Gutwin, C. (1997). **Workspace Awareness in Real-Time Distributed Groupware.** Ph.D. Thesis, Department of Computer Science, University of Calgary, Calgary, Alberta, Canada. December. Extract: Chapter 1.

Chapter 1 Introduction

Real-time distributed groupware systems are computer applications that allow people to work together at the same time, but from different places. These systems are becoming more common as network connectivity increases and organizations move towards mobile computing, telecommuting, and distributed work teams. Unfortunately, groupware has usability problems: when compared with working face-to-face, collaboration through groupware seems clumsy, stilted, and artificial. When people use a groupware system, they often cannot interact in the ways that they do in front of a whiteboard or over a table. Even simple things like pointing to an object or watching another person work are often difficult or impossible in groupware.

One of the problems with current groupware systems is that they make it hard for people to stay aware of one another. It is a simple fact that awareness is an important part of collaboration. One has only to watch a group activity to see people checking up, taking notice, and keeping track of others, whether over a jigsaw puzzle, in front of a chalkboard, or on a basketball court. Awareness is taken for granted in everyday face-to-face environments, but when the setting changes to distributed groupware, many of the normal cues and information sources that people use to maintain awareness are gone. Groupware systems provide only a fraction of the information about other people that is available in a face-to-face situation, and as a result, interactions and behaviours that were once effortless can become stilted and formal.

I believe that helping people maintain awareness of one another can improve the usability of groupware. In this research, I explore one kind of awareness called *workspace awareness*, and investigate techniques for supporting it in groupware interfaces. My perspectives are those of the computer scientist and the system designer: I am interested in the problem of groupware usability, and awareness is a potential design requirement that can improve multi-user systems. My goals are to show that the concept of workspace awareness can be framed in a

sufficiently operational form to be useful to groupware designers, and to show that support techniques drawn from that framework can improve groupware usability in a variety of ways.

The remainder of this chapter will provide some background on the computer science context into which the research fits, introduce the idea of workspace awareness, state the problem and goals of the research in more detail, and outline what is to come in the rest of the dissertation.

1.1 Setting the scene

This research falls within the field of computer-supported cooperative work (CSCW), an area in the study of human-computer interaction that attempts to understand and provide technological support for group activity (e.g. Baecker 1993; Greenberg 1991). CSCW is a wide-ranging field, containing research into the behavioural foundations of group activity (e.g. McGrath 1984; Galegher, Kraut, and Egido 1990), group interaction in natural settings (e.g. Tang 1991; Heath and Luff 1992), asynchronous communication and electronic mail (e.g. Malone et al. 1987; Borenstein and Thyberg 1988), synchronous communication through audio and video (e.g. Harrison and Minneman 1990; Abel 1990), support for electronic meetings (e.g. Pinsonneault and Kraemer 1990), and multi-user interfaces (e.g. Stefik et al. 1987a).

One area of CSCW research concentrates on real-time distributed groupware: multi-user computer systems that allow people to work together at the same time but from different places (e.g. Sarin and Greif 1985; Dewan and Choudhary 1991; Roseman and Greenberg 1996). This kind of groupware includes applications like shared editors (e.g. Leland et al 1988), drawing programs (e.g. Greenberg et al 1992), multiplayer games (e.g. Rohall et al 1992), and distributed control systems (e.g. Hughes et al 1994). Many real-time groupware systems provide an environment for collaboration called a *shared workspace*, a bounded space where people can see and manipulate artifacts related to their activities[1]. In the physical world, shared workspaces include chalkboards, tabletops, and control panels; the artifacts in

[1] I assume throughout this research that shared-workspace systems are used in conjunction with audio and perhaps video links that provide basic communication facilities.

them might be objects in a design drawing, paragraphs in a document, or gauges and dials on a control panel.

Groupware systems create virtual shared workspaces that are often patterned on these physical spaces. For example, Figure 1 and Figure 2 show simple drawing applications used by two people. The shared workspace here is akin to a large sheet of paper. As each person draws, their actions are communicated to the other machine, so both participants' workspaces contain the same objects. When the whole sheet of paper can fit on one screen, both participants see exactly the same thing at the same time (Figure 1). However, when the workspace is bigger than the computer screen (Figure 2), the groupware system can only show a part of the paper at once, and the participants have to scroll around to see and work on other areas. If both people scroll to different parts of the workspace, as in Figure 2, they

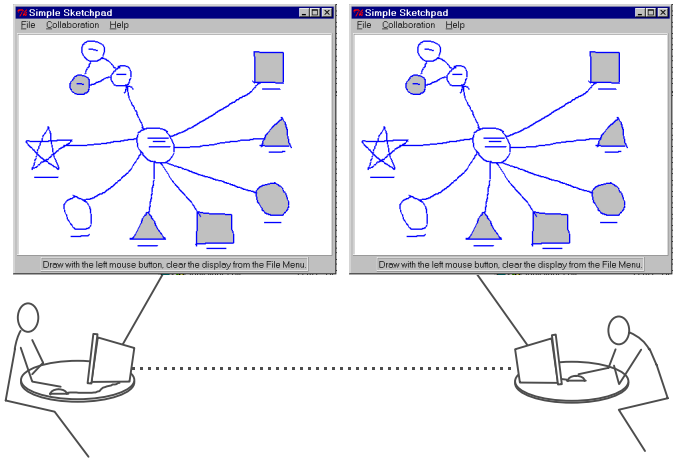

Figure 1. A groupware system where the workspace fits on one screen

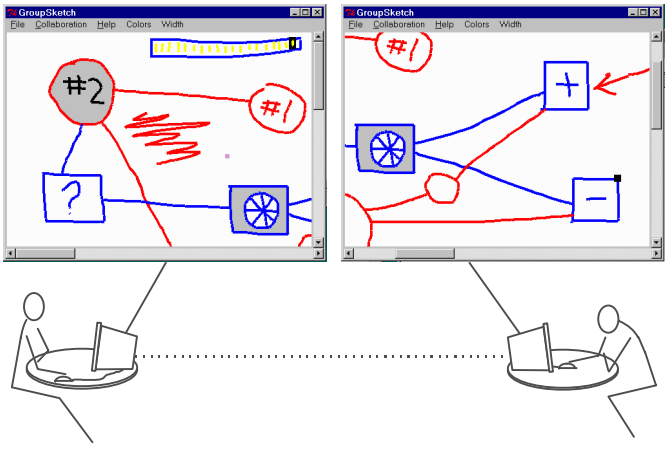

Figure 2. A system where the workspace is larger than the screen

cannot tell where the other person is working or what they are doing.

The question of whether real-time distributed groupware systems are usable or not is itself a recent one. Although shared workspace were first demonstrated in 1968 (Englebart 1968), they received little attention until about 1985. In the early years, researchers had to be concerned about basic technical issues like system architectures (e.g. Lauwers and Lantz 1990), message-passing mechanisms (e.g. Sarin and Greif 1985), and ways of tracking participants as they entered and left groupware conferences (e.g. Roseman and Greenberg 1992). Recently, though, tools such as GroupKit (Roseman and Greenberg 1995) have appeared that simplify the construction of real-time groupware by handling the low-level details of putting together a distributed system. Although technical issues still exist (e.g. Greenberg and Marwood 1994; O'Grady 1996), the tools let researchers focus more closely on the human factors of system design, which includes groupware usability.

In this research, I examine a kind of awareness that is specific to shared workspaces, and show how the maintenance of awareness can be supported in groupware. *Workspace awareness* (WA) is the up-to-the-moment understanding of another person's interaction with a shared workspace (Gutwin and Greenberg 1996a). Workspace awareness involves knowledge about where someone is working, what they are doing, and what they are going to do next. This information is useful for many of the activities of collaboration—for coordinating action, managing coupling, talking about the task, anticipating others' actions, and finding opportunities to assist one another. When people are able to maintain awareness of one another, these activities are more natural, spontaneous, and unforced. If these attributes can be brought to groupware, the quality and productivity of distributed collaboration can be greatly improved.

In order to keep this research effort manageable, and to narrow the focus of the investigation, I will constrain three aspects of the situation under study. First, I will restrict my conception of a workspace to medium-sized flat surfaces like chalkboards and tabletops. Second, since these kinds of spaces are used primarily by small groups, I will presume groups of between two and five people. Third, I will for the most part assume that groups engage in mixed-focus collaboration, where people move back and forth between individual and shared activities during a work session. These restrictions rule out certain kinds of activity,

such as large formal meetings and team sports, but still leave a rich variety of small-group collaboration. Typical examples could include two people organizing slides on a light table, a research group generating ideas on a whiteboard, or the managers of a project planning a timeline of project tasks. Figure 3 shows the context and focus of the research.

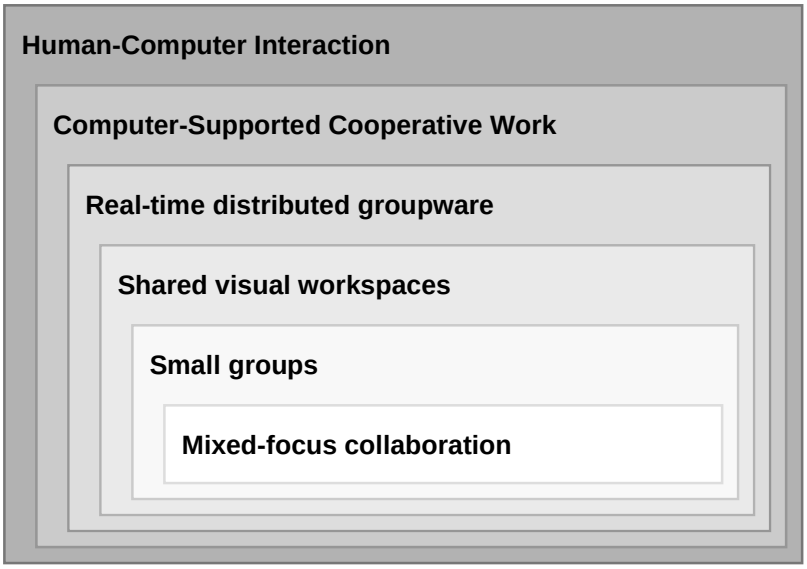

Figure 3. Research context.

1.2 Problem statement and research hypothesis

The problem addressed by this research is that current real-time distributed groupware systems are awkward and cumbersome in comparison with face-to-face shared workspaces. These usability problems are caused in part by the difficulty people have in maintaining workspace awareness. Workspace awareness is an important part of collaborative activity, whether the setting is a physical or a virtual workspace, and if workspace awareness is difficult to maintain, collaboration becomes more difficult.

Furthermore, difficulties in maintaining workspace awareness can be traced to the paucity of information that groupware systems provide about other people. Workspace awareness is maintained through a rich variety of sensory cues, but current groupware systems provide only a fraction of this information; they lack many of the natural affordances that exist in face-to-face settings, and artificial support is limited in current systems. As Liam Bannon has

said, "we find that most of the available facilities do not provide a very rich information space—especially if the focus is on 'real time' facilities" (in Robinson 1991, p. 41).

From this situation, I derive my research hypothesis: that *support for the maintenance of workspace awareness improves the usability of real-time distributed groupware systems.*

I will investigate and test this hypothesis through three research activities: I will operationalize the concept of workspace awareness, apply the concept to the design of groupware interfaces, and evaluate the usability of the resulting systems. Although the evaluation will be the eventual test of the hypothesis, the first two stages are required to connect the experimental results back to the concept of workspace awareness in a logical and traceable way. These three activities form the research process shown in Figure 4, and also define the three research objectives described below.

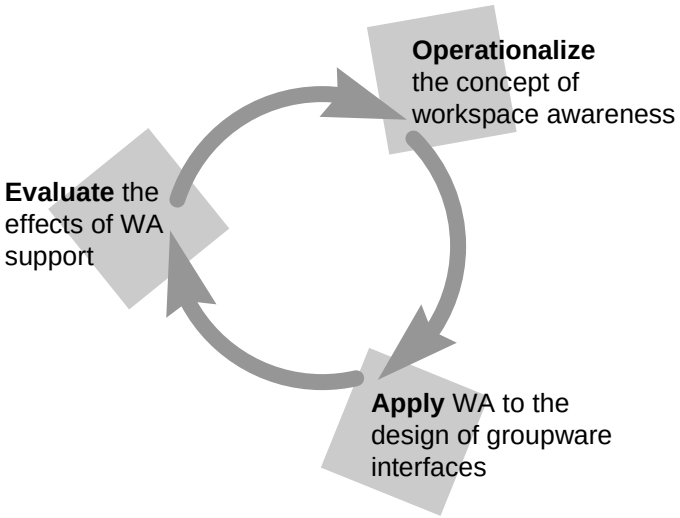

Figure 4. The research process.

Objective 1. I will frame the concept of workspace awareness in operational terms that are useful in groupware design. This objective will be met in three ways: by constructing a conceptual framework that synthesizes knowledge about workspace awareness, by identifying areas where groupware systems require awareness support, and by detailing the steps that designers must undertake to provide that support. My operationalization of the

concept will be successful if it can organize the design space of awareness support and can explain evaluation results in terms of workspace awareness.

Objective 2. I will show that the concept of workspace awareness can be applied in the design of groupware interfaces. This objective will be met by designing and building several displays that provide workspace awareness information, and by collecting additional techniques from existing research. The objective will be complete when I can demonstrate techniques in each area where groupware requires awareness support, and when I can organize the displays in terms of the conceptual framework.

Objective 3. I will show that adding awareness information to a groupware system can significantly improve aspects of groupware usability. This objective will be met by carrying out an experimental comparison of people's collaboration in two groupware systems that provide different levels of support for workspace awareness. I will consider the objective to have been successfully met if the experiment is reliable, valid, generalizable, and shows statistically significant results.

1.3 Contributions of the research

This research contributes original ideas, knowledge, and practices to CSCW and HCI. There are three major contributions.

1. I identify and define workspace awareness as a design requirement for real-time distributed groupware. CSCW research has previously recognized that awareness is important, but the nature and mechanics of workspace awareness have not been articulated before.

2. I construct operational descriptions of workspace awareness that can be used in designing and building groupware systems. I construct a conceptual framework that identifies elements of workspace awareness, mechanisms by which it is maintained, and its uses in collaboration. I also identify primary areas where groupware systems need to provide awareness support, and set out the steps that designers must take to provide it. These operational descriptions will help designers to organize the problem of awareness support, assess work situations, and compare designs and interface techniques.

3. I show that groupware systems can be made significantly more usable by support for workspace awareness. This experimental evidence puts intuitions about the importance of workspace awareness on an empirical footing, and increases knowledge about when and where awareness support will be effective.

There are several other minor contributions. First, I design and build a set of novel techniques for presenting workspace awareness information, several of which have not been seen before in groupware systems. Second, I increase understanding of the definition and measurement of groupware usability. Third, I gather experimental evidence that replicates and reinforces earlier observations about the role of shared workspace in collaborative interaction. Fourth, I investigate the tension between designing for individuals and designing for groups. These major and minor contributions will be discussed further in later chapters.

1.4 Overview of the dissertation

The remainder of the dissertation is organized into parts which follow the research process of Figure 4. To set the scene, Chapter 2 provides background on the two fundamental elements of this research—awareness, and shared workspaces. I consider basic questions of what awareness is and how it works, describe other kinds of awareness studied in previous research, and explore the affordances and constraints of medium-sized workspaces that affect the creation and maintenance of awareness.

Once this foundation is laid, the first part of the research process is operationalization. Chapter 3 focuses on the concept of workspace awareness, and introduces the conceptual framework that will be used in the rest of the dissertation. The conceptual framework explores three high-level questions that designers need to address in supporting awareness in groupware: what information do people need about each other in a shared space, how do people gather and maintain this information, and what do people use the information for? The framework answers these questions by describing three aspects of workspace awareness in operational terms:

- the elements of knowledge that make up workspace awareness,
- the process of maintaining WA and the mechanisms that people use to maintain it, and
- the uses of workspace awareness in mixed-focus, small-group collaboration.

Chapter 4 then explores specific issues of supporting workspace awareness in groupware. I first identify areas where current groupware systems hinder the maintenance of workspace awareness: they reduce perception of the workspace, they reduce the expressiveness of bodies in the workspace, and they reduce the amount of information that actions and artifacts can convey. I then detail the steps that a designer must take to provide awareness support. I consider issues in collecting workspace awareness information in a groupware setting, distributing it to others, and displaying it in the groupware interface.

The second part of the research process is application. Chapter 5 presents a number of display techniques for presenting workspace awareness information in the groupware interface. I demonstrate original and existing displays that address each of the areas in which groupware limits the maintenance of awareness. I describe techniques for workspace embodiment, methods for making actions and artifacts more expressive, and techniques for improving visibility of the workspace. I organize the design space in terms of what awareness information is presented, where in the interface it is presented, and how it is presented. Chapter 6 then looks more closely at one class of awareness displays called the radar view, a display used in both the evaluations of later chapters. I detail the origin and evolution of the display through three design cycles.

The third part of the research process is evaluation. Chapters 7, 8, and 9 evaluate the hypothesis that support for workspace awareness improves the usability of groupware systems. Chapter 7 introduces the methodology that will be used in later chapters, and discusses issues in evaluating and measuring groupware usability. Chapter 8 reports on an exploratory usability study of a shared-workspace groupware system, built to examine the effects of several kinds of awareness displays. Chapter 9 then discusses a laboratory experiment that looked more closely at how certain types of awareness information added to a groupware interface affect product, process, and participant satisfaction.

Finally, chapters 10 and 11 look back at the research and draw conclusions. Chapter 10 considers connections between the findings of the research and larger issues in the design of highly usable groupware. These include the role of workspace awareness in groupware usability, the importance of redundancy, the relationship between workspace awareness and rich workspace interaction, and the tension between design for individuals and design for

groups. Chapter 11 then summarizes the main findings and contributions of the research, assesses progress on each of the three objectives stated above, and suggests a number of directions for further study.

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.

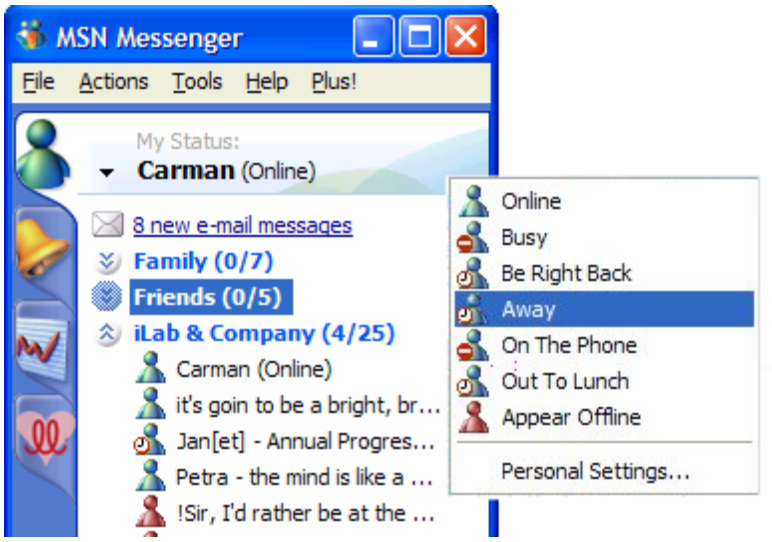

Figure 1.1: Availability states in MSN Messenger.

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.

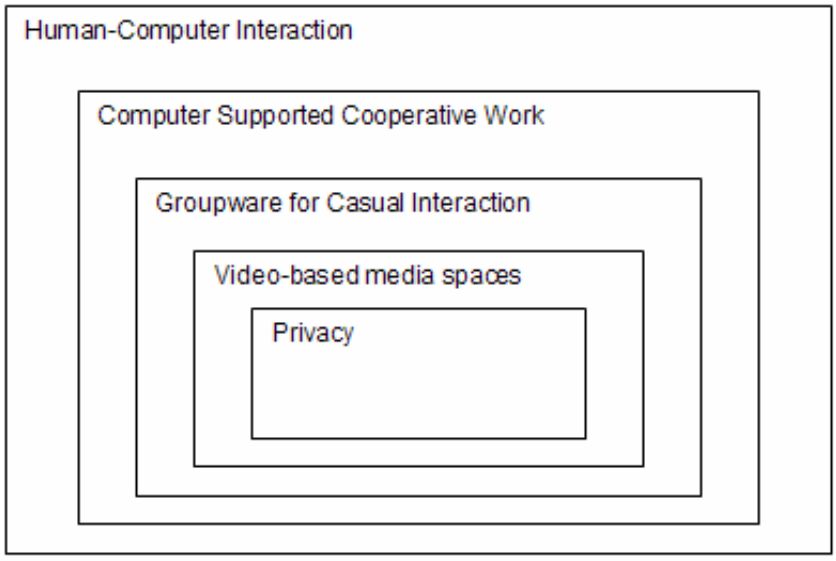

Figure 1.4: The context and scope of my research.

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.

Petra Neumann, 2004
Focus+Context Visualization of Hierarchical Data
Universität Magdeburg. Extract: Chapter 1.

CHAPTER 1

Introduction

Over the recent years, the combination of visualizations and digital media has become increasingly important. Much of today's communication involves visualizations in digital form. Aiding users with different tasks by making the underlying data visual is one important research area in computer science. Such research in the area of visualization is largely based on a characteristic of the human visual system: visual information can be processed in parallel and with a high bandwidth into the human cognitive centers (WARE, 2000). This is also the reason why an effective information display can often lead to insight quicker and more memorably than a few pages of written text. This thesis is concerned with the creation of such an effective information display. As one of the most common data types, hierarchical data structures will be visualized with an emphasis on the display of relations between data items. Motivation for the research presented here will be discussed in this chapter. The discussion will give insight into the inherent problems in the development of the visualization and briefly summarize the further development process discussed in subsequent chapters.

1.1 Motivation

Knowledge obtained from visualizations often comes from insight gained by recognizing that data items are related in some way. In graphs, for example, edges can represent any conceivable kind of relation including those of temporal, causal, or functional nature. The display of relations is, therefore, one of the most essential tasks in information visualization. Understanding relations between items in a visualization helps the viewer to build a mental model of the underlying data. This mental model is needed to understand the scheme or situation to which the presented data refers. For making decisions based on the visualization of data the interpretation of this internal

model is essential. This thesis will introduce a new visualization technique specifically tailored to recognize relations between data elements in a hierarchical data structure.

1.2 Problem Statement

In traditional displays of tree structures, edges represent parent-child relations. If more than the inherent parent-child relations needs to be visualized for nodes in a tree, different encoding techniques are necessary. Figure 1.1 affirms this proposition. The figure displays a tree structure in a traditional node-link diagram. Additional links are introduced to represent binary relations between nodes in the tree. From the display parent-child relations and the additional relations cannot be distinguished. In addition, edge-edge intersections are introduced that lead to a cluttered display.

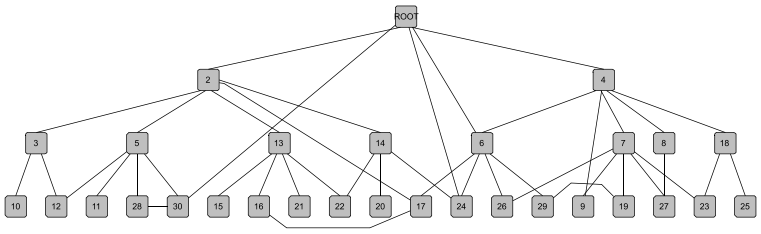

Figure 1.1: Additional relations introduced in a traditional tree layout. Different types of relations cannot be distinguished and edge intersections are introduced.

Typical encoding techniques to circumvent misinterpretations include using color, texture, size, or different node shapes to emphasize additional relations instead of direct links. However, psychological research has found that connectedness can be a more powerful grouping principle than proximity, color, size, or shape (PALMER and ROCK, 1994). The development of a visualization for relations in hierarchical data draws from these findings. This thesis introduces an attempt to solve the mentioned problems of a typical display of direct links on tree structures.

One of the most severe problems in visualizing tree structures is the limited screen space offered by common desktop displays. Trees easily require large aspect ratios which force parts of the tree to be cut from the display. The field of focus+context presentation deals with the display of large information spaces by taking the interest of viewers in parts of the data into account. Research from this field will be integrated into the visualization presented in this thesis to alleviate the display of large hierarchical data structures.

1.3 Results

This thesis introduces a novel visualization for the display of relations in hierarchical data structures. The chosen tree and relation layout solve the mentioned problems of displaying direct links between two data items in a tree structure. A particularly novel aspect is the consideration of relations in the development of focus+context navigation and presentation.

A nested tree layout was chosen for the visualization of tree structures to be able to restrict the visualization in space. An effective use of screen real estate is also achieved through the visualization of relations as arc shaped glyphs that can be adapted in height and shape. Zoom and filter mechanisms help to interactively explore tree content. Focus+context navigation offers a user-centered automation for interaction on the tree. Encodings for attributes of the data and affordances of displayed data items were carefully chosen. An example of the developed *arc tree visualization* can be seen in Figure 1.2.

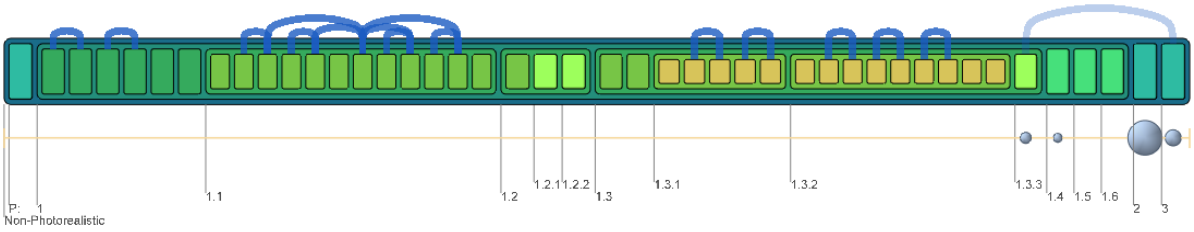

Figure 1.2: An overview of the developed *arc tree visualization*.

The work presented in this thesis was implemented as a JAVA3D application. JAVA3D is a high level, scene graph based API which offers sophisticated graphics rendering utilities. The underlying tree structure was implemented using the composite design pattern as introduced by COOPER (1998). The composite pattern allows the definition of a class hierarchy of simple objects and more complex composite objects so that they appear to be the same to the client program. All screenshots which serve as examples for the developed visualization were taken from the implemented application.

1.4 Thesis Organization

The following sections briefly introduce the contents of subsequent chapters.

Chapter 2

An overview of the research field that forms the framework for this thesis is presented in Chapter 2. It introduces the field of visualization with particular emphasis on information visualization. Further research presented in this thesis is influenced by several research areas in the information visualization community. These areas are concerned with the presentation of different data sets: linearly ordered data, hierarchical data, and graphs. The discussion of visualization techniques from these areas particularly emphasizes the presentation of relations within the data.

Chapter 3

Chapter 3 begins with the introduction of the specific challenges for the development of a visualization. In subsequent sections a layout for linearly ordered hierarchical data structures is developed. Emphasis is placed on the display of additional relational information between nodes in the tree. As a prerequisite to the development of the visualization a set of definitions for the underlying data structure including the relational information as well as the drawing and visualization of the data are discussed.

Chapter 4

Interaction and navigation are two important aspects of an information visualization. This chapter is concerned with the development of interaction techniques for the visualization developed in the previous chapter. First, general operations on the data set are described. Focus+context techniques extend these operations to aid navigation in the visualization. The presented techniques act as a filter mechanism to eliminate unwanted information and avoid information overload.

Chapter 5

Case Studies introduce possible application domains for the developed visualization. It is shown how the requirements specified for the visualization task were implemented in context to a particular application. The remainder of the chapter establishes how an evaluation of the visualization can be carried out without being restricted to a specific application domain. In addition, aspects for which an evaluation would be most appropriate at this stage of the development of the visualization are highlighted.

Chapter 6

This chapter concludes the thesis and summarizes the key contributions of the presented work. Areas for future work are highlighted in the remainder of the chapter.