space participants to view.

If media space participants are aware of these core cultural issues and abide by each of them, then they are able to rely on social protocols to mediate privacy concerns while still gaining awareness. However, the problem is that cultural norms must evolve over

time and privacy threats will be heightened until a point when the culture has become well defined. There is also the chance that participants will not always abide by established cultural protocols. In such cases, however, just like in real life, there are social consequences for not abiding by the given social norms.

5.5 Summary

In this chapter, I have outlined a framework for the design of a home media space. This framework contains a set of privacy-protecting strategies that can be used within a home media space to provide the user with control and feedback of privacy. First, I use social-psychological theory to define culture and discuss how each culture uses various *privacy mechanisms* to regulate interaction and privacy. Privacy mechanisms are the behaviors and actions that humans employ to control privacy, and can be categorized into four groups:

1. ***Verbal behaviors****:* the use of the content and structure of what is being said;
2. ***Non-verbal behaviors****:* the use of body language, e.g., gestures and posture;
3. ***Environmental mechanisms****:* the use of physical artefacts and features of an environment, e.g., walls, doors, spatial zones, timing; and,
4. ***Cultural mechanisms****:* the use of cultural practices and social customs.

These four categories form a coherent system for regulating privacy. Depending on the culture, mechanisms can be used from any number of these categories and are often substituted depending on the social environment and situation.

When people engage in the use of a home media space, they are building a home media space culture where privacy mechanisms presented by the design of the space along with culture norms can be used to regulate privacy. Privacy mechanisms in a home media space entail privacy *feedback*, which informs people when and what information is being captured and to whom it is made available (Bellotti, 1998). Just the same, privacy *control*—the power to regulate what information is available for others and who can gain access to it (Bellotti, 1998)—is a crucial component for privacy mechanisms.

For each of the four categories of privacy mechanisms, I have identified various strategies for providing control over privacy and feedback of the current privacy level in a home media space:

1. ***Verbal behavior:***
    - ***i.*** verbal instructions between media space participants, e.g., continuous or optional audio links for distance-separated participants;
    - ***ii.*** verbal instructions to devices within the media space, e.g., voice activated software; and,
    - ***iii.*** verbal instructions or sound cues from devices in the media space to media space participants, e.g., an audio file is played with a sound or voice command.
2. ***Non-verbal behavior:***
    - ***i.*** gesture-based input for devices within the media space, e.g., recognized hand gestures to alter what is captured; and,
    - ***ii.*** non-verbal instructions between media space participants, e.g., gestures through the video channel, which are similar to face-to-face situations.
3. ***Environmental mechanisms:***
    - ***i.*** lightweight mechanisms for altering the media space's physical environment, e.g., implicit or explicit controls to adjust camera state or video fidelity, LEDs for feedback;
    - ***ii.*** self-appropriation for controlling physical appearance and behavior, e.g., use video fidelity to appropriate oneself for both home and the office;
    - ***iii.*** adjustable personal space, e.g., place the camera a certain distance away from the participant, choose a room for the media space; and,
    - ***iv.*** use of timing, e.g., adjust one's work schedule.

4. Cultural mechanisms:

- ***i.*** define the purpose of the home media space, e.g., awareness between intimate collaborators separated by distance, not malicious use;
- ***ii.*** define who is allowed to view what is captured, e.g., the telecommuter and collaborator, specific family members or none at all, specific co-workers; and,
- ***iii.*** what content is appropriate to be seen, e.g., the telecommuter working, but not working shirtless.

Each of these categories presents viable solutions for balancing privacy and awareness in a home media space, albeit some are better than others. Each category also comes with its problems and depending on the current situation it may be suitable to use certain privacy mechanisms rather than others, or to combine mechanisms.

North American culture normally relies on environmental mechanisms, such as walls, doors, and personal space, to control privacy. As such, it is apparent that the design of a home media space for use within North America should reflect this and thus provide a plethora of environmental mechanisms for controlling privacy. In instances when environmental mechanisms are not able to balance privacy and awareness adequately, the onus is placed on strategies from other categories of mechanisms, such as cultural norms defined by the home media space culture.

In the next chapter, I use the design framework presented in this chapter to design a home media space. The design employs user interface design techniques for presenting the privacy-protecting strategies from this framework.

Chapter 6. The Design of a Context-Aware Home Media Space

In this chapter[5], I discuss the design of a context-aware home media space (HMS) using the framework presented in Chapter 5. In the home media space, I leverage the framework by presenting users with privacy-protecting strategies using the same four categories of privacy mechanisms that are used by humans in everyday life. Using context-aware technology and dedicated physical controls, users are provided with *explicit* and *implicit* control over privacy, along with *audio* and *visual feedback*.

First, I give an overview of the design philosophy of our context-aware HMS, including the design principles we feel a HMS should be based on and the specific HMS elements we use to support these design principles. Second, I describe how our HMS design uses each element within the HMS, along with a set of rules, to balance privacy and awareness for the telecommuter and others in the home. Third, I discuss how we have leveraged the framework from Chapter 5 by designing privacy-protecting strategies for a HMS using the same four categories of privacy mechanisms. Fourth, I discuss the software and hardware implementation of the HMS.

The home media space design presented in this chapter is not formally evaluated; however, it is important because it presents one approach for the design of such a space and the use of the framework presented in Chapter 5.

---

[5] A version of this chapter is published as:

Neustaedter, C., and Greenberg, S. (2003) **The Design of a Context-Aware Home Media Space for Balancing Privacy and Awareness.** Report 2003-722-25, *Department of Computer Science,* University of Calgary, May, 2003.

6.1 The Design Philosophy for a Context-Aware HMS

This section outlines the five principles behind the design of our context-aware HMS. First, we explain each of our design principles and why they are included in our design philosophy based on the design framework from Chapter 5. Second, to set the scene of our design, we describe the design elements that arose from our five principles.

6.1.1 Design Principles for a Context-Aware HMS

The results of our study in Chapters 3 and 4 highlighted the importance of providing user control over information conveyed through a video media space. For this reason, in Chapter 5, we began investigating how humans regulate privacy in everyday life through various behaviors and actions called privacy mechanisms (Altman and Chemers, 1980). When individuals require more privacy, they use these mechanisms to let others know they desire less interaction. Just the same, when individuals require more interaction, they use these mechanisms to let others know they desire less privacy. These privacy mechanisms are very natural and often form an unconscious act (Altman, 1975). Based on this research and the framework in Chapter 5, we believe the design of a HMS should use the following design principles:

1. Existing privacy mechanisms should be leveraged for home-based video conferencing systems;
2. Implicit actions using context-aware technology can regulate privacy;
3. No implicit action should ever decrease the amount of privacy without first warning the user and providing the opportunity to stop the operation;
4. Explicit actions using dedicated physical controls and gesture recognition can regulate privacy; and,
5. Visual and audio feedback makes the state of the system easily discernable at any time.

The first principle helps to create privacy mechanisms for a HMS that are both easy to understand and natural to use because they are based on techniques already familiar to humans. Our design supports this principle by providing users with privacy-protecting

strategies from the same four categories used by humans in everyday life (discussed in more detail later) using the design framework from Chapter 5.

Privacy regulation in real life is lightweight and often transparent. Such implications should also be available to HMS users. Thus, as the second principle states, privacy-protecting strategies in a HMS should also be lightweight and transparent. Our design supports this principle by using context-aware computing as a tool for balancing privacy and awareness through implicit means. Unlike previous work in context-aware computing (Want et al., 1992, Schilit and Theimer, 1994, Schilit et al., 1994), we enable one specific location—a home office/spare bedroom—with technology that senses who is around and then infers privacy expectations through a simple set of rules.

There is still a considerable gap between human expectations and the abilities of context-aware systems (Dey, 2001, Erickson, 2002). Context-aware systems can make mistakes and it is important that these mistakes do not increase privacy threat; the third design principle addresses this problem. Our design supports this principle by first warning users that an implicit action has initiated a privacy-decreasing operation; and second, by providing an opportunity for users to override this operation.

The fourth principle also addresses the previously mentioned problem by recognizing that we need to keep the user in the "control loop." Our design supports this principle by providing users with dedicated physical and graphical controls, where explicit actions such as adjusting a physical slider or gesturing towards the camera will alter the privacy level. We recognize that explicit control must absolutely be lightweight and executed with almost trivial effort.

The fifth principle is important because users must be able to fine tune the privacy/awareness balance as desired. To do this fine tuning, they must know how much privacy is currently being maintained. Our design supports this principle by providing feedback of the achieved privacy level through audio and visual cues, rendered on both physical displays (such as LEDs) and on the screen. This feedback is both understandable and continually available.

Figure 6.1: The HMS GUI: the telecommuter (top) and colleague (bottom).

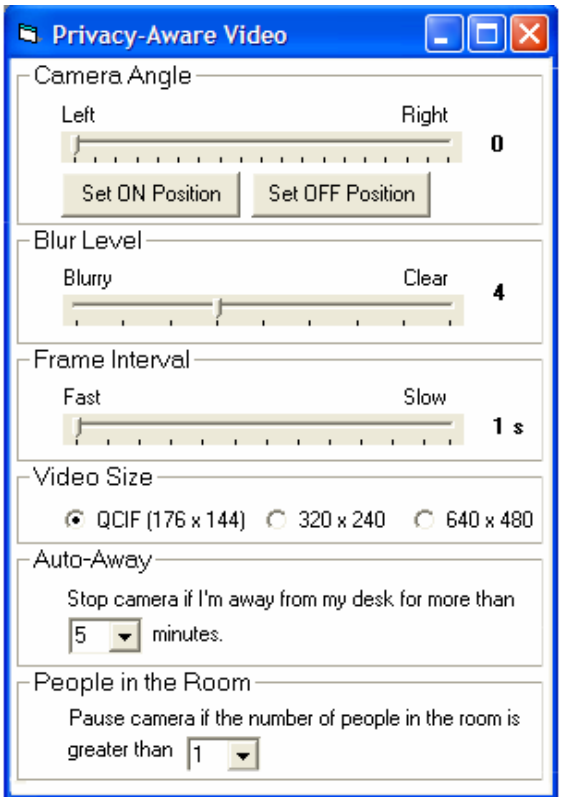

Figure 6.2: A configuration window to adjust various HMS attributes.

6.1.2 Elements of a Context-Aware HMS

To foreshadow the details of our design, this section outlines the elements of our HMS that arise from the five design principles. The subsequent section describes their importance by outlining how they work together to regulate privacy.

Figure 6.1 shows the HMS's graphical user interface (GUI) as seen by the telecommuter: the top window shows a mirrored image of the telecommuter as it is captured, and the bottom window shows the telecommuter's colleague. A third window contains additional options (Figure 6.2) and is displayed by clicking the options button in the telecommuter's toolbar (Figure 6.1, top). The other graphical controls are described below. Figures 6.3 and 6.4 show the layout of the HMS in the home office/spare bedroom of a telecommuter. The design is specific to this room layout, but the ideas presented can be applied to a variety of home settings. However, other rooms without a direct entrance, such as a living room, may be more problematic for designs.

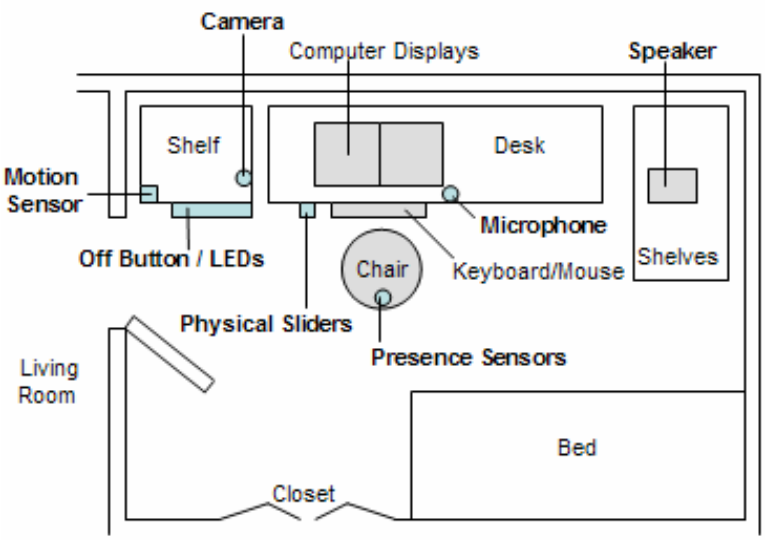

Figure 6.3: An overview of the HMS layout within the home office/spare bedroom.

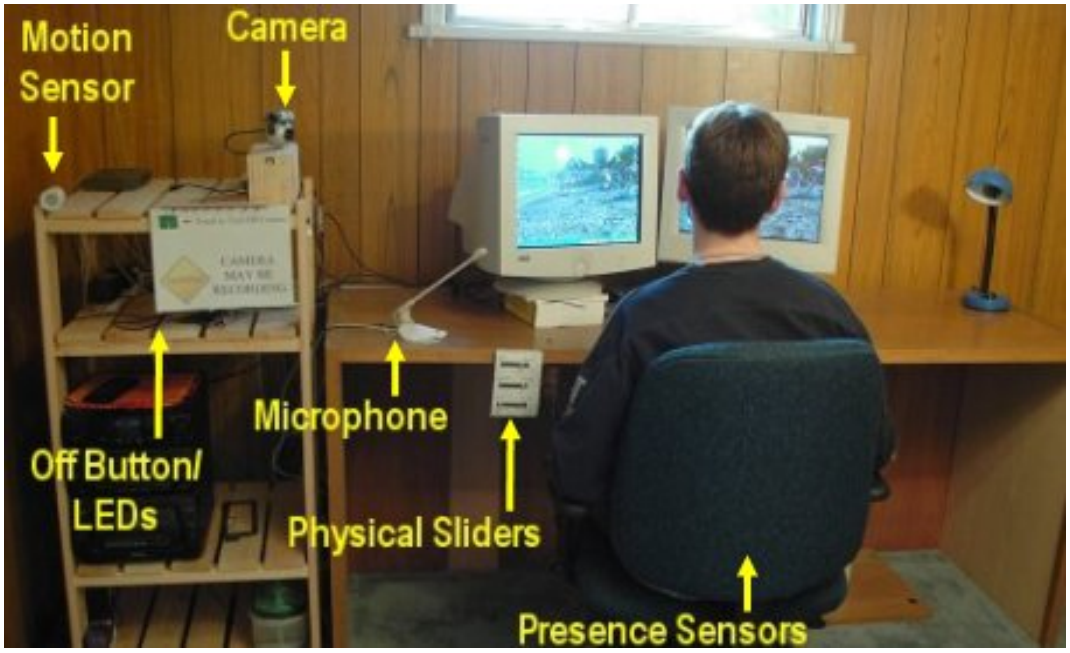

Figure 6.4: The layout of the HMS within the home office/spare bedroom

We support the five HMS design principles, discussed previously, by including specific elements within our design:

Camera state. The camera can be in one of three states: Play (Figure 6.1), Pause (Figures 6.5, 6.6), and Stop (Figure 6.7). In the play state, the camera is capturing and broadcasting video to other HMS participants (Figure 6.1). In the pause state, the camera no longer captures and broadcasts video to other HMS participants; however, other

Figure 6.5: The HMS paused with the telecommuter leaving his chair.

Figure 6.6: The HMS paused with multiple people in the room.

Figure 6.7: The HMS stopped and camera facing the wall.

availability information is sent including the last video frame captured of the user and a count of the number of people in the room (Figures 6.5, 6.6). In the stop state, like the pause state, the camera no longer captures video and the last image broadcast is of the wall (Figure 6.7), after the camera has rotated away from the user (discussed below). The major difference between the pause and stop states is that it is more difficult to move out of the stop state (described in more detail later). Users can explicitly move between states by clicking the play, pause, and stop buttons (three leftmost buttons, respectively in Figures 6.1, 6.5, 6.6, and 6.7).

Capturing angle. The camera, mounted on a rotating motor (Boyle et al, 2000), is placed near the door and, given the desired camera angle, can capture any region of the room, except the doorway (Figures 6.3, 6.4: Camera). This is important as the living room is not visible (Figure 6.3). We provide the user with dedicated physical sliders (Figure 6.3, 6.4: Physical Sliders, Figure 6.8-top) and graphical sliders (Figure 6.2) to explicitly alter the capturing angle.

Video fidelity. Users can adjust the captured video's fidelity by explicitly adjusting the level of blur filtration used (Figures 6.1, 6.2, 6.8-middle), the camera's frame rate (Figures 6.2, 6.8-bottom), or the camera's frame size (Figure 6.2). We provide the user with dedicated physical (Figure 6.3, 6.4: Physical Sliders) and graphical controls to explicitly adjust these three components of video fidelity.

Gesture-activated blocking. Users can easily turn off the camera by explicitly blocking it with their hand. We detect this gesture with a proximity sensor mounted on

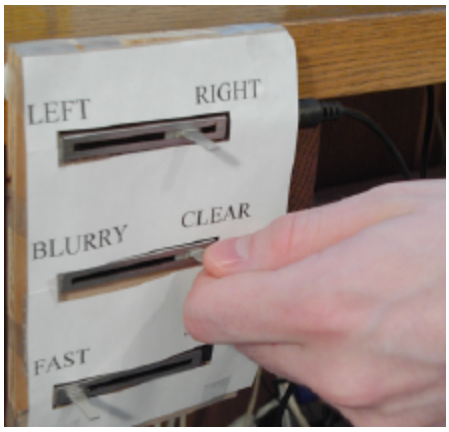

Figure 6.8: A user adjusts the blur level with a dedicated physical slider.

Figure 6.9: A user blocks the camera with his hand to turn it off.

top of the camera (Figures 6.3, 6.4: Camera, Figure 6.9). This can also be done using computer vision techniques (Boyle et al., 2000).

Gesture-activated voice. Users can easily open an audio channel by explicitly moving their hand over a microphone (Figures 6.3, 6.4: Microphone, Figure 6.10). Moving one's hand away from the microphone closes the audio channel. We detect this gesture with a light sensor mounted on top of the microphone. This can also be done (perhaps more accurately) using other sensors, such as proximity or capacitive sensors.

Easy-off button. Users can easily turn off the camera by touching an off button (Figures 6.3, 6.4: Off Button, Figure 6.11). We detect this explicit action with a capacitive sensor acting as the button, but this could also be done (and appear more

Figure 6.10: A user moves his hand over the microphone to open an audio link.

Figure 6.11: A sign containing LEDs (circled at the top) and an off button (circled below the LEDs).

Figure 6.12: The RFID reader and light sensor to detect the telecommuter's presence.

Figure 6.13: The telecommuter's RFID tag.

realistic) with a control resembling a real-world push button (Jancke et al., 2001).

Telecommuter detection. We know if the telecommuter is present at the computer by detecting (with a light sensor, Figures 6.3, 6.4: Presence Sensors, Figure 6.12) the implicit act of someone sitting down in or standing up from the desk chair. We use a radio frequency identity (RFID) tag in the pocket of the telecommuter (Figure 6.13) and a RFID reader (Figures 6.3, 6.4: Presence Sensors, Figure 6.12) in the chair to identify if the individual sitting is the telecommuter. If the telecommuter is not present, we can tell how long she has been away from the computer. Our *telecommuter detection* is not a realistic solution because of limits imposed by our RFID reader, yet it works for our prototype. Other approaches could include using Active Badges (Want et al., 1992) or embedding RFID tags within 'work' shirts worn by the telecommuter. This helps because it can ensure the telecommuter is appropriately dressed before the HMS can be used; however, now people must wear this special garment.

Figure 6.14: The infrared motion sensor used to detect the presence of people in the home office/spare bedroom.

Family/friend detection. We know if people other than the telecommuter are present in the room by using an infrared motion detector (Figures 6.3, 6.4: Motion Sensor, Figure 6.14) to detect the implicit act of walking into and out of

the room. This could be done more accurately with computer vision techniques; however, our solution does not require a camera to always be capturing the room's activities.

Visual feedback. We use several visual cues to let the user know how much privacy is currently being maintained, e.g., a sign (Figure 6.11), LEDs (Figure 6.11-top), the camera's direction, mirrored video (Figure 6.1, top), 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, e.g., the sound of a camera clicking and the sound of the camera rotating (Gaver, 1988).

There are many ways to create each of these elements and more accurate sensors exist than the ones we have chosen to use for our prototype. We have chosen methods and sensors that allowed us to rapidly and inexpensively prototype each element.

In the next section, we describe how these elements work together, along with a set of rules, to reduce privacy threats. We demonstrate this with a series of scenarios based on real telecommuting situations.

6.2 Rules for Balancing Privacy and Awareness

While the individual privacy control and feedback elements are suggestive, it is how they interoperate that is important. Our HMS design uses each element within the HMS, along with a set of rules, to balance privacy and awareness for the telecommuter and others in the home. Table 6.1 summarizes how the design elements are either: controlled, used for explicit or implicit control, or used as feedback. Each row in the table describes how one media space attribute (column 1) is controlled either explicitly (column 2) or implicitly (column 3). The fourth and fifth columns describe the *audio* and *visual feedback* that indicate to the users that the attribute in column 1 has changed and what its current value is. The first five rows of the table describe the transitions between the three *camera states*. The remaining three rows describe other HMS attributes that can be controlled.

	1	2	3	4	5
	Attribute Controlled	**Explicit Control**	**Implicit Control**	**Audio Feedback**	**Visual Feedback**
1	Stop to Play	Click play button	None	Camera clicking; Camera rotating	LEDs on; Camera rotates to face you; Mirrored video
2	Pause to Play	Click play button	Telecommuter sits in chair; Family/friend leaves room	Same as above; Camera Twitches	Same as above; Camera Twitches
3	Play to Stop	Click stop button; Block camera with hand; Touch off button	None	Camera rotating	LEDs off; Camera rotates to face the wall; Mirrored video
4	Play to Pause	Click pause button	Telecommuter stands up out of chair; Family/friend enters room	Same as above	Same as above
5	Pause to Stop	Click stop button; Block camera with hand; Touch off button	Telecommuter leaves the room for an extended period of time	None	Mirrored video
6	Capturing angle	Adjust physical or graphical slider	Change in camera state	Camera rotating	Slider position; Camera position; Mirrored video
7	Video fidelity	Adjust physical or graphical control	None	None	Control position; Mirrored video
8	Audio link	Moves hand over microphone base	None	Own voice	None

Table 6.1: Control and feedback mechanisms found in the HMS

We now present a series of scenarios that detail the privacy risks involved with using a HMS, the set of privacy rules we have created to address them, and how the HMS implements each rule to balance privacy and awareness.

6.2.1 Providing Awareness While Masking Embarrassing Acts

The first scenario illustrates one typical use of the HMS by a telecommuter, named Larry, who is working at home and using the media space to provide awareness to a close-working colleague at the office. Larry enters his home office/spare bedroom, dressed in casual pants and a golf shirt. While Larry is working at his computer, he suddenly sneezes. Naturally, he proceeds to blow his nose. Forgetting that the camera is capturing him, Larry begins to pick his nose at great length.

Privacy Risks: Larry is dressed appropriately to be seen at an office, yet he does not want his colleague to view him doing embarrassing, unconscious acts like picking his nose.

Rule 1: If just the telecommuter is present at the computer, the HMS assumes more awareness and less privacy is desired.

Design: This is Larry's first use of the HMS today and the *camera state* is Stop when Larry sits down at the computer. To turn the *camera state* to play (Table 6.1: Row 1), Larry must explicitly click the play button (Figure 6.1, leftmost button). Once the *telecommuter detection* has identified that it is indeed Larry at the computer, the HMS provides more awareness by moving the *capturing angle* away from the wall to record Larry. *Visual* and *audio feedback* lets Larry know the camera is now capturing (Table 6.1: Row 1): the LEDs turn on (Figure 6.11), the mirrored video updates (Figure 6.1), the computer plays the sound of a camera clicking, and, Larry sees and hears the camera rotate. Larry can fine tune the awareness information and attempt to mask embarrassing acts with *video fidelity* (Table 6.1: Row 7), e.g., by adjusting a dedicated physical slider to blur his image or adjust the frame rate.

6.2.2 Providing Privacy When Others Use the Computer

The second scenario illustrates what happens when the telecommuter leaves his desk and others use the computer. Larry is working at his computer when he leaves to get a coffee from the kitchen. Larry's wife, Linda, who is still in her pajamas, comes in to the home office to quickly check her email. Linda leaves the room just as Larry returns. Larry sits down and continues working.

Privacy Risks: Larry is appropriate to be viewed on camera and faces no privacy risks. Linda is not appropriate to be viewed, nor does she want to be viewed: Linda faces a threat/benefit disparity.

Rule 2: If someone other than the telecommuter is present in the room, the HMS assumes more privacy and less awareness is desired.

Design: The *telecommuter detection* knows that Larry has left his desk chair and changes the *camera state* to paused (Table 6.1: Row 4). *Visual* and *audio feedback* lets

Larry know the camera is no longer capturing (Table 6.1: Row 4): the LEDs turn off (Figure 6.11), the mirrored video updates (Figure 6.5), and, Larry sees and hears the camera rotating. The colleague maintains awareness by seeing Larry leave his chair in the last image broadcast (Figure 6.5).

When Linda enters the room, the *family/friend detection* flashes the LEDs (Figure 6.11) and plays the sound of the camera clicking to warn Linda to make sure the camera is off. *Visual feedback* shows her that the *camera state* indeed remains paused (Table 6.1: Row 4). Linda checks her email and is not captured on camera.

When Larry returns to his desk chair, the *telecommuter detection* unpauses the camera, but first warns Larry this is about to happen by twitching the camera left and right (Table 6.1: Row 2); just as people signal their intentions, so does the camera. This complies with our third design principle. *Visual* and *audio feedback* shows Larry that the *camera state* is again Play (Table 6.1: Row 2).

6.2.3 Using Gestures to Regulate Privacy

The third scenario illustrates how the telecommuter can use gestures to control HMS attributes, which in turn affect his privacy. Larry is working at his computer composing an email and drinking his coffee. Just then, Larry knocks his mug and coffee spills all over his shirt! Larry removes his shirt and then notices the camera facing him. Larry blocks the camera with his hand then tells his colleague (through the HMS) that he has to go get a new shirt.

Privacy Risks: Larry does not want to be seen shirtless, yet he still wishes to maintain a level of awareness with his colleague.

Rule 3: The HMS must provide simple lightweight means to immediately disable the capturing device, yet still maintain awareness through alternate channels.

Design: Larry can choose one of two explicit methods to instantly stop the camera: *gesture-activated blocking* or *easy-off button* (Table 6.1: Row 3). *Visual* and *audio feedback* lets Larry know the *camera state* has changed (Table 6.1: Row 3). Larry wants to maintain awareness and tell his colleague of his predicament without using the video channel so he uses *gesture-activated voice* to open the optional audio link.

6.2.4 Providing Privacy When Others Enter the Room

The fourth scenario illustrates what happens when multiple people enter the home office/spare bedroom. Larry is working at his computer in the home office/spare bedroom when Linda, who has just finished taking a shower in the bathroom next door, walks into the room to retrieve her bathrobe from the closet. Linda puts on her bathrobe and leaves the room.

Privacy Risks: Linda does not want to be captured on video, especially while she is naked! Linda again faces a threat/benefit disparity, while Larry still wants to provide awareness information to his colleague.

Rule 4: If more than just the telecommuter is present in the room, the HMS assumes more privacy and less awareness is desired.

Design: The *family/friend detection* knows that Linda has entered the room and moves the *camera state* to paused (Table 6.1: Row 4). *Visual* and *audio feedback* indicates that the *camera state* has changed (Table 6.1: Row 4). Larry's colleague maintains a level of awareness with the presentation of alternate awareness information when the camera is paused: the number of people in the room, and the image of Larry sitting at his desk (Figure 6.6). Using these two pieces of information, it is possible for Larry's colleague to infer that Larry is still working at his desk. If the capturing angle of the camera is set properly, the still image shown to Larry's colleague while the camera is paused will only contain Larry.

Once the *family/friends detection* knows that Linda has left the room (Table 6.1: Row 2) and the *telecommuter detection* indicates that Larry is still at the computer, the *camera state* will return to Play once it first warns Larry with *visual* and *audio feedback* (Table 6.1: Row 2).

6.2.5 Finishing Work and Leaving the Space

The fifth scenario illustrates what happens when the telecommuter finishes working and leaves the HMS. Larry has finished working for the day and leaves the home office.

Privacy Risks: The HMS is still active when the telecommuter is finished working; future use of this room may threaten privacy.

Rule 5: If the telecommuter is away from the computer for an extended period of time, the HMS will move to a permanent, non-recording state.

Design: The *telecommuter detection* notices Larry leaving and the *camera state* pauses. After being away from his desk for five minutes, the *camera state* moves to Stop and now the last image shown to Larry's colleague is of the wall (Figure 6.7). This timeout interval can be customized in Figure 6.2. The non-recording state is permanent in the sense that to start working again, Larry must explicitly click the play button (Figure 6.1). Until this time, the camera will not turn on and no video will be captured; thus, no privacy violations will occur while Larry is not working.

6.3 Supporting Privacy Mechanisms

We now describe how we have leveraged the design framework from Chapter 5 and the four categories of privacy mechanisms by designing privacy-protecting strategies for a HMS that fall into the same categories of mechanisms used by humans for privacy regulation in everyday life.

6.3.1 Verbal Behavior: Sound and Voice

Verbal behavior consists of the use of content and structure of what is said to control privacy (Altman, 1975). For example, if a family member approaches the home office while the telecommuter is currently working she may say, "I'd like to be left alone," if she would like to have more privacy or alternatively, "please come in," if she desires interaction. We use verbal behaviors in two ways within our design:

1. verbal instructions between media space participants; and,
2. verbal instructions or sound cues from devices in the media space to media space participants.

The first approach is trivially supported in the HMS's design for co-located HMS users (e.g., the telecommuter and others in the home): they can simply speak to others in the same location. Distance-separated users of the HMS must rely on a voice channel for this approach. The tradeoff is that we want an audio link, yet not the additional privacy threats found with a continuous audio link (Hudson and Smith, 1996). For this reason,

our design provides an optional audio link where *gesture-activated voice* allows users to easily engage and disengage the audio link.

The second approach offers a crucial component of privacy feedback. Feedback of the level of privacy being attained is most easily presented through visuals or with audio. In the case that visuals go unnoticed, *audio feedback* becomes vital. We use the sound of the camera rotating and the sound of a camera clicking as audio feedback mechanisms.

6.3.2 Non-Verbal Behaviors: Presenting and Using Gestures

Non-verbal behavior consists of the use of body language, such as gestures and posture, to control privacy and can either be implicit or explicit (Altman, 1975). For example, in an exam situation, people may try to block or cover their test paper, indicating their desire for privacy. We use non-verbal behaviors in two ways within our design:

1. gesture-based input for devices within the media space; and,
2. non-verbal instructions between media space participants.

The first approach can compliment verbal behaviors much like in face-to-face situations. Gesture-based input offers a lightweight means to control devices; users can give the media space explicit instructions using recognized hand or body motions. Our HMS uses *gesture-activated blocking* and *gesture-activated voice.*

The second approach is simply a replication of that which is done in face-to-face situations where people implicitly or explicitly use body language to control privacy. Co-located users (e.g., the telecommuter and others at home) should have little trouble with this, yet users separated by distance must rely on the video channel for presenting their non-verbal behaviors. *Video fidelity* must be high enough for other participants to easily interpret gestures and postures.

6.3.3 Environmental Mechanisms: Virtual Fences, Blinds, and Doors

Environmental mechanisms consist of the use of physical artifacts and features of an environment to control privacy (Altman, 1975). For example, to limit neighbors from viewing one's backyard, a fence may be built or a large row of trees could be planted. Just as individuals can control their own environment in the physical world, they should

be able to control their environment in a HMS. The environmental mechanisms for a HMS that we support can be grouped into three categories:

1. lightweight mechanisms for altering the media space's physical environment;
2. self-appropriation for controlling physical appearance and behavior; and,
3. adjustable personal space.

In the first approach, providing users with lightweight mechanisms to alter the environment, allows for easy and simple privacy regulation. Our design allows explicit control over *camera state, capturing angle,* and *video fidelity*; and implicit control over *camera state* and *capturing angle.*

The second environmental approach lays in the hands of media space users. Self-appropriation involves creating an appearance and behavior suitable for the current situation (Bellotti, 1998). We provide users with continual and understandable *visual* and *audio feedback* of the level of privacy currently being attained so that users have the power to control their own privacy by simply appropriating themselves correctly (Bellotti, 1998). To help home participants appropriate themselves for both home and the office, we provide lightweight controls to help users adjust *video fidelity.* 'Work' shirts with embedded RFID tags could also provide an interesting solution to this problem.

The third environmental approach allows HMS users to utilize personal space for controlling privacy, just like in face-to-face situations. Personal space is inherently supported by allowing users to place the camera at a location desirable for supporting awareness; our camera is placed approximately 1.5 to 2 feet away from the user with a side view of the user.

6.3.4 Cultural Mechanisms: Social Solutions

Cultural mechanisms consist of the use of cultural practices and social customs to control privacy (Altman, 1975). Although it may often go unnoticed, each culture contains a set of learned social practices and customs that have evolved and developed over time (Altman, 1975). Since the HMS has yet to be extensively used by individuals, we are not able to describe the use of cultural mechanisms to regulate privacy. The importance,

however, is that given an established set of social protocols, users can rely on them to regulate privacy when technology does not suffice. In the case that social norms are not followed, social ramifications may be in order.

6.4 Software and Hardware

The HMS is designed as an ActiveX® Control, which can be easily used with languages supporting Microsoft COM technologies, e.g., Visual C++, C#, Visual Basic. Currently, our research laboratory uses Greenberg and Rounding's (2000) Notification Collage (NC) for supporting casual interaction and informal awareness (Figure 6.15). The NC allows users to post various media items such as "sticky notes" (containing text), video snapshots, or web page thumbnails. Because of the popular use of the NC in our research lab, the HMS ActiveX® Control is used as a media item that users of the NC may post for others to see (circled in Figure 6.15). In our lab, the NC is typically used by multiple colleagues; however, a NC can be created for connecting groups as small as two people, which may be desirable for a HMS.

Two toolkits, developed at the University of Calgary, were used to develop the

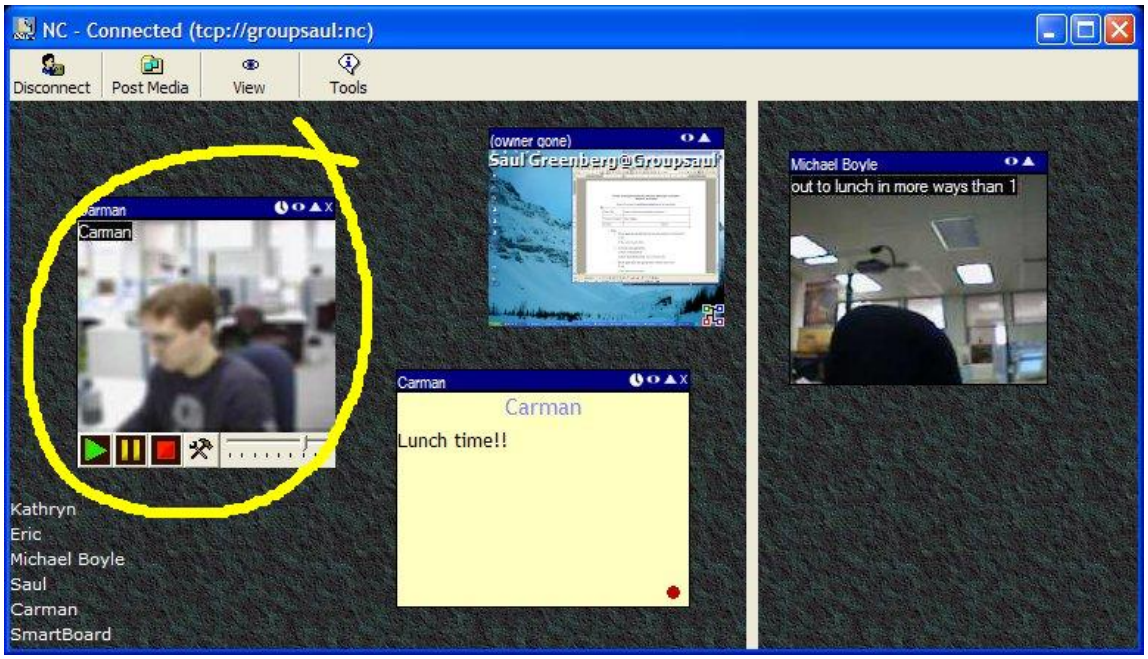

Figure 6.15: The use of the HMS within Greenberg and Rounding's (2001) Notification Collage.

HMS prototype. The first, Collabrary, makes it easy to create software with video and audio links and alter attributes such as video fidelity (Boyle and Greenberg, 2002). In the HMS, the Collabrary's shared dictionary component is used to capture and transmit video and audio between users of the HMS. The second toolkit, Phidgets™, which contains pre-packaged physical devices and a corresponding software Application Programming Interface (API), makes it easy to rapidly prototype physical interfaces and sensing environments (Greenberg and Fitchett, 2001). In the HMS, all of the sensors and controls are Phidget™ devices and are accessed using the Phidget™ API.

The importance of these two toolkits is that, in addition to the implementation of the context-aware HMS, we were also able to focus our research on design issues. As such, we were able to decide and explore how context-aware computing could be used, what its effects would be, and if our techniques were appropriate given our research goal of balancing privacy and awareness.

6.5 Design Experience

Our home media space was designed over a period of several months. During this process, I routinely worked at home as a telecommuter and used the home media space within my own home office/spare bedroom. At the time, I lived alone, yet my fiancé would periodically come over to visit. While it was very easy to regulate my privacy using the home media space when I was at home by myself, the situation was much more interesting when my fiancé was present. She did not like the camera capturing the home office and would often avoid the room because of it, whether it was on or not. Prior to the home media space being set up, my fiancé would turn the camera to face the wall to ensure she would not be captured when using the computer. Once the home media space contained features to regulate privacy for family/friends, she was more likely to enter the home office and liked the fact that the camera would automatically rotate away when she was present. Whether actually recording or not, if the camera was facing my fiancé, her perception was that it was recording.

6.6 Summary

This chapter presents the rationale and prototype design of a home media space (HMS). The HMS is designed specifically for the telecommuter who chooses to work at home, but who still wishes to maintain a close-working relationship with particular colleagues at remote office environments. This chapter contributes a set of five design principles for a HMS and a prototype HMS which illustrates these principles. Specifically, I explain how and why:

1. Existing privacy mechanisms are leveraged for use in home-based video conferencing systems;
2. Implicit actions using context-aware technology can regulate privacy;
3. No implicit action should ever decrease the amount of privacy without first warning the user and providing the opportunity to stop the operation;
4. Explicit actions using dedicated physical controls and gesture recognition can regulate privacy; and,
5. Visual and audio feedback makes the state of the system easily discernable at any time.

Using these five design principles, we created a set of HMS elements to support them. These elements are then combined with a set of privacy rules to balance privacy and awareness for the telecommuter and others in the home:

Rule 1: If just the telecommuter is present at the computer, the HMS assumes more awareness and less privacy is desired.

Rule 2: If someone other than the telecommuter is present in the room, the HMS assumes more privacy and less awareness is desired.

Rule 3: The HMS must provide simple lightweight means to immediately disable the capturing device, yet still maintain awareness through alternate channels.

Rule 4: If more than just the telecommuter is present in the room, the HMS assumes more privacy and less awareness is desired.

Rule 5: If the telecommuter is away from the computer for an extended period of time, the HMS will move to a permanent, non-recording state.

When more privacy is needed, the camera moves into a non-recording state (Pause/Stop) and alternate awareness information is presented, e.g., a still image showing availability, a count of the number of people present in the room. When more awareness is needed, the camera moves into a recording state (Play) and elements, such as video fidelity, can be adjusted to fine tune the amount of awareness/privacy that is gained. Visual and audio feedback continually shows users the current privacy/awareness balance being maintained. In the case of a "privacy emergency," the capturing device can be easily disabled using explicit actions like blocking the camera, or touching an off button.

The home media space design presented in this chapter is our first theoretically based prototype of a home media space. As such, we do not claim that our home media space is well designed. The actual use of context-aware software and dedicated physical controls in our HMS has yet to be formally evaluated for its effectiveness in balancing privacy and awareness. In retrospect, our design seems complex and not all of our control and feedback mechanisms are entirely natural. We realize now that the camera is the most important source of information for users; thus, control and feedback centred around the camera should be the focus for future redesigns. Despite this, the ideas presented in this chapter provide a general approach for integrating the framework from Chapter 5, including the privacy mechanisms used by people in their physical environments, into a HMS. By using two toolkits, including a set of pre-packaged physical devices and sensors, we were able to focus our research on understanding how context-aware computing can be used in real-world applications.

Chapter 7. Conclusions

I conclude this thesis by summarizing the research contributions. First, I reiterate the thesis problems from Chapter 1. Second, I describe my research contributions by outlining how I have solved each of my thesis goals from Chapter 1. Third, I suggest areas of future work for privacy/awareness and home media spaces.

7.1 Thesis Problems

Chapter 1 outlined three research problems within the area of privacy in video media spaces, dealing with telecommuters who use home media spaces:

1. **We do not know if blur filtration is able to balance privacy and awareness in a home media space.** Previous research (Boyle et al., 2000) has shown that distortion filters, such as the blur filter, are able to balance privacy and awareness for benign office situations. Yet we do not know if this balance is achievable for home use of video, as home situations present far riskier situations than office environments.

2. **We do not know what other privacy-protecting strategies, if any, are appropriate for balancing privacy and awareness in a home media space.** Research on privacy-protecting strategies for video media spaces has again primarily focused on office settings, rather than homes. It is unclear what other privacy-protecting strategies, aside from distortion filters, may be suitable for balancing privacy and awareness in home settings.

3. **We do not know what user interface techniques are appropriate for presenting users with privacy-protecting strategies in a home media space.** Privacy-protecting strategies for balancing privacy and awareness in a home media space must be presented to users in a simple, lightweight user interface. Research has previously focussed on designing video media spaces for office situations rather than home-settings

7.2 Thesis Contributions

This thesis presents solutions to each of the problems with the following research contributions:

1. **An evaluation of blur filtration for its effectiveness in balancing privacy and awareness in a home media space.** In Chapter 3, I defined a controlled experiment to evaluate blur filtration for its effectiveness in balancing privacy and awareness for home-based video conferencing where both a telecommuter and others in the home face privacy threats. My experiment tested blur filtration for typical home situations that varied in the amount of perceived privacy risk presented, from little or no risk to very high risk. In Chapter 4, I presented the results of the experiment and clearly showed that for home-based video links blur filtration by itself does not suffice for privacy protection; other privacy-protecting strategies and technologies are required. I also found that as privacy risk increases, people begin to abandon filtration as a strategy for preserving privacy and prefer other privacy-protecting techniques that offer direct control over their privacy, e.g., being able to position the camera, control the blur level, turn the camera on/off. Many researchers (e.g., Zhao and Stasko, 1998, Crowley et al., 2000) have pursued avenues where filtration is used as a technique for balancing privacy and awareness. However, my thesis demonstrates that blur filtration, and by implication other filtration techniques, are not suitable techniques for balancing privacy and awareness in home-based video conferencing because the camera remains facing the user and people simply do not trust this. *(Problem 1)*

2. **An investigation of other privacy-protecting strategies for balancing privacy and awareness in a home media space.** The results of my study on blur filtration highlighted the importance of providing user control over information conveyed through a video media space. To provide natural mechanisms for users to control this information, I took a step back and investigated social-psychological theory to understand how humans regulate privacy in everyday life through various behaviors and actions called privacy mechanisms. In Chapter 5, I present this social-psychological theory and describe how I have used it to develop a framework for the design of a home media space. This framework uses the same categories of privacy

mechanisms that humans in everyday life to present plausible privacy-protecting strategies for balancing privacy and awareness in a home media space. *(Problem 2)*

3. **The rationale and design of a home media space that presents users with privacy-protecting strategies.** In Chapter 6, I present the rationale and prototype design of a home media space, which uses context-aware computing and dedicated physical controls to present users with privacy-protecting strategies outlined in the framework from Chapter 5. The design rationale presents a set of principles for the design of a home media space articulating how user interface techniques can be used to present privacy-protecting strategies. Using a set of home media space elements based on these principles, users are able to control privacy through explicit and implicit actions, and, visual and audio feedback presents the level of privacy currently being maintained. The prototype design has not been formally evaluated and we are not yet sure if the strategies we use are appropriate; however, the design's importance is that it illustrates how to employ user interface design techniques to present users with lightweight strategies for balancing privacy and awareness and how to use the design framework from Chapter 5. *(Problem 3)*

7.3 Future Work

In this thesis, I evaluated and explored techniques for balancing privacy and awareness in home media spaces; however, it is not yet clear if the strategies employed in my prototype home media space are in fact appropriate and suitable for balancing privacy and awareness in an actual home setting where privacy risks are very real. As such, future work within privacy and home media spaces should involve redesigning and formally evaluating a home media space design. First, one should redesign the prototype home media space from Chapter 6 to ensure the techniques presented to users are both natural and lightweight. Second, one should deploy the home media space design to a small number of colleagues who regularly work from home as telecommuters. Third, based on the findings from this deployment, the home media space should again be redesigned. Fourth, the home media space should be deployed within a small group of non-biased participants' homes for an extended period of time (ranging from two to six months).

This type of design and evaluation process would serve several purposes. First, one can ensure the privacy-protecting strategies presented are lightweight and natural for people to use. Second, one could identify what privacy-protecting strategies people *actually* choose to use when they are really faced with the privacy issues that I have identified in this thesis. While the privacy-protecting strategies I present in Chapter 5 and the user interface techniques I use in Chapter 6 are plausible solutions for balancing privacy and awareness, it is unclear if they would prove effective for people in actual practice. Third, given a set of privacy-protecting strategies that people actually use, one can then evaluate them for their effectiveness in balancing privacy and awareness when people are faced with *actual* privacy violations. Fourth, we are still not sure how people would be affected by long term usage of a home media space. It may be the case that home inhabitants adapt to having technology such as always-on video and become less sensitive to privacy issues that they face. It could instead be the case that privacy violations become so frequent that eventually a majority of people become adamantly opposed to using video in their homes.

Naturally, problems will arise when researching the use of a home media space within an actual home setting. Such real-world investigations are very difficult; they are costly, time consuming, and hard to coordinate. As well, privacy investigations in real-world settings have the potential to create additional privacy risks for participants. Many people do not wish to have their privacy compromised even if the investigations will prove to be valuable for the greater good; thus, limiting future privacy violations in home-based video conferencing.

7.4 Conclusion

The issues presented in this thesis are very real as video cameras and media spaces are rapidly moving into everyday use in many settings including homes. This is made easy with the declining cost of PC cameras and several companies are offering free video conferencing software (e.g., Webcam for MSN Messenger, Yahoo! Messenger). While I have concentrated on one specific use of video in homes, the ideas contributed in this thesis have a broader significance for home-based videoconferencing in general. Regardless of the specific use of video in a home, people need and desire methods to

regulate their privacy; existing video conferencing systems ignore these user requirements. It is of vital importance that we as a community understand that privacy risks exist when people use video from home and people desire techniques that are able to preserve their privacy while still allowing them to benefit from the technology. This understanding can then be packaged into a form that manufacturers of home-based video conferencing systems can use to guide future design in order to provide users with techniques to mitigate privacy concerns.

Appendix A. References

1. Altman, I. (1975), **The Environment and Social Behavior: Privacy, Personal Space, Territory, Crowding**, *Wadsworth Publishing Company*, pp. 1-51, 194-207.

2. Altman, I., and Chemers, M. (1980), **Culture and Environment**, *Wadsworth Publishing Company*, pp. 1-12, 75-119, 155-214.

3. Bellotti, V., and Sellen, A. (1993), **Design for Privacy in Ubiquitous Computing Environments**, *Proceedings of the Third European Conference on Computer-Supported Cooperative Work (ECSCW'93),* Kluwer Academic Publishers, Milan, pp. 77-92.

4. Bellotti, V. (1996), **What you don't know can hurt you: Privacy in Collaborative Computing,** *Proceedings of HCI '96,* Springer, pp. 241-261.

5. Bellotti, V. (1998), **Design for Privacy in Multimedia Computing and Communications Environments**, *Technology and Privacy: The New Landscape,* Agre and Rotenberg eds., MIT Press, pp. 63-98.

6. Bellotti, V., Back, M., Edwards, K., Grinter, R., Hnderson, A., and Lopes, C. (2002), **Making Sense of Sensing Systems: Five Questions for Designers and Researchers**, *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2002)* [CHI Letters 4(1)], ACM Press, pp. 415-422.

7. Bly, S., Harrison, S. and Irvin, S. (1993), **Media spaces: Bringing people together in a video, audio, and computing environment,** *Communications of the ACM 36(1),* ACM Press, pp. 28-46.

8. Boyle, M., and Greenberg, S. (2003), **A Lexicon for Privacy in Video Media Spaces,** Report 2003-724-27, *Department of Computer Science,* University of Calgary, May.

9. Boyle, M., Edwards, C. and Greenberg, S. (2000), **The Effects of Filtered Video on Awareness and Privacy,** *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW'00),* [CHI Letters 2(3)], ACM Press.

10. Boyle, M., and Greenberg, S. (2002), **GroupLab Collabrary: A Toolkit for Multimedia Groupware,** J. Patterson (Ed.): *ACM CSCW 2002 Workshop on Network Services for Groupware*, November.

11. Chen, M. (2002), **Leveraging the Asymmetric Sensitivity of Eye Contact for Videoconference,** *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2002)* [CHI Letters 4(1)], ACM Press, pp. 49-56.

12. Coutaz, J., Bérard, F., Carraux, E., Crowley, J. (1998), **Early experience with the mediaspace CoMedi,** *IFIP Working Conference on Engineering for Human-Computer Interaction (EHCI98)*, Heraklion, Crete, Greece.

13. Crowley, J.L., Coutaz, J., and Bérard, F. (2000), **Things That See**, *Communications of the ACM,* ACM Press, Vol. 43, No. 3, pp. 54-64.

14. Dey, A. (2001), **Understanding and Using Context**, *Personal and Ubiquitous Computing*, January 2001, Vol. 5(1).

15. Dourish, P. (1993), **Culture and Control in a Media Space**, *Proceedings of the Third European Conference on Computer-Supported Cooperative Work (ECSCW '93)*, pp. 125-137.

16. Dourish, P., and Bly, S. (1992), **Portholes: Supporting Awareness in a Distributed Work Group**, *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '92),* pp. 541-547.

17. Erickson, T. (2002), **Some problems with the notion of context-aware computing**, *Communications of the ACM,* Vol. 45(2), February 2002, pp. 102-104.

18. Fish, R.S., Kraut, R.E., and Chalfonte, B.L. (1990), **The VideoWindow System in Informal Communications**, *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW'90),* Los Angeles, pp. 1-11.

19. Fish, R., Kraut, R., Root, R., & Rice, R. (1992), **Evaluating video as a technology for informal communication,** *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '92),* New York: ACM Press, pp. 37-48.

20. Fish, R.S., Kraut, R.E., Rice, R.E., and Root, R.W. (1993), **Video as a Technology for Informal Communication**, *Communications of the ACM*, ACM Press, Vol. 36, No. 1, pp. 48-61.

21. Gaver, W. W. (1988), **Everyday listening and auditory icons**, Doctoral Dissertation, University of California, San Diego.

22. Greenberg, S. (1996), **Peepholes: Low Cost Awareness of One's Community**, *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '96), Companion Proceedings*, pp. 206-207.

23. Greenberg, S. and Fitchett, C. (2001), **Phidgets: Easy Development of Physical Interfaces through Physical Widgets,** *Proceedings of the UIST 2001 14th Annual ACM Symposium on User Interface Software and Technology*, ACM Press, pp. 209-218.

24. Greenberg, S. and Kuzuoka, H. (2000), **Using Digital but Physical Surrogates to Mediate Awareness, Communication and Privacy in Media Spaces**, *Personal Technologies*, 4(1), January.

25. Greenberg, S. and Rounding, M. (2001), **The Notification Collage: Posting Information to Public and Personal Displays**, *Proceedings of the ACM Conference on Human Factors in Computing Systems* [CHI Letters 3(1)], pp. 515-521, ACM Press.

26. Gutwin, C., Stark, G., and Greenberg, S. (1995), **Support for Workspace Awareness in Educational Groupware**, *Proceedings of the ACM Conference on Computer Supported Collaborative Learning*, LEA Press, pp. 147-156.

27. Harper, R. (1996), **Why People Do and Don't Wear Active Badges: A Case Study**, *Computer Supported Cooperative Work*, Kluwer Academic Publishers, pp. 297-318.

28. Huang, W., Olson, J., Olson, G. (2002), **Camera Angle Affects Dominance in Video-Mediated Communication,** *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2002),* ACM Press, pp. 716-717.

29. Hudson, S.E., and Smith, I. (1996), **Techniques for Addressing Fundamental Privacy and Disruption Tradeoffs in Awareness Support Systems**, *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW'96),* Cambridge, MA.

30. Jancke, G., Venolia, G.D., Grudin, J., Cadiz, JJ, and Gupta, A. (2001), **Linking Public Spaces: Technical and Social Issues**, *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2001),* Seattle, pp. 530-537.

31. Kraut, R., Egido, C., and Galegher, J. (1988), **Patterns of contact and communication in scientific observation**, *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW '88),* pp. 1-12.

32. Lee, A., Girgensohn, A., Schlueter, K. (1997), **NYNEX Portholes: Initial User Reactions and Redesign Implications**, *Group '97,* ACM Press, pp. 385-394.

33. Mantei, M., Baecker, R., Sellen, A., Buxton, W., Milligan, T., and Wellman, B. (1991), **Experiences in the use of a media space,** *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '91),* New York: ACM Press, pp. 203-209.

34. Nardi, B., Whittaker, S., Bradner, E. (2000), **Interaction and Outeraction: Instant Messaging in Action.** *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW 2000),* pp.79-88.

35. Patterson, A., and Chiswick, N. (1981), **The Role of the Social and Physical Environment in Privacy Maintenance Among the Iban of Borneo**, *Journal of Environmental Psychology,* Vol. 1, pp. 131-139.

36. Pirhonen, A., Brewster, S., and Holguin, C. (2002), **Gestural and Audio Metaphors as a Means of Control for Mobile Devices**, *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2002)* [CHI Letters 4(1)], ACM Press, pp. 291-298.

37. Schilit, B., Adams, N., Want, R. (1994), **Context-aware computing applications,** *Proceedings of the Workshop on Mobile Computing Systems and Applications*, Santa Cruz, CA, December, IEEE Computer Society, pp. 85-90.

38. Schilit, B., and Themier, M. (1994), **Disseminating Active Map Information to Mobile Hosts**, *IEEE Network 8(5)*, pp. 22-32.

39. Tang, J., Isaacs, E., and Rua, M. (1994), **Supporting Distributed Groups with a Montage of Lightweight Interactions,** *Proceedings of the ACM Conference on Computer-Supported Cooperative Work (CSCW '94)*, ACM Press, pp. 23-34.

40. Want, R., Hopper, A., Falcão, V., and Gibbons, J. (1992), **The Active Badge Location System**, *ACM Transactions on Information Systems*, Vol. 10, No. 1, January 1992, ACM Press, pp. 91-102.

41. Whittaker, S., Frohlich, D., and Daly-Jones, O. (1994), **Informal workplace communication: What is it like and how might we support it?** *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '94)*, Boston, pp. 131-137.

42. Whittaker, S., and O'Conaill, B. (1997), **The Role of Vision in Face-to-Face and Mediated Communication**, *Video Mediated Communication*, Finn, Sellen, and Wilbur eds., Lawrence Erlbaum Associates Inc., pp. 23-50.

43. Zhao, Q.A., and Stasko, J.T. (1998), **Evaluating Image Filtering Based Techniques in Media Space Applications,** *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW'98)*, Seattle, pp. 11-18.

Appendix B. Pilot Study

This appendix discusses a pilot study, which was an initial investigation into the effectiveness of blur filtration for balancing privacy and awareness for home situations. The pilot study was carried out prior to the controlled experiment described in Chapter 3 and aided in the experiment's design. First, I briefly describe the study's methodology, including the materials and questionnaires. Second, I summarize the important results of the study and the issues found with the study's design.

B.1 Methodology

The goal of the pilot study was to evaluate blur filtration for its effectiveness in balancing privacy and awareness for typical home situations, which may be captured in a home media space. In the study, participants role-played either a *telecommuter* (trying to preserve his/her own privacy) or a *viewing colleague* (trying to gain awareness of the telecommuter) during the study. Those imagining themselves as a telecommuter rated the privacy threat of each scene and chose a blur level that made them comfortable in having their colleague see the scene. Those imagining themselves as a viewing colleague saw video scenes at each of eleven blur levels (one level was unfiltered) and had to identify how available the telecommuter was.

B.1.1 Materials: Video Scenes

To decide which video scenes to record and eventually show to study participants, an initial informal survey was created and deployed to colleagues and friends through email and an instant messenger. The survey contained a list of home scenarios ranging in the appearance of someone in the home, the same person's activity, and the appearance of the same person's location. Colleagues and friends were asked to rate how much they felt each scenario would threaten their privacy if they were the person at home being viewed by a co-worker.

Based on the findings of this informal survey, a total of 70 video scenes were recorded, 35 containing a paid male actor and 35 containing a paid female actress. The actor and actress were both over 18 years of age. All scenes were recorded in various rooms within a home using a high-quality Canon XL-1 digital video (DV) camera. Video scenes contained one of 16 different activities and also varied in the appearance of the actor or actress (clean and casual clothes *vs.* slightly unkept or dirty/wrinkled clothes) and the appearance of the location filmed (clean *vs.* messy). Each recorded sequence was approximately one minute in length. After recording, all 70 video scenes were edited into approximately 10 to 20 second video clips at the same quality (720 x 480 pixels) and frame rate (30 frames per second) as recorded with the DV camera.

After reflecting upon the initial survey's findings, a set of ten video scene pairs were chosen, which appeared to represent a broad range of privacy threats from no threat to very high threat. One video in each pair contained the male actor, while the other contained the female actress. The twenty scenes are described in Figure B.1 (female scenes) and Figure B.2 (male scenes). Each scene was also pre-processed at eleven


1: Dancing


2: Drinking alcohol

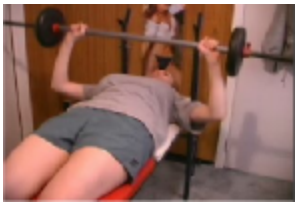
3: Lifting weights


4: Picking nose

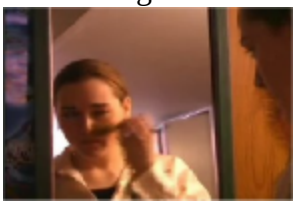
5: Putting on make-up


6: Reading a book


7: Scratching rear


8: Sorting laundry

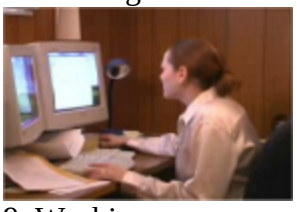
9: Working

Figure B.1: Female video scenes shown in the study. The tenth scene showing the actress changing clothes (and in underwear) is not shown.

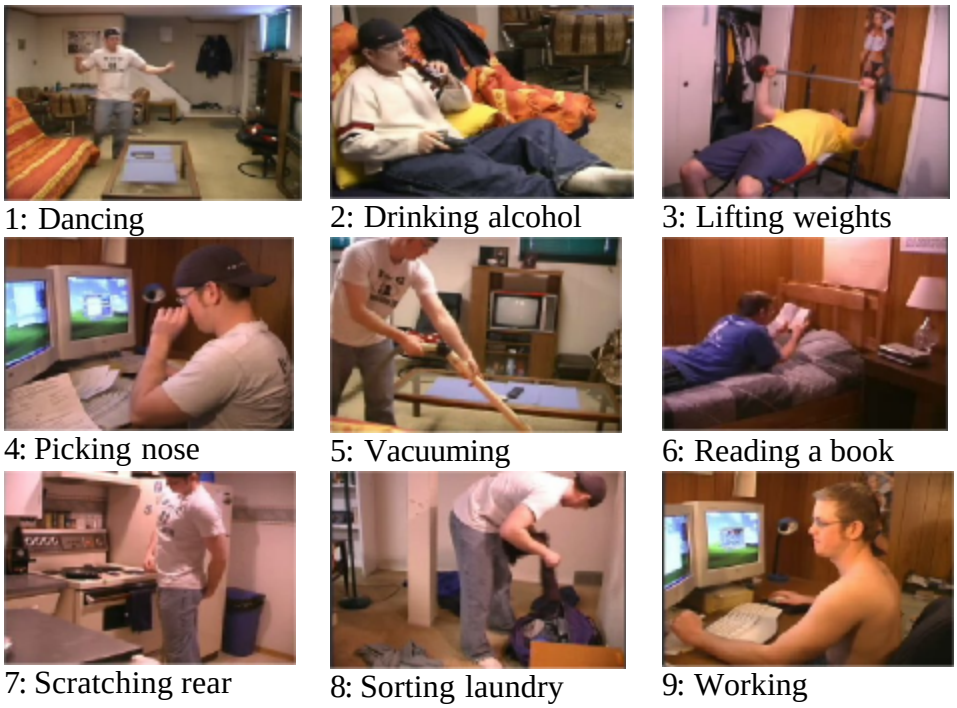

Figure B.2: Male video scenes shown in the study. The tenth scene showing the actor changing clothes (and in underwear) is not shown.

different levels of blur filtration. The range of blur levels went from level 10—completely blurred to level 0—unfiltered.

B.1.2 Materials: Questionnaires

Four questionnaires were used for the study: a pre-test questionnaire, two during-test questionnaires, and a post-test questionnaire. The pre-test questionnaire gathered demographics and asked about experience using video conferencing software. The first during-test questionnaire was used for the "telecommuter" group and asked participants to:

- describe the co-worker who they imagined was watching them in a video, e.g., co-worker's gender, co-worker's relationship to the participant
- rate the appropriateness of the person's activity, the person's appearance, and the location's appearance to be viewed by the same co-worker at the office

- rate how comfortable they would feel to have the previously described co-worker see the person's activity, the person's appearance, and the location's appearance if the participant was the person in the video
- rate how much the video would threaten their privacy if they were the person in the video
- choose a blur level that would make them feel comfortable to let the same co-worker see the video
- describe what aspects of the video they were trying to mask by blurring the video

The second during-test questionnaire was used for the "viewing colleague" group and asked participants to:

- describe the activity, position, and appearance of the person in the video
- rate the availability of the person in the video
- describe the location in the video, e.g., type of room, room's appearance
- describe visible objects in the video
- rate the appropriateness of the person's activity, the person's appearance, and the location's appearance
- rate the effect of the current blur level (not blurred enough *vs*. blurred unnecessarily too much)
- describe aspects of the video they felt should be blurred more

The "viewing colleague" group was also asked to rate their confidence in their ability to accurately describe each aspect of the scene.

A post-test questionnaire gathered each participant's opinion of blurring video to preserve privacy 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.

B.2 Participants

The pilot study was conducted with eleven participants (6 females, 5 males), containing both graduate and undergraduate students from the University of Calgary ranging in age from 19 to 34 years old. Ten of the participants were performing studies in Computer

Science, while one participant was performing studies in another area of the Sciences. All participants were recruited through email and were offered a donut or cookie for their participation in the study. Once recruited, participants were first randomly split into one of two groups: "telecommuters" or "viewing colleagues."

B.3 Method

Participants were given a scenario where they had a need and desire to work with another colleague very closely. The "telecommuter" group was told they had recently started working from home several days a week and wanted to stay in contact with a close-working colleague at the office. The "viewing colleague" group was told that their close-working colleague had recently become a telecommuter and they wanted to maintain contact with this person. Both groups were then told that a video link would be used to provide awareness between colleagues and may occasionally capture non-work related activities at the home. The "telecommuter" group then:

1. viewed one of the ten video scenes at random
2. answered questions about the scene (during-test questionnaire)
3. chose a blur level for the scene and gave a reasoning
4. repeated steps 1-3 for each of the remaining nine scenes

The "viewing colleague" group:

1. viewed one of the video scenes at random, starting with a completely blurred video
2. answered questions about the blurred scene (during-test questionnaire)
3. repeated steps 1 and 2 for the same scene at each of the remaining blur levels
4. repeated steps 1-3 for two more video scenes (chosen at random)

Participants saw only scenes containing an actor/actress from the same sex. Both groups of participants completed the study by answering the post-test questionnaire.

B.4 Results

The study was only designed as a pilot study, therefore this section discusses only a portion of the results. No statistical analysis is performed.

B.4.1 Privacy Threat

When asked how threatening each scene was to the telecommuter's privacy, responses from the male "telecommuter" participants created the following privacy threat groupings for scenes:

- ***No Threat***: Vacuuming
- ***Low Threat***: Dancing, Drinking alcohol, Lifting weights, Reading
- ***Moderate Threat***: Picking Nose, Scratching rear, Sorting laundry
- ***High Threat***: Changing, Working with no shirt

Responses from female "telecommuter" participants created the following privacy threat groupings for scenes:

- ***No Threat***: Picking Nose, Working
- ***Low Threat***: Dancing, Lifting weights, Putting on make-up, Reading
- ***Moderate Threat***: Drinking alcohol, Sorting laundry
- ***High Threat***: Changing, Scratching rear

When asked what made each scene threatening, the descriptions given by both male and female participants fell into one of three categories: the activity, the person's appearance, or the bcation's appearance. Participants noted that some aspects were embarrassing such as exercising, scratching oneself, wearing only underwear, or having posters of women wearing only a bikini visible. A total of 65 problems were identified by participants for all the video sequences. The location's appearance was identified as being a problem 16 times, the person's appearance was identified 20 times, and the person's activity was mentioned 29 times.

B.4.2 Identifying Awareness Cues

I wanted to know at what blur level the "viewing colleague" participants were confident in their ability to describe the scene and identify various awareness cues. Only the confidence is being analyzed because many aspects of the scene descriptions are open to interpretation; for example, one person may view a room as being messy while another may feel it is clean. As well, I am interested in what people perceive they are seeing, meaning if someone is confident they saw (say) a person undressing in a blurred video, then they will indeed think they saw that person undressing whether they actually did or not.

The following is a list of awareness items and the corresponding mean blur level over all scenes at which participants were confident in their ability to correctly identify the cue. The list is sorted by blur level so awareness cues identified first are at the top of the list:

- ***activity of the person:*** mean=6, *s.d.*=1.5
- ***availability of the person***: mean=5.5, *s.d.*=2
- ***appearance of the person***: mean=5, *s.d.*=2
- ***objects***: mean=5, *s.d.*=1.5
- ***location***: mean=4, *s.d.*=1.5

These means show that generally all awareness cues are identifiable around the same blur level, between levels 4 and 6.

B.4.3 Post-Test Questionnaire

When asked what they liked or disliked about blurring video to help preserve privacy many noted that they liked the idea of hiding details in the video, but some remarked that it was difficult to gain awareness at highly blurred levels . When asked if they would use an open video link in an office, 8 out of 11 said they would and 3 out of 11 said they would not. Of those that answered yes, most said they would because their activities would be appropriate for others to see and the video would provide an awareness of others. Those that answered no generally felt that video was intrusive. When asked if

they would use an open video link at home, 7 out of 11 people said they would, but they required control over where the recording was taking place, what was being recorded, and how much the video was being blurred. Several said they would prefer to have the video record in a home office dedicated to work.

B.5 Discussion

Several problems and insights were found when designing and running the pilot study. First, not all pairings of scenes between males and females were the same as would be the ideal situation. Some scenes that contain the same activity do not contain the same background scenario, such as the Drinking Alcohol scene. Both portray the telecommuter drinking beer while watching TV, but the room is messy for the female and clean for the male. Another scene contains a female applying make-up to her face, yet there is no male equivalent of the scene for obvious reasons; the male scene instead contains the actor vacuuming. Other scenes that seem to be quite similar like the two Sorting Laundry scenes are actually slightly different. Both scenes show sorting laundry as the activity; however the male's version of the scene contains him smelling clothes while sorting them and the female's version does not. While it is impossible to have identical male and female scenes, for a full study more attention needs to be paid to ensure each scene pair is as similar as possible.

Second, during the planning of the pilot study, it was hypothesized that problems in using video within a home setting would fall into one of three categories: the person's activity, the appearance of the person, and the appearance of the location (which included the type of room). By looking at what participants said was threatening in each scene, it is clear that this was the case, for all descriptions fell into one of the three categories. The person's activity was the problem most mentioned, followed by the person's appearance, and then the location's appearance. The ordering suggests that it may be desirable to have the location's appearance as a controlled factor in future studies for it seems to contribute less to privacy threat. Thus, the privacy risk for scenes should be determined by the person in the scene, rather than the location.

Third, the pilot study contained too many scenes that may be considered irrelevant to telecommuting scenarios. The scenes contain a very broad range of situations, many of which would likely not be captured on camera. This mostly stems from recording scenes throughout various rooms in a home. For example, it is unlikely that someone would place a camera in a living room or a kitchen, and activities such as vacuuming aren't likely to be captured. As well, we can predict that the camera would be turned off for some other scenes such as lifting weights or reading a book in bed. As mentioned, participants did suggest they felt the most appropriate place for a camera would be in a spare bedroom containing a home office. Many of the scenes used in the pilot study would then have a very low chance of being captured in actual practice. It is clear that for a future study, scenes should better typify home telecommuting situations that may occur in a home office/spare bedroom.

B.6 Conclusion

In this appendix, I have presented the design of a pilot study aimed at evaluating blur filtration for balancing privacy and awareness. Several problems and issues were found in the design of the pilot study; thus, future study designs should:

1. ensure scenes are as similar as possible between males and females;
2. control the background information in the scenes, e.g., use a typical home office/spare bedroom as each scene's location; and,
3. use scenes that *actually* represent situations facing telecommuters who work from home and use video for awareness.

The problems identified with the pilot study's design are used to aid in the design of the controlled experiment described in Chapter 3.

Appendix C. Experiment Materials

C.1 Protocol for the Experiment

Before you begin, you should have a pre-test questionnaire ready for the subject and the software up and running. No other windows should be visible.

Introduce yourself.

- My name is _____, and I will be giving you instructions on what to do and will answer your questions.

- We're researching privacy issues faced by telecommuters when using video to connect to the office. You're helping us by evaluating a technique that helps to preserve privacy and allowing us to understand what perceived privacy issues occur in a home setting when using video.

Tell them about the experiment.

- During the study, you will be shown a number of video clips. The clips portray one or more actors/actresses at home doing typical home activities in a home office/spare bedroom. While none of the video clips will be intended to offend you, they will be designed to portray the actors/actresses in situations that could be considered threatening to their privacy. None of the scenes will include sexually explicit material, but may include nudity. Actors and actresses will be persons 18 years of age or older.

- After seeing the video clips, you will be asked privacy and awareness related questions about them.

Tell the participant that it's OK to quit at any time.

- If you feel uncomfortable you are free to quit at any time without repercussions.

- Do you have any questions at this point?

Give them the consent form to sign. If it is not signed, do not proceed.
Record the subject ID. Hand the participant a pre-test questionnaire.

- Before we begin the study, I would like you to answer a few questions found on this form. They will give us some background information about your computer experience and privacy expectations.

Participant should answer the pre-test questionnaire.

- You will now begin the study where you will see various video clips and answer questions about them.

- I will be sitting here and will be available if you have any questions. I may also be recording various observations during the study on paper.

Participant should answer the during-test questionnaire.
Participant should answer the post-test questionnaire – supply them with the pictures of each scene to force sort by risk.

C.2 Consent Form

Research Project Title: Mediating Privacy in the Home

Investigator: Carman Neustaedter

This consent form, a copy of which has been given to you, is only part of the process of informed consent. It should give you the basic idea of what the research is about and what your participation will involve. If you would like more detail about something mentioned here, or information not included here, you should feel free to ask. Please take the time to read this carefully and to understand any accompanying information.

Experiment Purpose:
The purpose of this research is to understand privacy issues faced by telecommuters when using video to connect to the office. We will also be evaluating a technique that blurs video to help preserve privacy.

Participant Recruitment and Selection:
To be a recruited for this study, we ask that you allow us to use and analyze your results from the study.

Procedure:
The study should require no more than 1.5 hours of your time. You will be asked to perform the following activities:

1. View a number of video scenes portraying actors/actress performing various activities within a home.
2. You will be asked privacy and awareness questions about the scenes.
3. For portions of the study, you will also be asked to imagine yourself as the person in the video.

Confidentiality:
Your anonymity will be strictly maintained. Reports and presentations will refer only to a participant identification number and will be in a secure filing cabinet or on a secure computer.

Risks:
The only risk is the possibility that the video scenarios may cause you to feel uncomfortable or embarrassed. You are free to quit at any time without repercussions. All information collected from a person that withdraws will be destroyed.

Investigators:
Carman Neustaedter is a M.Sc. student in the Department of Computer Science at the University of Calgary under the supervision of Dr. Saul Greenberg, Professor in the Department of Computer Science.

Your signature on this form indicates that you have understood to your satisfaction the information regarding participation in the research project and agree to participate as a subject. In no way does this waive your legal rights nor release the investigators, sponsors, or involved institutions from their legal and professional responsibilities. You are free to withdraw from the study at any time. Your continued participation should be as informed as your initial consent, so you should feel free to ask for clarification or new information throughout your participation. If you have further questions concerning matters related to this research, please contact:

Carman Neustaedter (carman@cpsc.ucalgary.ca)

If you have any questions or issues concerning this project that are not related to the specifics of the research, you may also contact the Research Services Office at 220-3782 and ask for Mrs. Patricia Evans.

Participant's Signature	Date
Investigator's/Witness's Signature	Date

A copy of this consent form has been given to you to keep for your records and reference.

C.3 Pre-Test Questionnaire

Age: _____ Gender: _____
Occupation: _____

1. How often do you use a computer?

1	2	3	4	5
Never used	Use monthly	Use at least once a week	Use several time a week	Use daily

2. Have you ever been a telecommuter (work from home and connect to an office with some form of technology, e.g., phone, email, instant messaging)?

1	2	3
Never	Previously	Currently

If you have been a telecommuter, how frequently did you/do you telecommute?

1	2	3	4	5
Once a month	A few times a month	Once a week	Several days a week	Always

3. How often do you use video conferencing where others see your video image?

1	2	3	4	5
Never used	Used once	Used 2-5 times	Used many times	Use daily

4. How would you describe your personality?

1	2	3	4	5
Very Shy	Shy	Neither shy nor outgoing	Outgoing	Very Outgoing

5. How concerned are you about a co-worker's perception of the following (1-Not concerned to 5-Very concerned):

	Not Concerned				Very Concerned
a) You in general:	1	2	3	4	5
b) Your work performance:	1	2	3	4	5
c) Your actions at work:	1	2	3	4	5
d) Your physical appearance at work:	1	2	3	4	5
e) Your social status:	1	2	3	4	5

C.4 Post-Test Questionnaire

1. If you were a telecommuter and wanted to stay in close contact with a colleague at the office, how willing would you be to let him/her see:

a.) an ***unblurred*** video image of you at home while you are working (remembering the video may accidentally capture you doing other things)?

1	2	3	4	5
Not willing	Slightly Willing	Moderately Willing	Willing	Very Willing

b.) a ***blurred*** video image of you at home while you are working (remembering the video may accidentally capture you doing other things)?

1	2	3	4	5
Not willing	Slightly Willing	Moderately Willing	Willing	Very Willing

2. What did you **like** or **dislike** about blurring the video to preserve privacy while still trying to stay in contact with your close colleague*?*

_____

_____

_____

3. Would you use an open video link in **an office** to stay in contact with a close colleague if the video was blurred? Why or why not?

_____

_____

4. Would you use an open video link in **your home** to stay in contact with a close colleague if the video was blurred? Why or why not?

_____

_____

Appendix D. Ethics Approval

UNIVERSITY OF CALGARY

CERTIFICATION OF INSTITUTIONAL ETHICS REVIEW

This is to certify that the Conjoint Faculties Research Ethics Board at the University of Calgary has examined the following research proposal and found the proposed research involving human subjects to be in accordance with University of Calgary Guidelines and the Tri-Council Policy Statement on *Ethical Conduct in Research Using Human Subjects*:

Applicant(s): Carman G. Neustaedter / Michael J. Boyes

Department/Faculty: Department of Computer Science, Faculty of Science

Project Title: Mediating Privacy in the Home

Sponsor (if applicable):

Restrictions:

This Certification is subject to the following conditions:

1. **Approval is granted only for the project and purposes described in the application.**
2. **Any modifications to the authorized protocol must be submitted to the Chair, Conjoint Faculties Research Ethics Board for approval**
3. **A progress report must be submitted 12 months from the date of this Certification, and should provide the expected completion date for the project.**
4. **Written notification must be sent to the Board when the project is complete or terminated**

Chair
Conjoint Faculties Research Ethics Board

Date: 5 March 2002

Distribution: (1) Applicant, (2) Supervisor (if applicable), (3) Chair, Department/Faculty Research Ethics Committee, (4) Sponsor, (5) Conjoint Faculties Research Ethics Board (6) Research Services

09/00

UNIVERSITY OF CALGARY

CONJOINT FACULTIES RESEARCH ETHICS BOARD

c/o Research Services
Telephone: (403) 220-3782
Fax: (403) 289-0693
Email: plevans@ucalgary.ca

To: Ms. Carmen Neustaedter
Department of Computer Science

Date: June 10, 2002

From: Dr. Janice P. Dickin, Chair
Conjoint Faculties Research Ethics Board (CFREB)

Re: **Ethics Proposal Modification: Mediating Privacy in the Home**

The Certificate of Ethical Approval issued on March 5, 2002 continues in force and extends to the modifications as set out in your email request for approval, dated June 10, 2002. Your request, as detailed, is approved. You should attach a copy of the documentation you provided in order to request this amendment, together with a copy of this memorandum, to the original Ethics Certification in your files.

Sincerely,

Janice Dickin, Ph.D., LLB., Professor
Faculty of Communication and Culture and
Chair, Conjoint Faculties Research Ethics Committee

cc: Dr. S. Greenberg, Supervisor
Chair, Department/Faculty Committee

Appendix E. Co-Author Permission

UNIVERSITY OF CALGARY

May 1, 2003

University of Calgary
2500 University Drive NW
Calgary, Alberta
T2N 1N4

I, Saul Greenberg, give Carman Neustaedter permission to use co-authored work from our papers, "Balancing Privacy and Awareness for Telecommuters Using Blur Filtration" and "The Design of a Context-Aware Home Media Space for Balancing Privacy and Awareness," for Chapters 3, 4, and 6 of his thesis and to have this work microfilmed.

Sincerely,

Saul Greenberg

UNIVERSITY OF CALGARY

May 1, 2003

University of Calgary
2500 University Drive NW
Calgary, Alberta
T2N 1N4

I, Michael Boyle, give Carman Neustaedter permission to use co-authored work from our paper, "Balancing Privacy and Awareness for Telecommuters Using Blur Filtration," for Chapters 3 and 4 of his thesis and to have this work microfilmed.

Sincerely,

Michael Boyle

Michael Boyle