

**THE UNIVERSITY OF CALGARY**

**Privacy in Video Media Spaces**

by

**Michael John Boyle**

**A THESIS SUBMITTED TO THE FACULTY OF GRADUATE STUDIES**

**IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE**

**DEGREE OF DOCTOR OF PHILOSOPHY**

**DEPARTMENT OF COMPUTER SCIENCE**

**CALGARY, ALBERTA**

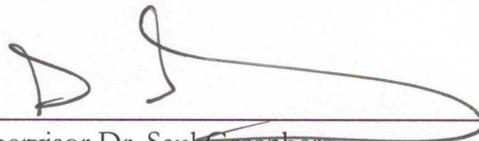
**APRIL, 2005**

**©Michael John Boyle 2005**

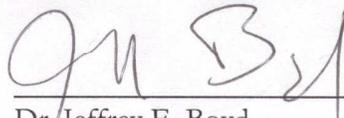


**THE UNIVERSITY OF CALGARY**  
**FACULTY OF GRADUATE STUDIES**

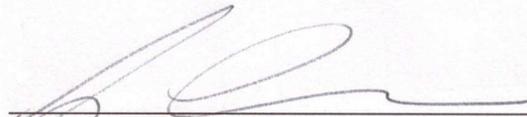
The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies for acceptance, a thesis entitled "Privacy in Video Media Spaces" submitted by Michael John Boyle in partial fulfillment of the requirements for the degree of Doctor of Philosophy.



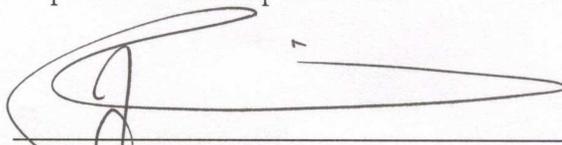
Supervisor Dr. Saul Greenberg  
Department of Computer Science



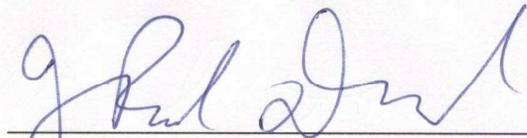
Dr. Jeffrey E. Boyd  
Department of Computer Science



Dr. Sheelagh Carpendale  
Department of Computer Science



Dr. Jeffrey K. Caird  
Department of Psychology



External Examiner, Dr. Paul Dourish  
University of California, Irvine

April 25, 2005  
Date



# Abstract

This thesis informs the design and development of video media spaces that enrich distributed collaboration and preserve privacy. Video media spaces are fully interconnected multi-person audio/video networks that are usually on all of the time. They have the potential to benefit distributed teamwork by providing rich informal awareness and casual interactions to distance-separated colleagues. However, the problem is that media spaces may be invasive to privacy. Although there are techniques proposed to preserve privacy in video media spaces, e.g., video blurring, their effectiveness and usability are not known. Furthermore, there is insufficient theoretical knowledge about what privacy is and how it should be supported to guide the design of privacy-preserving video media spaces.

In this thesis I describe four major contributions that fundamentally change our understanding of the nature of privacy, its role in people's daily lives, and how it may be realised in computer supported cooperative work (CSCW) systems.

First, I developed the COLLABRARY toolkit to rapidly prototype video media spaces and privacy safeguards so that I could understand the human factors of their design, implementation, and use.

Second, I evaluated the blur and pixelize distortion filtration techniques. These studies suggest that although they are widely suspected to be useful for mitigating privacy issues in video media spaces, these techniques fail to adequately balance awareness and privacy in high-risk scenarios.

Third, I developed a comprehensive descriptive theory of privacy in video media spaces based on the social, psychological and CSCW literature about privacy. It decomposes privacy into three normative controls for regulating interpersonal boundaries in an embodied dialectic: solitude, confidentiality and autonomy. This theory yields a powerful vocabulary of terms that disambiguate the many interrelated and subtle meanings of "privacy."

Lastly, I developed a systematic method for applying the vocabulary of the descriptive theory of privacy to analyse and describe video media space design and use. This method reveals omissions in the way a media space prototype handles privacy and hidden assumptions regarding technology design, use, users, and their contexts. I illustrate this method in several case studies that generate further analytical knowledge about privacy in video media spaces.

# Publications from this dissertation

Materials, ideas, and figures from this dissertation have appeared previously in the following peer-reviewed publications.

## **Long papers:**

BOYLE, M., EDWARDS, C. & GREENBERG, S. (2000). The Effects of Filtered Video on Awareness and Privacy. In *Proceedings of the CSCW 2000 Conference on Computer Supported Cooperative Work* [CHI Letters 2(3)], p1-10, ACM Press.

BOYLE, M. & GREENBERG, S. (2005, in press) The Language of Privacy: Learning from Video Media Space Analysis and Design. In *ACM Transactions on Computer-Human Interaction (TOCHI)*.

# Acknowledgments

Esther and John Boyle, my parents. From infancy to adulthood you have given me inextinguishable support coupled with complete freedom to arrange my life and my education according to my desires. I sincerely acknowledge your role in producing my life.

Dr. Saul Greenberg, my advisor. I am so grateful for the faith you have demonstrated in me and my abilities. I did not think I would be good at HCI or would enjoy it until I took your CPSC 481 course. Your energetic personality and enthusiasm towards the field is a source of great inspiration to me. I was so fortunate you hired me to be a summer student working on the Shared Notes PDA-PC system. This opportunity opened the door to an honours thesis, an NSERC scholarship and MSc degree, two internships, a PhD, and one can only guess what's next. Sometimes it feels like everything I know about anything that matters I know because either you taught me it or you created furnished me with an environment that supported me as I learned it on my own. I sincerely acknowledge your role in producing this thesis.

Dr. Sheelagh Carpendale, Dr. Jeff Caird & Dr. Jeff Boyd, my committee members. You have been very generous with your enthusiasm and sharing of experience. I am proud of my thesis and the contribution I make in it and the path I have taken would not have come to me without your guidance. I sincerely acknowledge your role in producing this thesis.

Chris Edwards & Carman Neustaedter, my collaborators. I sincerely acknowledge your role in the studies described in Chapter 4. I am also grateful for your broader roles as colleagues who have provided support, sympathy, and solutions to problems both emotional and intellectual.

Edward Tse, my colleague. I sincerely acknowledge your role in producing the Collabrary software described in Chapter 4. I also wish to acknowledge your role as a friend and companion in the Interactions Lab at the University of Calgary. Also, Shymmon Banerjee, fellow student. I acknowledge your work on the sophisticated COM serialiser that is used in the shared dictionary.

Shaun Kaasten, my colleague. Thank you for motivating me to learn Visual Basic and Microsoft COM programming in C++ with ATL, the foundations of the Collabrary. Thank you also for your companionship both in school and at Microsoft. Also, Michael Rounding, my colleague. Thank you for motivating me to design the shared dictionary well and being a robust beta tester of the Collabrary toolkit and your friendship.

Gregor McEwan, my friend and colleague. Thank you for countless conversations during the most difficult part of my thesis research, when I was burned out and unable to understand the value of my work. I acknowledge your role as stalwart confidant and research peer in the production of my thesis.

Wendy Segelken, Chad Bryant, Nicole Trenholm, Steven & Katherine Weimer, my friends. I have tendencies towards over-work, and you have through social stimulation, kept me grounded. I think I would be something rather despicable had it not been for your copasetic sympathy, Wendy on programming matters, Chad on matters of the discipline of science, Nicole on supervisor matters, and the Weimers on interpersonal relationship matters. I sincerely acknowledge your role in producing my life.

Elena Fanea, Jeroen Keijser, Stacey Scott, Eric Pattison, Charlotte Tang, Tony Tang, Nelson Wong, and the other students of the Interactions Lab at the University of Calgary. Thank you for your friendship,

Darcy Grant, Brad Arlt, Debbie Mazurek and the other members of the Department of Computer Science technical support team, for very forgiving technical support and for many good times together.

Finally, I would like to acknowledge the support of my funding agencies NSERC, Microsoft Research, the Department of Computer Science and the University of Calgary.

# Dedication

*To mom.*

# Table of contents

Approval Page .....	iii
Abstract.....	iii
Publications from this dissertation .....	v
Acknowledgments.....	vi
Dedication .....	viii
Table of contents .....	ix
List of tables.....	xvi
List of figures .....	xvii
Chapter 1— Introduction.....	1
Preface to Act I .....	8
Chapter 2— Video media spaces for awareness and interaction.....	9
2.1 The motivation for video media spaces .....	10
2.1.1 Characteristics of informal interactions in the workplace.....	10
2.1.2 Informal awareness supporting casual interactions .....	12
2.1.3 Roles and characteristics of informal awareness .....	12
2.1.4 Vignettes: Informal awareness supporting interactions .....	13
2.1.5 Visual channel is vital for awareness and interaction .....	14
2.1.6 Summary.....	15
2.2 Video media spaces: Designs explored .....	16
2.2.1 Architecture models and goals.....	16
2.2.2 Community space model: connecting social spaces.....	18
2.2.3 Office share model: Connecting personal workspaces .....	21
2.2.4 Hallway model: Browsing the social environment.....	23
2.2.5 Telephone model: Making transitions smooth.....	25
2.2.6 Community place: combining office spaces into a virtual communal space.....	26
2.2.7 Non-video systems for awareness and interactivity.....	29

2.2.8 Summary .....	31
2.3 Technological barriers .....	32
2.3.1 Bandwidth as a cost/scalability constraint.....	32
2.3.2 Visibility factors: Image size, field of view, resolution, and compression quality....	33
2.3.3 Smoothness factors: Frame rate, jitter, and latency.....	35
2.3.4 Audio problems: Sound level and echo.....	38
2.3.5 Summary .....	39
2.4 Conclusion .....	39
Chapter 3— The COLLABRARY toolkit.....	41
3.1 Toolkit objectives and motivation .....	42
3.1.1 The “missing” toolkit.....	43
3.1.2 Prototypes as design artefacts.....	43
3.1.3 Prototyping toolkits.....	44
3.1.4 The remainder of this chapter: The COLLABRARY in action.....	45
3.2 Multimedia capture and playback .....	46
3.2.1 Hardware abstractions .....	47
3.2.2 Event-oriented architecture .....	48
3.3 Multimedia groupware programming.....	50
3.3.1 Centralised server network architecture .....	51
3.3.2 Organising and storing data in a hierarchical dictionary.....	55
3.3.3 Subscription notifications and Model-View-Controller architecture.....	56
3.3.4 Controlling the presence and distribution of keys and values .....	57
3.4 Multimedia analysis and manipulation .....	60
3.4.1 Provide rich, composable operations despite flawed implementations.....	60
3.4.2 Direct media access for prototyping custom analyses and manipulations.....	64
3.4.3 Summary of multimedia analysis and distortion .....	65
3.5 Toolkit features not demonstrated .....	66
3.6 Missing toolkit features .....	67
3.7 Evaluating the COLLABRARY .....	67
3.8 Conclusion .....	69

Chapter 4— Evaluating distortion filtration.....	71
4.1 Why evaluate distortion filtration?.....	72
4.2 Methodology .....	73
4.2.1 Hypothesis .....	73
4.2.2 Materials: Video sequences.....	74
4.2.3 Materials: Questionnaires.....	77
4.2.4 Experimental design .....	78
4.2.5 Procedure .....	80
4.3 Results .....	81
4.3.1 Pre-test questionnaire.....	81
4.3.2 Identifying the number of actors in a scene.....	82
4.3.3 Identifying posture.....	83
4.3.4 Identifying gender.....	84
4.3.5 Identifying objects in a scene .....	85
4.3.6 Identifying actor activity in a scene. ....	86
4.3.7 Identifying busyness, seriousness and approachability in a scene .....	87
4.3.8 Rating privacy.....	88
4.3.9 Post-Test questionnaire .....	89
4.4 Discussion .....	89
4.5 Neustaedter’s study .....	92
4.5.1 Materials .....	92
4.5.2 Method .....	94
4.5.3 Results.....	94
4.6 Conclusions .....	95
Preface to Act II.....	99
Chapter 5— Language of privacy in CSCW .....	102
5.1 Video media spaces: A crucible for studying privacy.....	105
5.2 Approaches to privacy research .....	107
5.3 Overview of CSCW perspectives on privacy .....	108
5.4 Deliberate privacy abuses: Issues of control.....	109

5.4.1 Methods for controlling media space access .....	109
5.4.2 Control: User interface issues and trade-offs .....	110
5.5 Inadvertent privacy violations .....	112
5.5.1 Disembodiment confounds self-appropriation.....	113
5.5.2 Presence in multiple places forces appropriation in multiple contexts .....	114
5.5.3 Feedback: User interface issues and trade-offs .....	114
5.6 Apprehension .....	116
5.6.1 Surveillance confounds impression management.....	116
5.6.2 Decontextualisation prompts apprehension.....	117
5.7 Reflecting on the problems.....	118
Chapter 6— Perspectives on privacy.....	120
6.1 Private/public dichotomy.....	120
6.2 Privacy as an attribute of places and people .....	122
6.3 Privacy as an interpersonal process .....	122
6.4 Privacy as a need, right, and freedom .....	124
6.5 Privacy as a balancing act.....	126
6.6 Summary: Focusing on an interpersonal process model.....	128
Chapter 7— An integrated vocabulary for privacy and video media space design.....	131
7.1 Solitude .....	134
7.1.1 Attention and distraction.....	134
7.1.2 Verbal and para-verbal solitude controls .....	135
7.1.3 Westin’s four privacy states.....	136
7.1.4 Affordances of space for solitude .....	136
7.1.5 Personal space .....	137
7.2 Confidentiality .....	139
7.2.1 Sensitivity .....	139
7.2.2 Fidelity.....	140
7.2.3 Direct controls .....	142
7.2.4 Computers and confidentiality.....	143
7.2.5 Indirect controls.....	144

7.3	Autonomy.....	146
7.3.1	Preserving and constraining autonomy .....	146
7.3.2	Autonomy-confidentiality-solitude symbiosis .....	149
7.3.3	Identity.....	150
7.3.4	Pseudonymity .....	151
7.3.5	Role Conflict.....	152
7.3.6	Focus and nimbus.....	153
7.4	Conclusion.....	154
	Preface to Act III .....	157
	Chapter 8— Analysing video media space design and privacy .....	158
8.1	Theories that inform design.....	159
8.2	Revealing assumptions hidden in design.....	163
8.3	Analysis vocabulary .....	164
8.4	Analysis procedure .....	166
8.4.1	Step 1: Partition.....	167
8.4.2	Step 2. Describe .....	168
8.4.3	Step 3. Reveal .....	168
8.4.4	Step 4. Summarise.....	169
8.5	Conclusion.....	169
	Chapter 9— Case studies (1).....	171
9.1	Case study #1: The distortion filtration technique.....	172
9.1.1	Step 1. Partition.....	172
9.1.2	Step 2. Describe .....	173
9.1.3	Step 3. Reveal .....	176
9.1.4	Step 4. Summarise.....	180
9.1.5	Conclusion to case study #1 .....	181
9.2	Case study #2: Evaluations of the distortion filtration technique .....	182
9.2.1	Step 1. Partition.....	183
9.2.2	Step 2. Describe .....	183
9.2.3	Reveal.....	186

9.2.4 Summarise.....	189
9.3 Case study #3: The COLLABRARY toolkit.....	190
9.3.1 Step 1: Partition.....	190
9.3.2 Step 2: Describe.....	190
9.3.3 Step 3: Reveal.....	196
9.3.4 Step 4: Summarise.....	200
9.4 Conclusions.....	200
Chapter 10— Case studies (2).....	202
10.1 Neustaedter & Greenberg’s HOME MEDIA SPACE.....	202
10.1.1 Partition.....	203
10.1.2 Describe.....	205
10.1.3 Reveal.....	213
10.1.4 Summarise.....	215
10.2 Hong’s (2004) privacy risk model framework.....	216
10.2.1 Partition.....	218
10.2.2 Describe.....	219
10.2.3 Reveal.....	225
10.2.4 Summarise.....	229
10.3 Conclusion to the 2 <sup>nd</sup> set of case studies.....	230
10.4 Reflecting back on all of Act III.....	230
Chapter 11— Conclusions.....	232
11.1 Progress on thesis problems, goals, and deliverables.....	233
11.1.1 Problem #1: Rapid prototyping toolkit.....	233
11.1.2 Problem #2: Distortion filtration evaluation.....	234
11.1.3 Problem #3: Descriptive theory of privacy in video media space design.....	235
11.1.4 Problem #4: Methods for describing privacy.....	236
11.2 Thesis contributions.....	237
11.2.1 Major contributions.....	237
11.2.2 Minor contributions.....	238
11.3 Reapplying the results.....	239

11.3.1 Group collaboration tools .....	239
11.3.2 Pervasive computing infrastructure and applications.....	240
11.3.3 Regulatory measures for protecting individual rights and freedoms.....	241
11.4 Future work.....	242
11.4.1 Near term: Improving on the results .....	243
11.4.2 Far term: Extending the results .....	244
11.5 Final words.....	247
Bibliography.....	248
Appendix A. Co-author permission .....	258
Appendix B. Glossary.....	262

# List of tables

Table 2.1 Persistency and directedness as characteristics of video media space architecture. ...	17
Table 4.1 Topic of questions asked in the during-test questionnaire.....	78
Table 4.2 Thresholds for identifying awareness cues and for preserving privacy.....	90
Table 5.1 Vocabulary that embodies the descriptive theory of Act II.....	104
Table 6.1—Negative aspects of insufficient control over privacy, and positive aspects of sufficient and necessary control over privacy. From Altman (1975), Brierley-Newell (1995). .....	124
Table 7.1—Interpersonal distances and the interactions supported at each (Hall, 1966).....	138
Table 10.1—Genres of disclosure in Neustaedter & Greenberg (2003) HOME MEDIA SPACE. ....	211
Table 10.2—Conditions in which HMS privacy preserving features will counter inattention slips in self-appropriation.....	215
Table 10.3—Questions used in Hong’s Privacy Risk Analysis method. ....	217
Table 10.4—Questions used in Hong’s Privacy Risk Management method. ....	218

# List of figures

Figure 2.1 The Microsoft vKITCHEN video media space (from Jancke et al, 2001). .....	19
Figure 2.2 The CAVECAT media space connected personal offices with conventional video media space node equipment (a) and miniaturised integrated units called HYDRA surrogates (b). Photo: Buxton (1997). .....	21
Figure 2.3 In the media space COMEDI (a), video snapshots for awareness are arranged on a perspective wall, while a hi-fi video link (bottom part) supports private 1-on-1 interactions. In the NOTIFICATION COLLAGE (b), snapshot video links for awareness and are shown beside “media items” for text- and artefact-based interactions, creating a Community place.....	27
Figure 3.1—C# source code (and screenshot of resulting application) for a “blurry mirror” that demonstrates how simple multimedia capture and playback is with the COLLABRARY. ....	46
Figure 3.2—User interface for the basic media space prototype developed in Section 3.3.....	51
Figure 3.3—Complete C# source code for an <i>n</i> -way video media space working system. Some text has been abridged with ellipses for display purposes. ....	54
Figure 3.4—In the prototype video media space, clients post audio and video data to a shared dictionary server and receive notification of audio/video posted by others. The shared dictionary maps hierarchically organised string keys to values. ....	56
Figure 3.5—Background subtraction can be performed with just a single method call in the Collabrary, but the implementation is not fully robust. Part of the person’s body (a) has been misclassified as background in (b). ....	61
Figure 3.6—Blurring can be easily combined with face tracking to produce interesting effects with the Collabrary. ....	62
Figure 3.7— C# code to implement Gutwin’s traces for video and an EKG-like activity history display and a representative image. (Contrast increased for print reproduction.)	63
Figure 3.8—Custom video analysis/manipulation technique to remove flesh-coloured pixels.	64
Figure 4.1—Video scenes used. ....	75
Figure 4.2 Filters and levels tested (top: blur; bottom: pixelize). ....	76
Figure 4.3 Screenshot of in-test questionnaire software. ....	79
Figure 4.4 Screenshot of in-test software showing last question. ....	80
Figure 4.5 Normalised number of actors identified <i>vs.</i> filter level. ....	82

Figure 4.6 Mean threshold level where participants correctly identified actor posture.....	84
Figure 4.7 Mean threshold level where participants could identify objects. ....	85
Figure 4.8 Mean threshold level where participants could identify activities.....	86
Figure 4.9 Mean threshold at which participants could confidently indicate busyness, seriousness, and approachability.....	87
Figure 4.10 Mean privacy rating by level: 1 is unprotected, 5 is highly protected.....	88
Figure 4.11—Scenes used in Neustaedter's second study were of much higher risk than those used in the first study.....	93
Figure 4.12—The median and range of blur levels chosen by participants for each scene. Blur level 0 represents choosing to turn the camera off. Figure reproduced from Neustaedter, Greenberg & Boyle (2005).....	95
Figure 4.13—The first study found that in mundane office scenes that have low privacy risk, privacy and awareness can be balanced with the blur filter around level 5. ....	96
Figure 4.14—The second study found that privacy and awareness cannot be adequately balanced in risky scenarios routinely expected for home telecommuters.....	97
Figure 5.1—A design space showing some previously explored techniques for preserving privacy in video media spaces. ....	110
Figure 7.1—The blur distortion filter can operate at a variety of levels. Each level affects fidelity and risk, which in turn affect awareness and one’s ability to control confidentiality. The left part of the figure shows a mundane scene used in the Boyle et al filter study for which privacy could be balanced with awareness. The right part of the figure shows a risky scene used in the Neustaedter et al filter study for which privacy and awareness could not be balanced.....	141
Figure 9.1—The first case study examines distortion filtration (like the blur filter) being used for this idea that privacy and awareness are opposing design goals that might be balanced. (Figure repeated from Chapter 4.) ....	172
Figure 9.2 Vocabulary for the Partition analysis step. ....	172
Figure 9.3 Vocabulary to describe confidentiality.....	173
Figure 9.4 Vocabulary to describe solitude and autonomy.....	174
Figure 9.5 Vocabulary to describe mechanics of privacy.....	174
Figure 9.6 Vocabulary to describe computers and privacy.....	175
Figure 9.7 The distortion filtration technique does not preserve privacy when the sensitive aspects of a scene are of low precision. ....	178
Figure 9.8 Some performances are social acceptable in the actor’s setting, but not the audience’s setting and even heavy filtration does not change this.....	180

Figure 10.1—The Home Media Space prototype included environmental sensors that were used to autonomously enable/disable capture. The camera was mounted on a servo motor so that it could be rotated to face the room or face the wall. (Figure reproduced from Neustaedter & Greenberg, 2003.)..... 203



## Chapter 1—Introduction

This thesis concerns video media spaces and privacy. Video media spaces are fully interconnected multi-person audio/video networks where the cameras, microphones, displays and speakers are usually on all of the time: see, e.g., Bly, Harrison, & Irwin (1993) for a chronology of media space design. Media spaces hold great potential to benefit distributed teamwork by providing distance-separated members of a group with rich informal awareness supporting the micro-coordination of casual interactions: see, e.g., Kraut, Egido & Galegher, (1988) for the motivation for focusing on informal awareness and casual interactions.

Yet, these benefits may not be realised because media spaces may be invasive to people's privacy. More importantly, many media spaces have been perceived as privacy-invasive, e.g., the CRUISER media space: Fish et al (1992). Many people do not wish to use video media spaces that prompt privacy violations, even if the media space benefits distributed collaboration: see, e.g., Jancke et al (2001) for a description of some reactions people have to video media spaces. While there are a limited number of techniques such as video filtration (Zhao & Stasko, 1998) that have been employed to preserve privacy, their effectiveness or usability is not known.

Therefore, this thesis seeks to inform the development of video media spaces that support privacy **and** enrich distributed collaboration. This issue is considered to be a human-computer interaction design problem in which individual, social, and technical factors weigh heavily. This is a widely held view motivated by Bellotti (1998) and others. Privacy is an incredibly complex subject to understand in the context of video media space design, as described by, e.g., Palen & Dourish (2003).

The over-arching goal of this thesis is to inform the design and implementation of a privacy-preserving video media space. Two central research questions emerge from this goal:

**Research Question #1:** What low-level technological factors need to be considered when building a privacy preserving video media space?

**Research Question #2:** What high-level social-psychological factors need to be considered when designing a privacy preserving video media space?

These questions mark out a long-term research programme that extends well beyond the end of this thesis. Within this large research area I have carved out a set of short-term problems and goals to specifically address in this thesis.

The first problem concerns rapid prototyping of video media spaces. This is important because in order to understand the significant and relevant human factors involved in the design and use of such complex socio-technical systems one must be able to rapidly prototype them (Kaplan, 1998). Addressing this problem contributes to answering research question #1.

**Thesis Problem #1:** It is hard to rapidly develop video media spaces because the programmatic interfaces for multimedia are complicated and require considerable programmer effort and expertise.

**Thesis Goal #1:** Develop a toolkit to support the rapid prototyping of video media spaces and the distortion filtration method for preserving privacy therein.

**Thesis Deliverable #1:** To achieve goal #1, I have produced the COLLABRARY toolkit and the discussion of its design in Chapter 3 of this thesis. The toolkit works with rapid application development tools like Microsoft Visual Basic to provide the low-level infrastructure needed to build a media space.

The next problem considers the effectiveness of the popular distortion filtration techniques that alter video images. It is important to evaluate these techniques because they are already being used to ostensibly preserve privacy. Yet, privacy violations may result if users are misinformed about the effectiveness of the media space's privacy safeguards. Addressing this problem by evaluating distortion filtration contributes to answering research question #1, dealing with a low-level technique for preserving privacy, even though the evaluation considers high-level human perceptions of the technique's performance.

**Thesis Problem #2:** It is widely suspected that distortion filtration may be useful for mitigating privacy issues in video media spaces but its usefulness has not been rigorously evaluated and there is no guidance as to how much filtration is ideal.

**Thesis Goal #2:** Determine if it is possible to use the distortion filtration technique to strike a balance between awareness and privacy in a video media space. If it is possible, determine at which levels a balance can be reached.

**Thesis Deliverable #2:** To achieve goal #2, I have conducted a semi-controlled laboratory user study in collaboration with Chris Edwards and Saul Greenberg to evaluate the blur and pixelize filters for balancing awareness and privacy in a video media space. This study was published as Boyle, Edwards, & Greenberg (2000). I also conducted a second study with Carman Neustaedter that addressed methodological problems in the first study. This was published as Neustaedter, Greenberg & Boyle (2005). The results of these studies suggest that while an adequate balance between awareness and privacy can be found for benign situations, this balance will not be found in risky scenarios using the distortion filtration technique alone. The studies and results are discussed in Chapter 4.

I originally hoped that the distortion filtration technique would balance privacy and awareness, and that I could exploit this to design and build a prototype of a privacy preserving video media space after completing thesis goals #1 and #2. Yet, because this was not the case, my progress was blocked. Even though I had a toolkit to help me implement things the lacklustre performance of distortion filtration under evaluation left me without a good privacy preservation technique to implement in a prototype. There was nothing in the studies to suggest that preserving privacy in video media spaces is an intractable design problem. Rather, the results suggest that there is more to the issue of privacy in video media spaces than can be explained by an awareness/privacy design trade-off or be effectively mitigated by the distortion filtration technique by itself.

On reflection, what really blocked my progress was the lack of a theoretical understanding of privacy itself that is needed to know how to design support for it. It proved hard to acquire this understanding. While there are theoretical descriptions of privacy in CSCW as it pertains to media spaces (e.g., Bellotti, 1998; Palen & Dourish, 2003) these descriptions do not account for the full range of suspected privacy issues in media space design and use. While there are rich theories of privacy developed in disciplines such as sociology and psychology, these theories are formed at a very abstract level that is hard to relate directly to media space design and use. Even after absorbing both categories of literature it is hard to apply these theories to video media spaces because the scope of privacy and the language used to talk about it differs with each researcher and discipline. In this context, two additional pairs of thesis problems and goals emerged.

The first additional problem is the lack of a comprehensive vocabulary for privacy. This is important because people who want to analyse the strengths and shortcomings of the privacy support in a video media space need a rich common vocabulary to articulate findings and discuss them with others. Addressing this problem contributes to the answering of research question #2 by giving labels to the high-level social-psychological factors relevant to privacy and video media space design.

- Thesis Problem #3:** There is no comprehensive vocabulary of privacy terms—one that integrates conceptions and theories of privacy from many disciplines—to support unambiguous description of how privacy is affected by video media space design and use.
- Thesis Goal #3:** Integrate privacy theories and observations from many disciplines of scientific inquiry to produce a vocabulary for describing privacy and a video media space’s effect on it in an unambiguous and comprehensive manner, accounting for at least the privacy issues reported in previous literature.
- Thesis Deliverable #3:** To achieve goal #3, I have produced a descriptive theory of privacy: one that supports designers’ dialogue by identifying the important privacy concepts and phenomena and providing a vocabulary of terms for them. This theory is described in Chapters 5 through 7 of this thesis.

As we will see, Chapter 5 synthesises the observations of numerous privacy researchers in CSCW, notably: Bellotti (1998); Palen & Dourish (2003); Grudin (1999); and Nardi et al (1997). Chapter 6 broadens this by incorporating perspectives of privacy borrowed from other disciplines such as law, sociology, and psychology: Altman (1975); Gavison (1980); Goffman (1959); and Schwartz (1968). Chapter 7 presents the kernel of the vocabulary: a tri-partite theory of privacy—informed by the many works cited but ultimately the unique product of my own integration and deliberation—that decomposes privacy into solitude, confidentiality and autonomy modalities of control and then relates these back to the video media space privacy issues discussed in Chapter 5.

I had hoped that the deeper understanding of privacy I gained through developing the privacy vocabulary would guide me directly towards the design of privacy safeguards for use in a video media space and to metrics for the user-centred evaluation of such safeguards. Through both internal struggle and external dialogue with my supervisory committee members, I came to understand that it is not possible to take a descriptive theory of privacy and design i.e., one that supports description of privacy-related phenomena, and use it directly

as though it were a prescriptive theory of privacy, i.e., one that provides guidelines and heuristics for design such as the U.S. Privacy Act of 1974's fair information practices (U.S. Congress, 1976).

While the development of a prescriptive theory of privacy in video media spaces is part of the long-term goal of building a privacy-preserving video media space, such a theory is beyond the scope of this thesis. Instead, I have focused my attention on exploring and illustrating how the descriptive theory of privacy informs the design of privacy-preserving video media spaces. The last problem/goal pair I have addressed in this thesis stems from this.

In particular, the vocabulary I developed to address thesis problem #3 identifies the set of concepts that are important to privacy and provides labels for talking about them—which are valuable contributions in their own right—but it does not provide a systematic method for applying these concepts to describing video media space design and use. Such a method is needed to think analytically about the merits and demerits of the privacy support in a media space. Addressing this problem contributes to answering research question #2 by providing a way to reveal hidden assumptions or omissions that constrain the video media space's support for privacy and circumscribe limits on the conditions in which the media space will actually preserve privacy.

**Thesis Problem #4:** There is no systematic method for applying the concepts in the privacy vocabulary to understand and inform the design of privacy-preserving video media spaces.

**Thesis Goal #4:** Develop a systematic method of applying the terms in the privacy vocabulary to describe and analyse the effect of a video media space's design and use on privacy.

**Thesis Deliverable #4:** To achieve thesis goal #4, I have produced a method that guides this kind of analysis. This method and case studies to illustrate its application are presented in Chapters 8 through 10 of this thesis.

The method will reveal hidden omission and assumptions in the media space's support for privacy and allow the analyst to set out limits on the circumstances in which privacy will be preserved. The object of analysis might be a video media space prototype but it could also be: an idea for a privacy safeguard; an evaluation of a privacy safeguard; or, guidelines for the design and implementation of privacy-sensitive systems. In this method, an initial description phase uses the broadly-based vocabulary to systematically describe the object of analysis, teasing apart the individual aspects of privacy and revealing omissions in the object of analysis's handling of these different aspects. A reflective phase has the analyst recombine these descriptions to reveal hidden assumptions regarding technology design, use, users, and their contexts.

Chapter 8 describes this method and its motivation, while Chapters 9 and 10 illustrate the application of the method through detailed case studies involving all of the aforementioned kinds of objects of analysis. While the descriptions produced and the omissions and assumptions revealed in these case studies are valuable and informative products of this research, the motivation for performing these case studies is to illustrate the analysis method which addresses thesis goal #4.

In addition to the chapters described, in Chapter 2 I provide a thorough survey of casual interaction, informal awareness, and video media space design literature; and an analysis of the limits on media space hardware and software in terms their effect on informal awareness and casual interactions.

These chapters are grouped into three-chapter collections (called Acts I, II, and III) organised in chronological order. I call these major thesis divisions 'acts' to accentuate the 'bursty' sort of evolution of my thinking on the greater research goal and questions. Act I collects the necessary background, my initial toolkit design, and my evaluation of distortion filtration. Act II collects the descriptive theory of privacy. Act III collects chapters that evaluate the theory by showing how it can be applied to the work in Act I. Finally, Chapter 11 concludes the thesis by revisiting the problems and goals stated here, reiterating how the various thesis products satisfy the goals and address the problems. It also revisits the larger research goal and questions to sketch an outline of the kinds of investigations and analyses that are needed beyond what has been accomplished in this thesis.

## PREFACE TO ACT I

LOW-LEVEL TECHNICAL FACTORS RELATED TO PRIVACY-PRESERVING  
VIDEO MEDIA SPACE DESIGN.

The three chapters bundled in this first act present my early work focused on low-level matters of technology and video media space design.

My original goal was to build a privacy preserving video media space. At this stage, I conceived of privacy as a system design goal that conflicted with awareness (part of the motivation for video media spaces). In order to have more privacy, it was necessary to have less awareness. My strategy was to use blurring to filter out some (but not all) awareness cues to increase privacy. My plan was to develop a toolkit to rapidly build video media space prototypes. In parallel, I planned to conduct a user experiment to determine how much I would need to blur video in order to adequately balance awareness and privacy. The chapters in this act tell this part of the story.

I anticipated that I would be able to use my toolkit to implement a video media space prototype and incorporate blurring as my study results indicate. This prototype would then be deployed to my research colleagues, and we would use it for an extended period of time to see if simple blurring sufficed or if more needed to be done.

In the end, these studies showed blurring was insufficient, and it was unclear if other techniques could be substituted in its place. Although I had the technological capacity to prototype a privacy preserving video media space, I had no idea how to design it. Consequently, I brought the thread of inquiry encapsulated in this act to a close, as I was dissatisfied with the “bottom-up” build-something-test-it-make-it-better approach.

While the material I discuss in this act may not have as far reaching effects as that covered in the next two, it is nonetheless a necessary and meaningful component of the overall contribution I make in this thesis. What I perceived at the time to be a failure ended up providing me with the context and motivation for the approach I take up in Act II.

## Chapter 2—Video media spaces for awareness and interaction

The first of the two over-arching thesis research questions put forth the previous chapter asks about the low-level technological factors that are relevant to building privacy preserving video media spaces. This chapter contributes to answering this question by providing a historical development of video media spaces from a technology-oriented perspective. Video media spaces were developed as a technological solution to a specific problem: distributed colleagues miss out on casual interaction because they lack rich cues for informal awareness. Thus a privacy preserving video media space must provide these cues if they are to benefit collaboration. Consequently, in this chapter I seek to first understand how a media space can benefit collaboration without worrying about the implications to privacy.

Section 2.1 examines casual interactions and informal awareness as characteristics of co-located collaboration that motivate the use of video media spaces for distributed groups. Section 2.2 describes video media space architecture models and the prototype media spaces that have implemented them. Section 2.3 examines media space hardware/software limitations and their impact on casual interactions and informal awareness. Section 2.4 summarises the contents of the chapter. Later, Chapters 3 and 4 will cover technical factors related to the rapid prototyping and evaluation of privacy-preserving video media spaces.

## 2.1 The motivation for video media spaces

Video media space design has historically been based on conclusions drawn from empirical investigations into how groups leverage physical propinquity to collaborate e.g., Kraut, Egido, Galegher (1988), and Ackerman (2000). When people work together in a co-located setting they have many opportunities for **casual interactions**. These are unplanned encounters that happen through the work-a-day world. They are vital to socially satisfying collaboration. Video media spaces are specifically designed to improve opportunities for distance-separated collaboration by supporting casual interactions. In this section I present the CSCW findings related to casual interactions that motivate and guide video media space design.

### 2.1.1 Characteristics of informal interactions in the workplace

Kraut, Egido, & Galegher (1988) conducted observational studies of communication and collaboration among researchers in the scientific community to establish the effect of physical propinquity on collaboration. In their study, they observed that researchers spend a large portion of their time in unplanned face-to-face encounters with others. Later, Whittaker, Frohlich, & Daly-Jones (1994) found that such interactions consume anywhere from 25% to 70% of office workers' time (depending on job type). These interactions serve as a kind of 'glue' that binds collaborators with regular contact between times of tightly coupled synchronous interactivity. The distinguishing characteristics of casual interactions articulated in these two works are summarised below.

- **Unplanned, brief, and frequent.** About 90% of casual interactions are unscheduled. They occur about once every 12.5 minutes and last for a mean of 2 minutes. The mean duration is less when the people involved frequently interact with each other. These descriptive statistics indicate that casual interactions are **lightweight**: it requires little effort to locate potential interlocutors and engage them for interaction. People can capitalise on serendipitous opportunities to spontaneously interact at the moment such casual interactions are relevant and appropriate.
- **Involve small groups of people familiar with one another.** About half of all casual interactions involve people who interact with each other more than twice daily. About 80%

of casual interactions begin with just two people and 80% of those subsequently gain a third participant.

- **Useful for accomplishing artefact-centric work and fostering social bonds.** About half of all casual interactions involve discussions around documents or other work artefacts. They often occur in places where people do their work, e.g., personal or shared offices. Tools—like pens, paper, and whiteboards—and artefacts—such as documents or specialised equipment—are readily available to augment the interactions. The topics of conversation shift smoothly between production-oriented ones—e.g., transfer of undocumented business information, making business decisions, or asking for and offering assistance—and social ones—e.g., inquiring about family life or recreation plans.

Hackman (1985) cites three goals that groups must accomplish to be successful: production, group maintenance, and member support. Casual interactions are typically well suited to supporting group production and coordination activities under conditions of uncertainty (Kraut et al, 1990). They help socialise new members: “teach them the ropes” and familiarise themselves with colleagues’ roles and habits and the unwritten rules about the way things are done in an organisation. They contribute to making work relationships satisfying and fostering a sense of group membership.

- **Affected by physical separation.** Ten times as many collaborative endeavours take place between pairs of researchers who have offices along the same corridor of a building than those pairs whose offices are in different corridors of the same floor. Collaborators whose offices are separated on different floors have about the same likelihood of collaborating as those whose offices are in different buildings! These observations applied equally to people with semantically related interests or organisationally related roles.

The final observation—that distance impairs distributed collaboration and casual interaction—strongly motivates the use of computers and telecommunication technology to overcome the distance barrier and enrich collaborations among distributed workgroups.

### 2.1.2 Informal awareness supporting casual interactions

Kraut, Egido, & Galegher (1988) also found that when collaborators inhabit the same physical space they conveniently maintain an up-to-the-minute understanding of the people around and the activities they are engaged in. Put more precisely, they maintain a keen sense of:

- **presence:** knowledge of the spatial location of others, in particular, if they are in their offices or in nearby common public spaces;
- **activities:** knowledge of the past, present, and planned future activities of others; and,
- **availability:** understanding of when it is appropriate to engage another for interaction.

They labelled the aggregate of these pieces of information—directly sensed or implicit inferred—**awareness**. Gutwin (1997) later specifically labelled it **informal awareness** to differentiate it from other kinds of awareness such as workspace awareness (of activities occurring in our shared workspaces) or conversational awareness (of the state and evolution of human dialogue). Bly, Harrison & Irwin (1993) and others have used the term **peripheral awareness** to emphasise that much of this is information people track in the periphery of their consciousness without ever actually attending to it. It may include other kinds of information, such as weather or traffic conditions, as in Cadiz et al (2002).

### 2.1.3 Roles and characteristics of informal awareness

Whatever the moniker used, informal awareness is vital to coordinating casual interactions.

- It serves as stimulus to remind participants of a need for conversation (Kraut, Egido, & Galegher, 1988).
- It makes appropriate times for casual interactions known to potential participants (Dourish & Bellotti, 1992).
- It provides context (i.e., background information) for those conversations (Dourish & Bellotti, 1992).

Further underscoring its importance, Kraut and colleagues reported that distributed collaborators' interactions become more formal (time, location and agenda planned in advanced) when they lack informal awareness. Informal awareness makes transitions into and

out of tightly-coupled interactivity smooth (Kuzuoka & Greenberg, 2000). This makes it possible for people to back-out gracefully (without making other participants feel bad), recover from premature abortion of a conversation, and ensure that the time, agenda, location, and the participants involved are mutually satisfactory. Informal awareness has several important characteristics described in various sources: Kraut, Egido & Galegher (1988); Gaver et al (1992), Dourish & Bellotti (1992); and, Kuzuoka & Greenberg (2000).

- **Continuous yet dynamic.** One's senses constantly monitor the environment for informal awareness cues and in doing so one maintains up-to-the-moment awareness of the presence and availability of others.
- **Lightweight and unobtrusive.** One monitors the environment for awareness cues as one moves about the workplace. Little extra effort is expended to gather awareness and gathering it does not take much attention away from one's other work activities.
- **Peripheral and implicit.** The tabulation of events and changes in one's environment and the inferring of changes in others' presence and availability often occur in the periphery of one's consciousness. One can rapidly make subtle leaps from observation to hypothesis. The whole process consumes so little cognitive effort that it usually goes unnoticed.

#### 2.1.4 Vignettes: Informal awareness supporting interactions

To illustrate how informal awareness is gathered, consider the following seemingly mundane scenario. John and Mary are co-workers. This morning, John is in his office when he hears the elevator doors open and footsteps along the corridor adjacent to his office's corridor. As the footsteps near him he instinctively turns his glance towards his open office door. He sees Mary walk by and becomes aware of her presence. Concurrently, Mary glances into open office doorways as she walks down the corridor. She becomes reciprocally aware of John this way. Even if they make eye contact, they might not greet each other or engage in conversation. Accruing this awareness did not interfere with John's work. For Mary, accruing this awareness was an effortless side effect of walking down the corridor towards her own office.

The workplace is populated with scores of rich cues signalling presence and availability. For example, how far open an office door is set could be used to infer the availability of the occupant: closed when she is absent; slightly ajar when busy; or wide open when available. A quick glance into the office can reveal further information about the occupant's availability. For example, an empty office with a colleague's coat draped around a chair on a chilly day could suggest that the office occupant is still somewhere about the building. Likewise, she may be considered busy if seen on the phone or with a visitor. More subtle cues may be sense from the environment: for example, the aroma of an ethnic cuisine wafting from the kitchen may indicate a particular person is having lunch.

An alternate scenario illustrates how these informal awareness cues support informal interaction. John needs to speak with Mary about amendments to a report they are writing together. Upon seeing her walking in the corridor, John is reminded of the need to talk with Mary. As Mary glances in to John's office while walking by, John merely holds up the report for Mary to see. She takes this as an implicit signal requesting her to step inside his office to talk. Mary merely nods and carries on her way, starting to take off her jacket. There is no miscommunication here: John can confirm Mary's meaning when he hears her open her door, hang up her coat, and promptly walk back to his office. This exchange shows how mutual awareness of presence and activities serves as an important cue to establish the context for conversation (e.g., the report John holds in his hands) and to time conversation for when they are most relevant and appropriate (e.g., after Mary hangs up her coat).

### 2.1.5 Visual channel is vital for awareness and interaction

One can see from the vignettes presented that many of the informal awareness cues people use to infer presence and availability rely on non-verbal visual and auditory cues. Both channels can be peripherally sensed and from afar. For example, the human visual perception system has a large peripheral visual cone (a field of view of about 180°) but at any instant can only see a small portion of that (about 0.05°) clearly. Humans are well adapted to taking in large quantities of visual information peripherally. These characteristics make the visual channel an ideal fit for informal awareness cues. Similar statements may be made of non-speech audio but there are some special cues for regulating interpersonal interactions that are particular to video.

Whittaker & O’Connail (1997) also identify four kinds of visible bodily behaviours.

- **Gaze** indicates locus of interest and signals attentiveness.
- **Facial expressions** augment the spoken word with affective information that disambiguates the meaning of sensitive or easily misinterpreted content.
- **Gestures** convey content or redundantly reinforce attitudinal messages in the spoken word.
- **Posture** signals affective information about the emotional state of the individual.

These visible behaviours are subtle but extremely significant. While people can adapt to their absence the resulting conversations tend to be more formal and less personable than face-to-face conversations. Several theoretical frameworks regarding conversation and interaction place great significance on them, as well. Whittaker & O’Connail (1997) ascribe to these visible bodily behaviours **process coordination** and **content coordination** roles. Process coordination is roughly analogous to informal awareness. Content coordination is much like Gutwin’s (1998) conversational awareness. This framework mirrors Clark & Brennan’s (1992) discussion of mechanisms for grounding in conversation. Similarly, Argyle (1972) ascribes to non-verbal communication two important roles: managing the immediate social situation (just like process coordination) and sustaining verbal communication (just like content coordination).

Consequently, always-on video that connects the spaces inhabited by distant collaborators is seen as a surrogate for the direct visual channel. Fish et al (1992) cite evidence that suggests video media spaces increase the frequency, spontaneity and richness of casual interactions. These precisely match the descriptions given in section 2.1.1 and are the motivation for using video in media spaces.

### 2.1.6 Summary

This section discussed the characteristics of casual interaction and informal awareness. An understanding of these characteristics is vital to determining the low-level technological factors that are critical to media space design (research question #1 posed in Chapter 1). As well, methods for preserving privacy must not undermine the informal awareness and casual

interactions that the media space seeks to support. Below I list the main observations put forth in the previous text.

- Spontaneous casual interactions and informal awareness are vital to successful group work [Section 2.1.1 and 2.1.2].
- Informal awareness helps people coordinate casual interactions [Section 2.1.3].
- Informal awareness and casual interactions are easy when people are co-located and hard when they are apart [Section 2.1.4].
- Video is a rich source of informal awareness cues [Section 2.1.5].

In the next section, I will draw upon this knowledge to explain how different media space designs either facilitate or impede casual interaction and informal awareness.

## 2.2 Video media spaces: Designs explored

A video media space is essentially a collection of **nodes**—a single camera, microphone, display, and speaker combination installed in one location—and the logical **links** between them. An important thing to consider is that nodes have costs associated with their installation and links have costs associated with their operation. In this section I will categorise video media space designs that have been produced over the last 15 years in terms of two fundamental characteristics—technical factors, if you will—of the links between nodes: **persistence** and **directedness**.

### 2.2.1 Architecture models and goals

A pair of nodes can be installed in distance-separated personal spaces so that the link can be established originating from one individual and directed to another specific individual. Alternatively, the nodes can be installed in shared spaces so that the link is shared by all people occupying that space: this is an undirected link. Both designs affect the size of the community of distributed collaborators that can be supported with a given number of links and nodes.

Although it became possible to deploy a great multitude of video media space nodes cheaply as the price of digital video equipment fell during the 1990's, the costs of operating

	<b>Undirected</b>	<b>Directed</b>
<b>Persistent</b>	<p><b>Community space</b> Persistent links between common rooms allow walk-up-and-use interactions between non-specific people in those rooms.</p> <p>Discussed in: Section 2.2.2</p> <p>Example system: vKITCHEN</p>	<p><b>Office share</b> Persistent links interconnecting a very small number of people to simulate sharing a single “virtual office.”</p> <p>Discussed in: Section 2.2.3</p> <p>Example system: CAVECAT</p>
<b>Intermittent</b>	<p><b>Hallway</b> Probe random locations to maintain awareness, like peering into open offices while walking down a hallway.</p> <p>Discussed in: Section 2.2.4</p> <p>Example system: CRUISER</p>	<p><b>Telephone</b> Establish a connection to a specific person just to engage in conversation, much like using the telephone.</p> <p>Discussed in: Section 2.2.5</p> <p>Example system: MONTAGE</p>

Table 2.1 Persistency and directedness as characteristics of video media space architecture.

links did not decrease by much. Instead, it became possible to use switching networks to turn some links on and others off from time to time so as to limit the total number of concurrent links in operation at any given time. That is, it became possible to use intermittent links in addition to persistent links to create an apparent increase in the total number of links without added cost to their operation.

The intersection of these two technical factors produces the grid of video media space architecture models shown in Table 2.1. In Sections 2.2.2~2.3.5 I will describe a different video media space prototype for each of these four different architecture models. Many video media space prototypes combine features to support more than one kind of architecture model. For example, the RAVE media space had the ability to browse the social environment using the Hallway model to find potential conversation partners as well as the ability to launch into an extended open link using the Office share model (Gaver et al, 1992). In Section 2.2.6, I discuss an additional model of media space architecture—built from combining personal offices spaces into a virtual communal space—that is entirely different from the other four. This architecture model is illustrated in the PORTHOLES (Dourish & Bly, 1993) and NOTIFICATION COLLAGE (Greenberg & Rounding, 2001) media spaces.

Putting the discussion of media space architectures aside for a moment, I wish to talk about the objectives for these architectures. I will use these objectives in both the media space prototype descriptions to follow as well as a discussion of alternatives to video media spaces for informal awareness and casual interactions (section 2.2.8).

A critical design goal for video media spaces emerges from descriptions of casual interaction and informal awareness given in the previous section: media spaces should be designed to be **ready-to-hand** (Dourish, 2001). This means that people do not become conscious of the fact that the media space is mediating their interactions. Two critical human factors that affect ready-to-handedness in a media space are **ubiquity** and **social transparency**. I will use these two human factors to evaluate the ready-to-handedness of previous video media space designs.

A ubiquitous video media space has nodes and links located in many places where casual interactions are likely to occur. While this makes them available to mediate spontaneous and serendipitous casual interactions, it also makes them more expensive to install and operate. A socially transparent video media space allows people to readily employ familiar social protocols for negotiating and conducting face-to-face casual encounters (Fish, Kraut, & Chalfonte, 1990). As a result, initiating casual interactions feels natural and lightweight. Social transparency requires high quality video links (section 2.3) that are more expensive to operate. While the combination of ubiquity (there when you need it) and social transparency (ease and naturalness of use) yields ready-to-handedness, achieving this design goal comes at a cost. The different architectures models mediate this ready-to-handedness/cost trade-off by varying persistency and directedness. This forms the basis of the evaluations of the models I make in the descriptions below.

## 2.2.2 Community space model: connecting social spaces

**Community space:** Persistent/undirected. Links between common rooms allow walk-up-and-use interactions of a predominately social nature between people in those rooms.

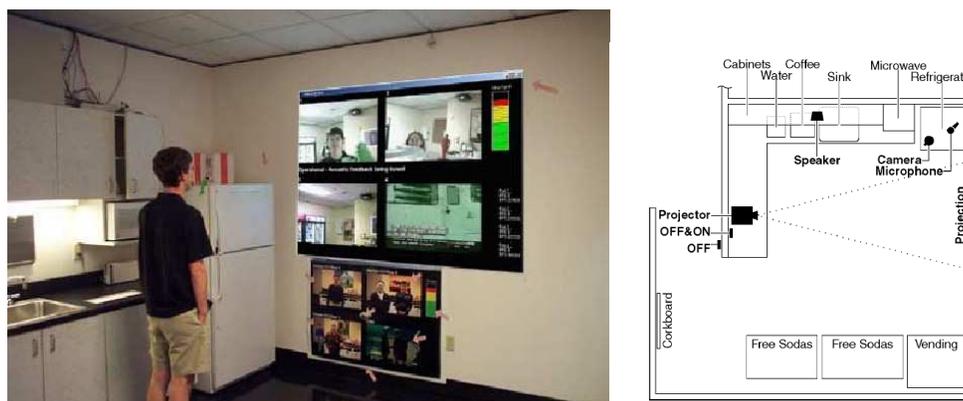


Figure 2.1 The Microsoft vKITCHEN video media space (from Jancke et al, 2001).

The earliest work in video media spaces consisted of attempts to extend social interactions across distances with video using the Community space model. The model extends social places—lounges, coffee rooms, and so forth—across distances. For example, the Xerox PORTLAND media space—one of the first media spaces—had its nodes installed in common areas that acted as hubs for foot traffic (Olson & Bly, 1991). Bellcore Lab’s VIDEOWINDOW system connected coffee rooms on two different floors of a building with large video displays, placing special emphasis on using near-life sized high-fidelity analog video with microphone arrays to make the technology social transparent (Fish, Kraut, & Chalfonte, 1990). In this section, I will use the Microsoft vKITCHEN (Jancke et al, 2001) system as a later example of a video media space designed using the Community space architecture model.

Microsoft vKITCHEN used near-life sized digital video links between four kitchens of different floors and buildings of the Microsoft Research campus. A typical vKITCHEN layout is pictured in Figure 2.1. At the time vKITCHEN was deployed, the entire research division had recently moved to new buildings. The old buildings had cavernous stairwells through which foot traffic from all departments flowed. Chance encounters in these stairwells were frequent. The newer buildings lacked any sort of similar central nexus and the move to the new buildings decreased serendipitous and spontaneous casual interactions. vKITCHEN was deployed to offset this decrease.

Common rooms, like kitchens (vKITCHEN) and lounges (VIDEOWINDOW, Xerox PORTLAND) are usually chosen as the venue for Community space model media spaces because both the locations and the systems are intended to support serendipitous social

encounters. However, the penalty for this placement is low use. For example, more than half the time when a person was in one vKITCHEN the other three were empty. The VIDEOWINDOW system had a similar problem despite the fact that designers moved the mailboxes of some 50 participants and set out free cookies and candies within the camera's field of view as incentives to encourage trips to the coffee room where it was installed. The designers of VIDEOWINDOW summed the situation up as: "The centralized nature of VIDEOWINDOW results in a sampling problem.... People are more likely, on average, to be someplace other than where the VIDEOWINDOW system is." That is, VIDEOWINDOW and vKITCHEN were not ubiquitous enough to be ready-to-hand for casual interactions.

A lack of ubiquity does not fully account for the use observed, though. Even if there was an opportunity to interact, it was rarely utilised. For example, in self-report survey data regarding vKITCHEN use, 81% of respondents said they noticed someone in another kitchen, but only 21% said they gestured to them and just 4% said they spoke with them. While a live CNN television news link was piped into the kitchens to serve as an impetus for conversation, it seemed to have minimal impact. These low use figures are hard to reconcile with the fact that 78% of vKITCHEN users reported a need for increased casual interactions and 87% reported that technology has a role in mediating them. While the VIDEOWINDOW system designers cited technological problems in achieving social transparency as a cause for the similar low use they observed, most vKITCHEN users found video quality in their system acceptable. Even though there were some audio problems, they were minor: 72% of users reported that improved audio/video quality would only somewhat have changed their experience or not at all.

While it seems people did not use these Community space model video media spaces, the cause is not due to the video quality. Instead, there is a mismatch between the place in which the media space is installed, the people present there, and their interaction needs. Interestingly enough, the same criticism could be levied against the stairwell after which vKITCHEN is patterned and the coffee room in which VIDEOWINDOW was installed. Although there are no observations of the pattern of contact made in the stairwell or coffee room which could be used to draw comparisons, there is a critical difference between a common room as realised in physical space versus one realised in virtual space. As people go to or from the



Figure 2.2 The CAVECAT media space connected personal offices with conventional video media space node equipment (a) and miniaturised integrated units called HYDRA surrogates (b). Photo: Buxton (1997).

physical common room, others notice and may join them. It is in the corridors one walks and in the offices in which one stops along the way to the common room that conversation partners are found and the contexts for casual interactions in the common room are set. While Community space model media spaces might be available for use in the destination, they fail to provide contexts to predicate this use because they do not “leak out” to these other important spaces. These observations form the rationale for using directed links as per the Office share model, discussed next.

### 2.2.3 Office share model: Connecting personal workspaces

**Office share:** Persistent/directed. Links interconnecting a very small number of people to simulate sharing a single “virtual office.”

The University of Toronto CAVECAT video media space (Mantei et al, 1991) used the Office share model for media space architecture. Media space nodes were installed into the researchers’ own personal offices (Figure 2.2a shows two such offices). As a custom of use, directed links between offices were left on for extended durations. This combination of persistent and directed links between a small number of personal offices gave the feeling that the media space occupants were all sharing a big “virtual office.” Many media spaces use the Office share model, e.g., the RAVE (Gaver et al, 1992) and CRUISER (Cool et al, 1992) media spaces both provided Office share modes of operation. I will discuss these prototypes later, focusing here on CAVECAT.

The CAVECAT researchers found that their media space was ready-to-hand for work-focused casual interactions because it was situated in personal offices: the places where people

spend most of their time *working*. They also found that it supported social encounters just like the Community space model media spaces. CAVECAT included shared drawing and collaborative text editing groupware tools (CAVEDRAW and SASSE) and multi-point video links wherein a group of three or four people, each in his or her own personal office, could simultaneously see and hear all the others. Thus, CAVECAT supported both serendipitous encounters and pre-planned group meetings. The researchers found it easy and natural to engage in casual interactions with CAVECAT because the persistent links provided rich informal awareness cues. Video simultaneously provides a channel for casual interactions and a channel for informal awareness to support the micro-coordination of these interactions (Mantei et al, 1991).

In a later iteration, the researchers produced prototype video media space nodes that they miniaturised down to the small integrated camera display units see in Figure 2.3b. Called HYDRA units, these behaved as physical **surrogates** for a person: the display is a surrogate face, the camera is a surrogate eye, the microphone is a surrogate ear, and the speaker is a surrogate mouth. The motivation for HYDRA units was to increase the social transparency of group meetings in CAVECAT by preserving the spatial interpretability of gaze and eye contact. These are used as important cues in social protocols for regulating speaker/listener role exchanges. (It is difficult to evaluate the HYDRA surrogates' actual social transparency because it is unclear if they were actually deployed.)

Office share model media spaces support all three of Hackman's group goals discussed in section 2.1.1 (Olson & Bly, 1991). The availability of groupware tools support work focused on computer-based artefacts coupled with improved accessibility to one's co-workers increases the *capacity* (but not the *efficiency*) of group work (Hackman's production goal). The intimate (small group) nature of the video channel supports socially oriented casual conversations (Hackman's group maintenance goal). And, the video channel provides continuous peripheral informal awareness that provides a feeling of being a member of a group (Hackman's member support goal). In particular, Mantei et al (1991) reported that the persistent links of the CAVECAT media space heightened the sense of social co-presence and tele-proximity that users felt.

Experiences with the CAVECAT media space (Mantei et al, 1991) also underscore the unanticipated social meanings that are tied to video media space design and use. For example, camera placement affects perceived status: a camera placed above eye height looking down onto a person's face makes him or her look like a junior or a supplicant while a camera placed below eye level looking up at the face makes the person look like a superior or aggressive. The relative sizes of faces indicate interpersonal distance: if the video is zoomed in too close on a person's face, it can make viewers feel as though they are unnaturally near to the person. Also, the field of view becomes a surrogate for the viewer's standing or sitting position within the personal office. For example, casual interactions in personal offices are almost always negotiated with the visitor standing at the office doorway before receiving acknowledgement and being ushered into the office for the actual conversation. For this reason, Buxton (1997) mounted a media space node atop his office door so that this social protocol for negotiating into casual interactions was supported.

Office share model media spaces have historically been implemented using analog video that does not scale cheaply to large communities of casual interaction partners. But, there are also scalability problems with modern-day digital video. Although digital video cameras are so inexpensive and easy to acquire, they are far from ubiquitous in corporate offices and are of quite poor quality. If higher quality cameras and microphones are used, the costs are still high enough to be a barrier towards ubiquity.

While the intimacy—the heightened sense of social co-presence generated by Office share media spaces—matches the relationships one has with those peers whom he or she interacts with most frequently, there are many others in the distributed workplace whom one routinely engages in spontaneous and serendipitous casual interactions. The video media spaces that used the Hallway model for connectivity, discussed in the next section, use this model to address these scalability and serendipity concerns.

#### 2.2.4 Hallway model: Browsing the social environment

**Hallway model:** Intermittent/undirected. Probe random locations to maintain awareness, like peering into open offices while walking down a virtual hallway.

The Bellcore CRUISER media space prototype (Cool et al, 1992) installed nodes in personal offices like CAVECAT, but used mostly intermittent links. It had several modes of operation, but two in particular—random-glance and autocruise—implemented the Hallway model of media space architecture by establishing brief (1 or 3 second) intermittent reciprocal connections to random individuals.

Even though these links, while active, involve one node connecting to another I consider these kinds of links to be undirected. Although one person initiates them the system determines who receives them. In fact, the recipient of a glance or autocruise does not need to take any action in response to it. In one iteration of CRUISER, the system established a link between two random individuals who both happen to have, say, printed a document recently. Thus, these features give the feeling of wandering down the corridor peeking into others' open offices gathering awareness and serendipitously encountering others along the way. This solution provides a scalable way of browsing (or polling) a large social environment and maintaining awareness of a larger community.

There are problems with the Hallway model that are reflected in the extremely low usage of these features observed during field trials of CRUISER. In particular, negotiating transitions into interactivity were made awkward: normal social protocol was often violated and CRUISER was not socially transparent. First, the undirected nature of the links meant that people often connected with others with whom they shared no immediate or obvious motivation to interact. The “approach” of another person was sudden and unexpected. A person was seemingly instantaneously connected to another with a full reciprocal connection. Although reciprocity was enforced for privacy reasons (to discourage surreptitious surveillance), reciprocity paradoxically became a privacy problem. Although the recipient of a glance did not have to acknowledge the link to the system, the lack of a **graceful** and smooth approach made the reciprocal video link very distracting and disturbing. Finally, in the CRUISER user interface there was no way to capitalise on serendipity and turn a glance—arguably intended only for accruing awareness—into a semi-persistent link suitable for casual interaction.

Although the Hallway model scales well (can be made ubiquitous more cheaply) because of the use of intermittent links, a media space which uses only this model will not be ready-to-hand for mediating casual interactions. Although undirectedness is important in this model to

support spontaneous encounters, this lack of directedness is precisely what impairs social transparency when it comes time to transition into conversation. In the next section I discuss the Telephone model, which uses directed intermittent links with careful attention to fostering socially transparent approach when establishing an intermittent connection.

### 2.2.5 Telephone model: Making transitions smooth

**Telephone model:** Intermittent/directed. Establish a connection to a specific person just to engage in conversation, much like using the telephone.

Nearly all forms of telecommunication use intermittent links, initiated by one person, directed to another specific person. It is the model used by the telephone (from whence it gets its name) as well as email, instant messaging, desktop video conferencing, and most groupware applications. Although it has appeared in many media space prototypes (.e.g., RAVE and even CRUISER), the media space prototype that I will use as an example of this model is Sun Lab's MONTAGE system (Tang, Isaacs & Rua, 1994).

In the Telephone model, one person initiates a connection to some other specific person—to poll for awareness or to have a conversation—and then disconnects when finished. This is just like placing a call with the telephone. The immediate and obvious problems with this model are that callers won't know the availability of a callee until a call is made and, as evidenced by the CRUISER experience, the calls themselves can be disruptive. Consequently, MONTAGE placed special emphasis on making these transitions into and out of tightly coupled casual interactions graceful. MONTAGE faded video windows in and out on the display smoothly over a period of a few seconds and used supplemental unobtrusive audio cues to complement these transitions. This fade-in mechanism is like the ring of an incoming telephone call, but it is more graceful and it identifies the caller.

To solve the problem of knowing when to call, MONTAGE provided on-line presence and availability indicators just like status icons in instant messaging systems. Thus, while a CRUISER user had to establish a full video link in order to determine if a person was available to call, a MONTAGE user could inspect the presence of a potential interlocutor without disturbing him or her with a reciprocal video connection. A video connection could be

established to get even more pre-conversation informal awareness: the smooth fade-in would be less disruptive than the same thing in CRUISER. Unanswered calls in MONTAGE faded out without any user intervention. These features made MONTAGE more socially transparent which in turn made it more ready-to-hand for mediating casual interactions. (Recall that ready-to-hand here means that the system is both easily accessed at the time it is to be used and is so effortlessly and naturally used that it “disappears” from thought.)

The Telephone model also helps make a media space more ready-to-hand for work-focused casual interactions. Groupware tools like shared white boards and screen sharing facilities can be easily started and connected to the right individual once the video link is established, such as in the ITU-T series of recommendations and in many instant messenger systems. Although MONTAGE used strictly two-person audio/video links (like a real telephone) it integrated desktop video presentation facilities to support multiparty discussions.

To increase ready-to-handedness by increasing ubiquity, MONTAGE was targeted for the next generation of desktop computers using digital video atop low-cost Ethernet networks. MONTAGE seriously stressed the limits of this class of hardware. Frequent software failures and long delays when establishing an intermittent video connection often stymied successful use of the system for spontaneous collaboration-focused interactions and diminished the ready-to-handedness of the system.

### 2.2.6 Community place: combining office spaces into a virtual communal space

The four interaction models discussed thus far can be combined. For example, the RAVE media space had several modes of operation (Gaver et al, 1992). The “Background” mode implemented the Community space model; the “Sweep” mode implemented the Hallway model; the “Office Share” mode implemented the Office share model; and the “Vphone” mode implemented the Telephone model. However, these modes were mutually exclusive: one could not activate the “Sweep” mode to for awareness of others while in a conversation with someone using a “Vphone” connection.

In this section, I wish to talk about a fifth media space architecture model that is completely different from the previous four.

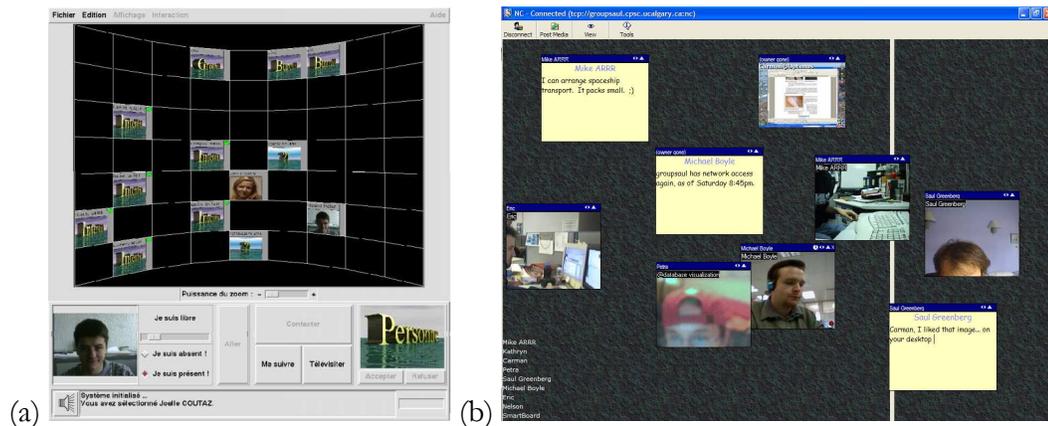


Figure 2.3 In the media space COMEDI (a), video snapshots for awareness are arranged on a perspective wall, while a hi-fi video link (bottom part) supports private 1-on-1 interactions. In the NOTIFICATION COLLAGE (b), snapshot video links for awareness and are shown beside “media items” for text- and artefact-based interactions, creating a Community place.

**Community place:** Video from personal offices is posted into a virtual communal space along with tools and artefacts for collaboration. The result is a locale (site and means) supporting informal awareness and casual interactions among members of a small distributed community.

The PORTHOLES system (Dourish & Bly, 1992; Lee, Girgensohn, & Schlueter, 1997) was an early example of this model of media space architecture. PORTHOLES was designed with the understanding that the episodes marking significant availability changes range in duration from a few seconds to many minutes or even hours. It displays **snapshot video**—with very low frame rates like 5 seconds per frame and longer—aggregated from a large community of sources. This technique scales very cheaply across the Internet. To support interaction, PORTHOLES integrates with a two-way desktop audio/video conferencing system. The media space COMEDI (Coutaz et al, 1998) shown in Figure 2.3a is a later iteration of the concept. In it, the video snapshots are aggregated using a perspective wall visualisation. The user may click on a snapshot to “take it down from the wall” and turn it into a real-time interactive video channel. This lightweight mechanism makes transitions from video-mediated awareness to audio/video-mediated interactivity smooth.

There are important differences between PORTHOLES and COMEDI versus a Telephone model system like MONTAGE. First, PORTHOLES uses persistent video for awareness, where as MONTAGE uses either non-video status icons or intermittent video for awareness. Second, the awareness displays in MONTAGE are designed to be visible at the time one wishes to strike up a conversation with another. In PORTHOLES, however, the video-mediated awareness display is designed to be peripherally visible at all times.

The NOTIFICATION COLLAGE system (Figure 2.3b) is similar to PORTHOLES in that it aggregates the display of undirected persistent snapshot video links for awareness over a community. NOTIFICATION COLLAGE users post automatically updating video snapshots to a shared “bulletin board.” These snapshots provide video-mediated informal awareness cues. Users also post into the same virtual space artefacts for interaction: e.g., text notes, web page thumbnails, desktop snapshots, photos, and more. In this way, interactions mediated with NOTIFICATION COLLAGE become undirected and are fundamentally public (multiparty) to the community. This is a significant distinction from PORTHOLES and COMEDI, in which conversations are strictly private (two-person). Thus, the NOTIFICATION COLLAGE system supports social protocols in which a third person overhears and subsequently joins in on an existing conversation.

PORTHOLES, COMEDI, and NOTIFICATION COLLAGE do not tax the hardware/software resources of commodity desktop computer systems. Moreover, they support heterogeneous setups in which some users have video equipment but others not. As a result, these systems can be made more cheaply ubiquitous (and thus more ready-to-hand for mediating casual interactions) than most of the previously discussed systems. Like MONTAGE and other Telephone model media spaces, social transparency for ready-to-handedness depends on how well the “approach” matches existing social protocol for negotiating casual interactions and also the ability for new protocols to be developed around it. There is clear indication that these systems are quite ready-to-hand. For example, many PORTHOLES users would use the system to check the availability of *co-located* colleagues before walking down the hall to meet them. Also, one PORTHOLES user remarked in field interviews that he had the feeling of “knowing” colleagues a continent away—recognising personal characteristics and habits—even though he had never met them.

### 2.2.7 Non-video systems for awareness and interactivity

As mentioned in Section 2.1, video is a very rich means of capturing and rendering awareness information and mediating interactions. Despite advances in technology that have substantially lowered the cost of powerful multimedia desktop computers, digital video cameras, and higher-capacity network links, it is still true that video—and in particular high quality video—is a comparatively costly means for supporting distributed collaboration. In this section, I briefly describe two lower-cost non-video mediated approaches to informal awareness and casual interaction.

In one approach, ambient cues are used to render awareness cues as unobtrusive characteristics of the surrounding physical environment. For example, Dahley, Wisneski & Ishii (1998) project awareness cues like presence and activity as subtly shifting light patterns cast using water lamps and pinwheels. These abstract and ambient awareness displays support the continuous, dynamic, lightweight, unobtrusive and peripheral aspects of awareness listed in Section 2.1.

The problem with this approach is that the devices that present awareness and the devices people use to control and mediate casual interactions are different. There is a gap in interaction and cognition and this gap makes the transition into and out of interactivity abrupt. This impairs ready-to-handedness and runs contrarily to the value of informal awareness described in section 2.1: that it helps make such transitions smooth and graceful. To help mediate transitions from awareness into interactivity, Kuzuoka & Greenberg (2000) used physical surrogates to display low-fidelity awareness cues and augmented the surrogates with sensors that form a tangible user interface for initiating a conversation. Alone, these surrogates do not make use of existing social protocols for conversation establishment but are sufficiently lightweight that it might be possible for people to easily develop new protocols incorporating them. When these surrogates were combined with a video media space, they became part of a ready-to-hand user interface for regulating media space activity that allowed existing social protocols for conversation regulation to be transferred to the media space. For example, the proximity sensors built into the ACTIVE HYDRA (Kuzuoka & Greenberg, 2000) surrogate allowed people to use interpersonal distance as a cue to control audio or video recording and display.

In a second approach, instant messaging (IM) systems couple real-time text chat and file sharing facilities with simple on-line presence indicators. Presence may be an offline/online/idle indicator estimated by measuring computer workstation input idle time and logon status, or a richer flag that conveys a greater sense of activity and availability, such as “busy”, “out to lunch”, or “on the phone.” Display names can often be re-appropriated to add further awareness information such as location, activity, psychological or physiological state, recent events, future plans, vent frustrations or engage in humour. Many modern systems also include the ability to set a display picture for awareness or the ability to see live video via a webcam. Users can visually scan their lists of IM contacts (buddies) and quickly get a sense of who is around and available. Many systems also provide unobtrusive notifications for when contacts sign in or become available. These notifications are typically presented as audio alerts (as in RAVE and MONTAGE) combined with small message windows called “toasts” that briefly appear and then disappear. Much as glances in CRUISER and MONTAGE, these toasts do not require any sort of explicit user acknowledgement.

These simple awareness facilities make it easy to identify opportune times and partners for casual interactions. As well, the text chat channel provided is particularly suited to casual interactions as they have been characterised in Section 2.1. Reading and writing instant messages and transferring computer files are almost effortless. This helps make IM ready-to-hand for most work- and social-oriented casual interactions. Emoticons, winks, display pictures, and other evolving features add a measure of fun and encourage free intermixing of both work and social topics. Because the system is slightly inaccurate in showing presence, **plausible deniability** (Nardi, Whittaker & Bradner, 2000) allows recipients the liberty to respond to incoming messages on their own terms. Instantaneous line-by-line transmission of messages and the lack of spelling or grammar checkers keep the tone informal. Instant messenger clients also act as hubs for all sorts of CSCW application add-ons: e-mail, on-line calendaring, mobile messaging, desktop videoconferencing, shared whiteboards, screen sharing, and telephony. Most can be added with only modest effort.

Instant messengers target the lowest common technological denominator and this increases their ubiquity. The awareness and interactions they provide are not as rich as video. While they have proven themselves in practice to be generally sufficient, there is evidence that

people seek the richer awareness and interactivity of video. For example, present-day IM clients now incorporate webcam sharing facilities. This development underscores the value of video as discussed in section 2.1.5.

### 2.2.8 Summary

This section sampled the history of video media space design focusing on how different models of media space architecture—patterns for media space design, if you will—address ready-to-handedness, ubiquity, social transparency, and costs-of-scale. This survey helps answer research question #1 because much of the knowledge we have about the low-level technological factors relevant to video media space design emerges from the experiences of the researchers who developed the systems discussed here.

- Video media space nodes and links are costly to install and operate [Section 2.2 introduction]
- Different models of media space architecture are produced by varying the persistency and directedness of the media space video links [Section 2.2.1].
- Video media spaces are designed to be ready-to-hand. Ubiquity and social transparency help make media spaces ready-to-hand [Section 2.2.1].
- Community space model media spaces lack contexts to predicate their use [Section 2.2.2].
- Office share model media spaces do not scale cheaply to large communities and are not appropriate for all relationships one might have with others in a large community [Section 2.2.3].
- Hallway model media spaces break social transparency during interaction negotiation [Section 2.2.4].
- Telephone model media spaces lack rich video-mediated awareness [Section 2.2.5].
- Community place model media spaces scale cheaply at a slight cost to the richness of video-mediated awareness [Section 2.2.6].
- Instant Messengers for informal awareness and casual interactions are increasingly incorporating video media space-like features [Section 2.2.7].

## 2.3 Technological barriers

The characteristics of informal awareness and casual interactions described in section 2.1 motivate ready-to-handedness through ubiquity and social transparency as design goals for video media spaces. But, as described in section 2.2, designers are forced to trade ubiquity off for social transparency (and vice versa) because there are costs associated with the installation and operation of video media space nodes and links. These costs are related to technological constraints on the quality and number of video links that can be handled concurrently in a video media space. Lifting these constraints to increase the scalability of a media space is a very costly measure and it is necessary to design within and around them.

### 2.3.1 Bandwidth as a cost/scalability constraint

The principal constraint on the cost-efficient scalability of video media spaces is network *bandwidth*. The term ‘bandwidth’ can refer to the total capacity of the telecommunications channels for signalling data or the required capacity needed for a video link. For example, a typical connection used by a business to access the global Internet might offer 1536 kilobits per second (kbps) bandwidth in each direction while a typical video media space link might require 384 kbps bandwidth in each direction<sup>1</sup>. Under such conditions, the network has enough bandwidth to support just four concurrent video media space links.

In the previous section, I discussed two technical factors that affect the ready-to-handedness of a video media space: persistency and directedness. These factors directly affect the bandwidth required by a video media space. Persistent links waste available bandwidth on links that might not be in use. Intermittent links allow available bandwidth to be shared among more nodes by deactivating nodes not in use. Directed links require more bandwidth simply because more links are needed to support a given community.

---

<sup>1</sup> A T-1 line provides 1536 kbps bandwidth in each direction. A CIF-sized video signal (352×288 pixels) at 30 frames per second (full motion) and voice-only audio can be compressed down to 384 kbps with the widely-used H.264 (video) and G.728 (audio) codecs with good quality.

In this section I will discuss other technical factors that affect the ready-to-handedness of a video media space. Some of the factors determine the bandwidth required by a video media space. These can be “tweaked” by designers to the lower bandwidth of an individual link and thereby trade some social transparency for ubiquity. Other factors however are not under designer control but also affect ready-to-handedness. In presenting this discussion in this section, I complete this chapter’s treatment of research question #1 discussed in Chapter 1, which concerns the low-level technical factors that must be considered when building a media space.

### 2.3.2 Visibility factors: Image size, field of view, resolution, and compression quality

Technical factors of a video media space can be adjusted its designers to lower bandwidth requirements, but these adjustments have significant consequences to the ready-to-handedness of the media space.

**Image size** refers to the dimensions of the rectangular array of pixels which represents a computer image: each pixel encodes the colour of a small area of the image. The more pixels in an image, the greater the bandwidth required to transmit it. Image size is mostly determined by the design and manufacturing of the light sensor used inside the camera. Typical upper bounds on a camera’s image size are  $640 \times 480$  or  $720 \times 480$  pixels (the later being used in digital camcorders and VCRs). However, a typical image size used in video media spaces is the CIF size:  $352 \times 288$  pixels.

**Field of view** refers to the area of the environment will be visible in the image. Field of view is usually fixed by the lens system used in the camera. While some expensive cameras have built-in pan and zoom features, the inexpensive “webcam” cameras used in video media spaces are stationary. While webcams are small enough that they can be picked up and moved around by their owners to change the field of view, they are usually tethered to either a computer or, in the case of wireless webcams, to an electrical outlet. Camera placement is a critical factor to consider because the camera’s field of view is so much less than the human eye’s that the camera will be “blind” to a wealth of peripheral informal awareness cues.

In addition to the problems reported in the CAVECAT experience (section 2.2.3), camera placement determines which parts of a person's body are visible when he or she is seated or standing and where in the field of view they are located. These things carry social meanings: for example, a person seen far off to the edge of the field of view may be perceived as evasive. Similarly, a camera mounted atop a monitor placed in front of a bookshelf may create an uncomfortably close view of a colleague's chest when she stands up and reaches to the bookshelf to retrieve something. A fixed-placement camera is simply unable to move around people as they move around in the space the camera views.

Field of view and image size factor significantly in achieving social transparency. A narrow field of view may break reciprocity: the characteristic of most face-to-face interactions such that in order for Person *A* to see Person *B*, Person *B* must also be able to see Person *A*. Without reciprocity, social transparency might be lost because mutual sighting afforded by reciprocity is a prerequisite in typical social protocol for initiating interactions. The reciprocity breakdowns due to constrained field of view in the VIDEOWINDOW system had a measurably negative impact on system use. Fish, Kraut & Chalfonte (1990) report cases where attempts to start conversations failed because although one person could see another at a remote site he was not reciprocally visible. The target for conversation was unaware of the other's presence and thus could not correctly interpret and respond to the attempt to start a conversation.

The two factors together determine the **resolution** of an image, i.e., the level of detail at which objects in an image will be visible. Assuming image size is constant (which it often is), video media space designers and end-users can trade off field of view against resolution. A wide field of view—even at the expense of resolution—might be preferred during loosely coupled work while a narrow field of view (focusing on the face, for example) might be preferred during conversation. If the resolution is too low, subtle informal awareness and conversation awareness cues like gaze or facial expressions might not be clearly discernable. Figure 2.4a shows an image and Figure 2.4b shows the same image, but with 80% fewer pixels leading to a substantial decrease in resolution. In (b) gaze is harder to discern, as are workspace artefacts held in the woman's hands. Recall from section 2.1.5 that these visible behaviours are vital to social transparency and also form the main motivation for using video as a medium for distance-separated collaboration. For this reason, video media space designers often seek to



Figure 2.4 Bandwidth requirements of high-quality video (a) can be reduced by decreasing the number of pixels in an image (b) or by using compression (c), but these methods affect social transparency.

employ video that has a large image size so that there is plenty of room to trade off field of view against resolution in achieving social transparency.

There are sophisticated audio and video compression algorithms that can greatly decrease the bandwidth requirements of a video link by discarding some visual information. Ideally, the differences between compressed and uncompressed video will be imperceptible. For example, the H.264 video codec (compressor-decompressor) is widely used in video telecommunications. It takes advantage of “deficiencies” in the human eye to discard visual information that humans are incapable of seeing well. It tracks the video over time using motion estimation to approximate scene changes and remove minor or redundant changes. Finally, the codec uses compact variable-length representations of an image (different from the rectangular array of pixels) that allow it to eliminate redundancy and also allow the video media space designer to specify peak and sustained bandwidth constraints. If the constrained bandwidth is too low, visible information (usually the details) will be discarded. This loss introduces errors—called compression artifacts—in image quality that can affect social transparency. Figure 2.5c shows the original image JPEG compressed at 20% quality. Important facial features, particular the fine features of the mouth visible in (a) are obscured by compression artefacts in (c).

### 2.3.3 Smoothness factors: Frame rate, jitter, and latency

**Frame rate** refers to how quickly the video display is updated. Motion video is produced by showing a sequence of frames (i.e., still images)—each captured at a slightly different time—many times per second. Frame rate is the number of frames displayed in a second. Reducing frame rate lowers the required bandwidth. For example, the NTSC television standard used in

Japan and North America prescribes a frame rate of about 30 interlaced frames per second (fps). In interlaced video, each image is divided into even- and odd-numbered lines drawn alternately at double-frame rate speeds. Thus it takes two  $1/60^{\text{th}}$  of a second intervals to show a complete 525-line NTSC frame: in the first  $1/60^{\text{th}}$  of a second, the even numbered lines are drawn on the display while in the second  $1/60^{\text{th}}$  of a second, the odd numbered lines are drawn. Interlacing is a way of making an apparent increase in frame rate without increasing bandwidth.

When the frame rate is high, the contents of the streaming of individual frames are perceived by the human visual sensory system as smooth and fluid motion. If the frame rate is too low, motion appears jerky. Small or rapid movements as in facial expressions or changes in gaze may be missed altogether. The jerky ‘animation’ of bodily movements may seem jarring such as in the case of lips synchronised with audio. For this reason, during casual interactions, most video media spaces strive to offer the highest frame rate the underlying technology will support. Commonly understood thresholds for frame rate are a minimum of 25 fps for “full-motion” video and a minimum of 12 fps for smooth conversations. Tang & Isaacs (1993) found 5 fps to be the lowest “usable”. As explained in section 2.2.6, PORTHOLES used extremely low frame rate video, on the order of  $0.01\sim 0.1$  fps, to provide low-bandwidth video-mediated informal awareness. With frame rates this low, a poorly timed video frame—for example, one that captures the participant in her office even though she was on her way out—remains visible for an extended duration that might not give an accurate picture of her present availability. This is analogous to plausible deniability in instant messenger systems (section 2.2.7).

In packet-switched networks, like today’s global Internet, a continuous video stream is broken up into small packets of data which are routed through the network individually. This scheme allows many computers to share network bandwidth, but causes available bandwidth to fluctuate from moment to moment according to use. To deal with these fluctuations in available bandwidth, many video-mediated communication applications that use the global Internet for interconnectivity use adaptive frame rates. Sometimes frame rate reductions are the result of the video system’s purposeful “dropping” of frames but often the network will

discard packets before they have reached the receiver: the video media space might just skip over the discarded video frame or it might attempt to resend it.

**Jitter** refers to tiny differences in the delay between frames. Packets of video—each flowing roughly independent of the others in a packet-switched network—freely intermix with other traffic for email, file transfers, and so forth. Jitter results as network congestion rises and falls with use and a variable number of packets get intermixed. Jitter in video channels makes motion appear irregular and jitter in audio channels produces stuttering, warbled speech. Jitter can be smoothed out by holding (buffering) audio samples or video frames for a while before playback to ensure they are displayed at a constant rate. These jitter buffers, however, introduce a delay between the time the audio/video is recorded and the time it is played back on the other end of the video link. There are, however, other sources of delay in the transmission of digital audio/video.

**Latency** is the total time involved in the capture, compression, transmission, decompression, and rendering of audio or video data. Some latency is due to the fact that electrical and optical signals take time to travel through copper wire, radio waves, and fibre optic cabling. There is latency due to packet switches and routers used in digital networks. Furthermore, a computer can only process compressed or decompress and display audio and video at a finite speed, introducing more latency.

Latency and jitter are related. When dealing with jitter in broadcast media, ensuring smooth playback is well worth the ten or fifteen second additional wait required to fill the jitter buffer. For interactive video to support spontaneous interactions, it seek to eliminate video jitter so long as there is no perceivable impact on audio jitter or latency. Inadequacies in the video channel are tolerated much more than inadequacies in the audio channel because people can adapt their conversations in the absence of a reliable video channel so long as it does not interfere with the basic mechanics of talking and listening (Tang, Isaacs & Rua, 1994). Audio latencies beyond 60~100 ms make casual interactions difficult because it is too hard to smoothly time speaker/listener role exchanges. People adapt by reverting to more formal (explicit) **floor control**: procedures for regulating the conversation through turn-taking.

There are also special latencies that occur only during connection establishment. If these latencies are long it becomes very hard to use the system to support spontaneous and serendipitous casual interactions. This is a significant problem for media spaces that use the Hallway, Telephone or Community place interaction models because these models rely on intermittent connections for casual interactions. Connection-time latencies interfere with social protocols for negotiating contact, reducing transparency and make the system less ready-to-hand for mediating casual interactions.

### 2.3.4 Audio problems: Sound level and echo

Audio is fraught with a number of problems. Mantei et al (1991) discuss the problem of sound levels. Office design, furniture, microphone placement relative to seating, microphone sensitivity and loudspeaker volume are all factors that weigh in on determining how loudly a person will be heard by others. Audio levels may differ from site to site: how loudly a person hears others does not indicate how loudly he or she may be reciprocally heard by them. But, people often expect audio quality problems to be reciprocal. They naturally raise their voices to talk with someone they cannot hear (even though that person may hear them just fine) and it further exacerbates audio level problems (Mantei et al, 1991).

Echo is a pernicious problem for full-duplex audio conferences. Sounds from one site played at the other site get picked up and sent back to the source site. If the echo is quite loud the speaker will naturally slow down his or her speech to an unnatural rate to overcome the confusion caused by it. Listeners, who might not be in a position to hear the echo, may then be left to wonder why the pace of conversation has slowed down!

The VKITCHEN media space is an example of one in which echo was a serious technical challenge. It connected relatively small rooms with hard, non-porous walls that exacerbated echo problems. The VKITCHEN designers also had to deal with persistent sources of sound like air-conditioning fans and refrigerator pumps. The echo from these kinds of sources never fades and can eventually crescendo into a painfully loud source of feedback that renders the audio link an environmental hazard. In earlier systems this type of echo cancellation required expensive high-speed hardware switches. In VKITCHEN these problems were overcome with

very sophisticated software echo cancellation algorithms that introduced only negligible latencies.

### 2.3.5 Summary

This section contributes to answering research question #1 by characterising the factors that are relevant to technological support for informal awareness and casual interactions in a video media space. Previous research experience with video media spaces has illustrated that these factors are extremely critical, affecting the success of a video media space. A privacy-preserving video media space must work within the bounds established by these limits and balance these low-level technological factors to support both informal awareness and casual interaction while respecting participants' privacy.

- Bandwidth is a significant factor limiting the ability of a video media space to support casual interactions [Section 2.3.1].
- Technical factors that determine the “visibility” of people, artefacts, and their environment in a media space can be adjusted (with constraints) to reduce bandwidth requirements, but these adjustments can affect casual interactions [Sections 2.3.2].
- Technical factors that determine the “smoothness” of actions and interactions can have great consequences on ready-to-handedness but are often beyond designers' control [Section 2.3.3].
- Technical factors related to audio can severely impair ready-to-handedness and are often the hardest to diagnose and solve [Section 2.3.4]

## 2.4 Conclusion

Chapter 1 stated, as part of the overarching research goal for this thesis, the following problem:

**Research Question #1:** What low-level technological factors need to be considered when building a privacy preserving video media space?

In this chapter I have begun the process of addressing this research question. I surveyed knowledge gained in over 15 years of CSCW research into video media spaces, casual

interactions, informal awareness, and the low-level technological factors that are critical to the design of video media spaces for distributed collaboration.

In Section 2.1, I discussed the motivation for always-on video media spaces as tools to support distributed collaboration. In particular, I explained how casual interactions are the backbone of everyday collaboration in co-located groups, and how always-on video is ideal for mediating informal awareness that helps people microcoordinate spontaneous casual interactions.

In Section 2.2, I surveyed a history of media space design and research. I discussed five different models of media space architecture meet objectives such as ready-to-handedness, ubiquity, and social transparency. I also discussed how designers must balance costs-of-scale with these design objectives.

Finally, in Section 2.3, I discussed important constraints on technology for video media spaces. In particular, I discussed how bandwidth constraints limit the ability of a video media space to support informal awareness and casual interactions. Some of these factors, like compression quality, can be tweaked by designers to reduce bandwidth at some cost to social transparency. Other important factors, like latency, are independent of bandwidth constraints and often cannot be controlled by designers.

This knowledge about low-level technology for video media space design provides the important context for the next chapter, which the COLLABRARY, a toolkit for rapidly prototyping these video media space design points.

## Chapter 3—The COLLABRARY toolkit

Chapter 1 asked:

**Research Question #1:** What low-level technological factors need to be considered when building a privacy preserving video media space?

In Chapter 2, I examined this question from the perspective of the video media space designer, emphasising critical factors related to architectures for media spaces and the characteristics of casual interactions and informal awareness that are the motivation for these architectures. In this chapter, I examine this question from the perspective of the video media space *developer*, i.e., the person who has the task of writing the software which runs the media space. This perspective is important as much of what has been learned about video media spaces, casual interactions, informal awareness, and (soon to be discussed) privacy comes directly from researchers' reflections and analysis of their own experiences living with the technology. The thesis problem and corresponding goal that I set out in Chapter 1 to provide this perspective are:

**Thesis Problem #1:** It is hard to rapidly develop video media spaces because the programmatic interfaces for multimedia are complicated and require considerable programmer effort and expertise.

**Thesis Goal #1:** Develop a toolkit to support the rapid prototyping of video media spaces and the distortion filtration method for preserving privacy therein.

To achieve this goal and solve this problem, I have produced the COLLABRARY toolkit. It is a Microsoft COM component library that facilitates the implementation of high-fidelity and

working system prototypes of video media spaces, and incorporate novel features intended to support privacy. This chapter will describe the COLLABRARY as a solution to thesis problem #1.

### 3.1 Toolkit objectives and motivation

I began the development of the COLLABRARY circa 1999 because at that time there was a lack of suitable toolkits which I could adapt to my research. I wanted a toolkit that would enable me to rapidly prototype different strategies for distorting video to preserve privacy in a video media space, such as one used in the distortion filtration study to be discussed in Chapter 4. Based on this, the following nine objectives emerged. Some, like #1, are self-evident; some, like #3 and #4, are general requirements for a toolkit; others, like #8, are peculiar to my own goals and purposes.

- Provide **access** to audio/video sources (e.g., cameras, microphones, disk files) and codecs.
- Be **compatible** with a conventional rapid application development tool and GUI toolkit, e.g., DHTML, Visual Basic, Tcl/Tk, or Java.
- Be **reliable** enough to be used in software for accomplishing real-world work tasks.
- Offer good **performance** on consumer-class Microsoft Windows 98/2000-based desktop personal computers using inexpensive ‘webcam’ digital video cameras.
- Be **accessible** so that an average programmer can quickly learn to use the toolkit.
- Be **lightweight** so that a programmer who is already experienced with the toolkit can rapidly create and deeply change programs that use the toolkit.
- Provide **n-way audio/video transmission** over local area networks (e.g., 100 Mbps Ethernet) and broadband Internet connections at to the home.
- Provide the end-programmer with **direct read-write access** to captured audio/video before it is compressed and transmitted.
- Be **flexible** and useful for prototyping unanticipated media space designs that I had not yet dreamed up.

### 3.1.1 The “missing” toolkit

Why build a multimedia toolkit? While several other toolkits were available, I faced important problems trying to use them. Two of the most promising ones were the Microsoft NETMEETING toolkit and the JAVA MEDIA FRAMEWORK.

The popular Microsoft NETMEETING (Microsoft Corporation, 1996) product included a toolkit for building atop it that met objectives #1 through #6 above. However, this toolkit provided only 2-way video conferencing (failing objective #7) and did not give end-programmer access to manipulate the audio/video (failing objective #8). Version 2.0 of the JAVA MEDIA FRAMEWORK (Sun Microsystems, 2002) was slated to meet all requirements, but was not expected to be released for over a year later.

Both of these toolkits are layered atop lower-level operating system APIs for audio/video capture and compression. While these APIs provide the greatest functionality and flexibility, they are not an appropriate solution. These APIs demand sophisticated C/C++ coding skills, require considerable effort to be used for even simple tasks, and are generally unavailable within rapid application development environments (failing objectives #2, #5, and #6).

### 3.1.2 Prototypes as design artefacts

In essence, I wanted a toolkit to be able to rapidly construct video media spaces prototypes, i.e., “artefacts that simulate or animate some but not all of the features of the system” (Dix et al, 1998). Prototypes vary in fidelity and purpose. A **low-fidelity** prototype might consist of ideas sketched in pencil on paper to quickly get a sense of the major pieces of the design. A **medium-fidelity** prototype might consist of a first-cut implementation of an important algorithm to evaluate performance and scalability factors. A **high-fidelity** prototype might consist of an extensive interactive user interface mock-up made with an animated vector graphics program that can be demonstrated to users, developers, and marketers to gather feedback and generate interest.

However, my needs are for **working system** prototypes: initial implementations of the system that can be deployed in to resilient users working in supportive situations. These

situations feature users (such as early adopters) who do not mind occasional glitches, crashes and restarts. They also utilise idealised hardware, software, and network platforms that behave within the parameters of the prototype, and expect secured networks (or benign social situations) so that security issues do not have to be considered. While the generalisability and overall deployability of a working system prototype does not match that which would be expected of the final system, the limited deployment is extremely valuable. It will help uncover subtle technical and social issues that are hard to detected unless the system is put to extended, real use. Buxton (1997) notes that working system media space prototypes permit “living with the technology” that is critical to identifying and solving the most pervasive and troublesome problems.

Within this setting, the COLLABRARY toolkit is intended to support medium fidelity prototypes of video distortion algorithms and working system prototypes of video media spaces that use those algorithms.

### 3.1.3 Prototyping toolkits

Prototyping toolkits must pay special attention to balancing objectives #5 (accessibility), #6 (lightweightness) and #9 (flexibility) above. Greenberg (in press) argues for the need for easy to program toolkits for novel interface areas: “By removing low-level implementation burdens and supplying appropriate building blocks, toolkits give people a ‘language’ to think about these new interfaces, which in turn allows them to concentrate on creative designs.” In particular they must follow advice that “simple things should be simple, and complex things should be possible” attributed to Alan Kay.

With a lightweight and flexible toolkit, the designer can iterate through and examine many designs, both promising and disappointing. In doing so, the designer gains superior knowledge about the design problem, and this should lead to superior designs. The designer must feel free to discard the prototype or make substantial changes to it without regretting the time already spent on it. Similarly, when the design area is complex like that of a privacy-supporting video media space, experimental features must be quickly implemented in a reasonably robust for so that they can evaluated either by trying it out, by actually using it in context, or by conducting formal studies.

Toolkits make prototyping less effortful by providing the developer with conveniently reusable bundles of functionality common to many applications. Toolkits are invaluable for rapid prototyping because they take care of the dirty work in software programming by implementing tedious or mundane aspects of the software and any complex functionality that is hard to implement reliably with good performance. This functionality is wrapped up in an application programming interface (API) which, in the case of the COLLABRARY, consist of a “class library” and a few administrative service applications.

Good toolkits not only speed up development, but also promote a good final product by steering designers and developers towards design patterns that have been established as successful. Most of what is known about the design of useful and usable toolkits comes by example. Consequently, I patterned the COLLABRARY after such toolkits as GROUPKIT (Roseman & Greenberg, 1996), ELVIN (Fitzpatrick et al, 2002) and the GUI toolkits that are native to Microsoft Visual Basic.

### 3.1.4 The remainder of this chapter: The COLLABRARY in action

In the following three sections, I explain the COLLABRARY by example, where I implement a prototype video media space and prototype several experimental features to support privacy. As we will see, the COLLABRARY offers two core sets of features to the media space programmer: multimedia capture and display (especially video), and the ability to distribute multimedia to various machines.

- Section 3.2 covers multimedia capture and playback. In this section, I build a blurry “mirror” in which video and audio captured on one computer is played back on the same computer but the video is blurred.
- Section 3.3 covers multimedia groupware programming. In this section I build a complete working  $n$ -way video media space prototype.
- Section 3.4 covers multimedia analysis and manipulation. In this section I prototype four experimental video effects that might be interesting or useful in a privacy preserving video media space.

```

1 using System.Drawing;
2 using System.Windows.Forms;
3 using Collabrary;

4 class MainForm : Form {
5     Collabrary.Camera camera=new Collabrary.CameraClass();
6     Collabrary.Microphone mic=new Collabrary.MicrophoneClass();
7     PictureBox pBox=new PictureBox();
8     Collabrary.Speaker spkr=new Collabrary.SpeakerClass();

9     MainForm() {
10        pBox.Dock=DockStyle.Fill;
11        pBox.SizeMode=PictureBoxSizeMode.StretchImage;
12        this.Controls.Add(pBox);
13        camera.Captured+=new CameraEvents_CapturedEventHandler(camera_Captured);
14        mic.Captured+=new MicrophoneEvents_CapturedEventHandler(mic_Captured);
15        camera.Size=CameraSizeStyle.CIF;
16        camera.FrameRate=15;
17        mic.Recording=true;
18    }

19    void camera_Captured(Collabrary.IPhoto Frame) {
20        Frame.Distort(0.5, PhotoDistortStyle.Blur);
21        pBox.Image=Image.FromHbitmap(Frame.Hbitmap);
22    }

23    void mic_Captured(Collabrary.IWaveform samples) {
24        spkr.Play(samples);
25    }

26    [STAThread] static void Main() {
27        Application.Run(new MainForm());
28    }
29 }

```



Figure 3.1—C# source code (and screenshot of resulting application) for a “blurry mirror” that demonstrates how simple multimedia capture and playback is with the COLLABRARY.

## 3.2 Multimedia capture and playback

Capturing and rendering multimedia is a vital programming task in a media space. In Figure 3.1, I show a complete C# program that illustrates the COLLABRARY being used to “mirror” live video and “echo” live audio by playing both back instantly on the same computer that captures them. This trivial example demonstrates two key concepts present in the COLLABRARY and several ways in which the COLLABRARY meets some of the nine requirements I had for it.

### 3.2.1 Hardware abstractions

The `COLLABRARY` provides the end-programmer with succinctly-named classes that encapsulate high-level abstractions of multimedia hardware. In the figure, video is captured by a `Collabratory.Camera` object (line 5) and audio is captured by a `Collabratory.Microphone` object (line 6). Video is rendered by a system-supplied `PictureBox` GUI widget (line 7). Audio is rendered by the `Collabratory.Speaker` object (line 8). By providing these classes, the `COLLABRARY` meets objectives #1 (access to multimedia capture sources) and #2 (compatibility with a rapid application development platform). The program runs indefinitely without error or aberrant performance (objective #3) and easily sustains the standard videoconferencing video fidelity 352×288 pixel video at 15 fps with minimal CPU usage on even low-end machines (objective #4).

When using the lower-level operating system APIs for multimedia capture, end-programmers work with the devices themselves. The abstractions used in the `COLLABRARY` make with the devices accessible and lightweight (objectives #6 and #7) by removing unnecessary programming complexity while at the same time adding robustness. First, the program will run without exception even if there is no camera attached to the computer. In such cases, the `Collabratory.Camera` object automatically inserts a ‘test pattern’ image in place of live video. Secondly, the code continues to function even if a ‘plug and play’ camera is detached, and automatically connects as soon as a new camera is attached. Third, multiple instances of the `Collabratory.Camera` object (e.g., multiple copies of this program running on the same computer) can simultaneously share access to the same camera device. (Similar functionality for the `Microphone` class was not implemented due to a lack of time.)

To make “simple things simple,” the configuration of most capture properties is made optional by giving them useful default values. For example, on line 16 of the figure, the `Camera.FrameRate` property is set to 15 fps. It must be configured to trigger automatic frame capture. On line 15 the `Camera.Size` property is set to CIF size (352×288 pixels). However, it could have been left unconfigured, in which case it uses QCIF size as a default. The only configuration property on the `Microphone` that is adjusted is a flag to enable recording (line 16). Because of this, `COLLABRARY` objects require little initialisation beyond instantiation before they may be used: properties are changed only if the programmer wishes a non-default

behaviour. Also, there is no ‘shutdown’ or ‘cleanup’ code in the example given: all COLLABRARY objects gracefully release held resources when they are garbage collected.

These niceties are important ways in which the COLLABRARY seeks to make the basic task of multimedia capture accessible and lightweight (objectives #6 and #7). They are not often found in other toolkits. For example, when using the JAVA MEDIA FRAMEWORK (an example of a toolkit that is more akin to the lower-level operating system APIs) programmers must write careful code: to handle error conditions during capture; migration between cameras; and graceful relinquishment of resources during shutdown. Failure to do this cleanup triggers exceptions that are hard to debug. Yet, this extra code contributes little to the real work of the prototype. As well, some things are not possible, with JMF. For example, the JMF does not support concurrent access to the same camera from different processes unless the underlying OS driver supports it. Also, the JMF does not handle ‘plug and play’ cameras with graceful failover to a ‘test pattern’ video sequence on disconnect, or automatically migrate to a different camera on reconnect.

### 3.2.2 Event-oriented architecture

The COLLABRARY manages multimedia capture using the same asynchronous programming paradigm that the underlying GUI toolkit uses for handling user input. When multimedia is captured by a COLLABRARY object, the object “raises an event.” The end-programmer can attach a callback method to handle the event.

Going back to the example in Figure 3.1, these event handlers for the video camera and microphone are attached on lines 13 and 14 of the `MainForm` constructor. The `camera_Captured` method handles the `Camera.Captured` event (lines 19~22), and is invoked each time a video frame is captured. The captured frame of video is wrapped in a `Collabratory.IPhoto` object that is passed as a parameter to the event handler (line 19). In this example, the video frame is distorted using a blur distortion filter. The `Photo.Distort` method used here (line 20) is merely an example manipulation that can be performed on the video. In the final iteration of the prototype, I will discuss other kinds of manipulations the Collabratory makes possible. Finally, the system-supplied `Image.FromHbitmap` method is used to convert the

video frame data into a form that can be displayed in the `PictureBox` widget (line 21). This is how video from the local camera is mirrored to the GUI display.

Audio is treated in a similar manner. The `mic_Captured` (lines 23~25) method handles the `Microphone.Captured` event. Periodically, after collecting a small block of audio data (by default, every 50 ms of audio) the `Microphone` object will raise its `Captured` event. The `mic_Captured` event handler method takes the block of digital audio data (wrapped in a `Collabrary.IWaveform` object) and sends it to the `Speaker.Play` method for playback (line 24). This is how audio from the microphone is echoed to the speaker. (For brevity, I omit an example audio manipulation.)

The flow of multimedia information from the camera and microphone (i.e., capture source) to the GUI display or speaker (i.e., render sink) forms a **pipeline**. The event-oriented pipeline architecture used in the `COLLABRARY` differs from the stream-oriented pipeline architecture that is most commonly used in multimedia APIs. Take, for example, the `JAVA MEDIA FRAMEWORK`. To use the `JMF` to create the same sort of blurry mirror in this iteration a programmer does the following:

1. Create a pipeline manager.
2. Create and configure the sources and add them to the pipeline. When capturing live media this will involve enumerating the available capture sources.
3. Create and configure an effect filter that will perform the blur. Add it to the pipeline.
4. Create and configure sinks for the audio and video and add them to the pipeline.
5. Instruct the pipeline manager to finalise the construction of the pipeline, adding needed codecs.
6. Tell the pipeline manager to start the pipeline, which runs autonomously.
7. When finished, tell the pipeline manager to stop the pipeline.

With a `JMF`-style stream-oriented pipeline, the pipeline manager frees the programmer of the burden of moving data from source, through filters and codecs, to sinks. While programmers using the `COLLABRARY` have to write code to do the same, it is fairly simple because of the event-oriented architecture.

A significant feature of a stream-oriented pipeline, however, is that media that flows through the pipeline is stamped with timestamps that come from a reference clock that is common to all streams in the pipeline. The pipeline manager uses these timestamps to ensure that audio and video streams are tightly synchronised: the end-programmer does not need to worry about this difficult task. The event-driven architecture in the COLLABRARY does not support synchronisation.

When using a stream-oriented architecture like the one in the JMF, it is easy to connect sources to sinks (objective #1), but it is difficult to manipulate the media itself (objective #9) when there is no ready-to-use effect filter to perform the manipulation desired. This is because writing a filter is conceptually difficult. Metaphorically, the filter is like a piece of a large puzzle that must be carefully shaped to fit the hole available for it. The programmer must create a class that implements interfaces required by the pipeline manager. These interfaces are extremely generalised, however: both audio and video media types are presented as byte arrays instead of a rich image or audio type. This makes it awkward to draw raster graphics or compose images within a filter and the programmer is forced to implement mundane code that is irrelevant to the real work of the filter. In contrast, the event-oriented pipeline architecture used in the COLLABRARY makes precisely these kinds of manipulations easy to accomplish.

### 3.3 Multimedia groupware programming

In this section, I develop a fully-functional video media space prototype that could be deployed to up to around eight resilient media space participants. While the previous iteration focused on illustrating the basics of multimedia capture and rendering with the COLLABRARY, this iteration focuses on multimedia transmission between distributed processes using the COLLABRARY shared dictionary system. The interface is shown in Figure 3.2 as it appears on screen. A large outer window holds video of the local participant and smaller inner video windows for each other participant. All 120 lines of C# code for this prototype, shown in Figure 3.3, illustrate a number of important concepts that I elaborate upon in this section.



Figure 3.2—User interface for the basic media space prototype developed in Section 3.3.

### 3.3.1 Centralised server network architecture

The transmission of audio/video and application program data is done using the *Collabratory shared dictionary system*. This system provides a centrally-coordinated data store. Data caches are placed at each client for rapid access. To the end-programmer, this data store looks like a hash table that maps hierarchically structured keys—‘/’ delimited text strings resembling paths in a conventional disk file system—to values. But, there is an important difference: multiple clients can connect to the same shared dictionary hosted on a server, where any client can add, change, or remove any of the entries in it, concurrently. Behind the end-programmer’s API, a communication infrastructure ensures that changes are serialised and replicated to all clients of a common shared dictionary server. This shared dictionary architecture incorporates elements from many sources: groupware toolkits e.g., GROUPKIT (Roseman & Greenberg, 1996); distributed shared memory systems e.g., LINDA TUPLESPACES (Carriero & Gelernter, 1989); and notification servers e.g., ELVIN (Fitzpatrick et al, 2000).

The `Collabratory.SharedDictionary` object maintains the connection to the server, a local cache of the data store, and a list of the other computers connected to the same shared dictionary server. In the `COLLABRATORY` API, each shared dictionary client connection gives itself a globally unique identifier string. Servers are identified by URLs (e.g., `tcp://localhost:mediaspace`).

52

```
1 using System;
2 using System.Collections;
3 using System.Drawing;
4 using System.Windows.Forms;
5 using Collabrary;

6 class InnerForm : Form {
7     PictureBox pBox=new PictureBox();
8     Collabrary.JPEG jpeg=new JPEGClass();
9     Collabrary.Speaker spkr=new SpeakerClass();
10    Collabrary.Subscription audio;
11    Collabrary.Subscription video;
12    Collabrary.Subscription name;

13    public InnerForm(SharedDictionary sd,
14        string clientId) {
15        this.Activated+=new
16            EventHandler(InnerForm_Activated);
17        this.Deactivate+=new
18            EventHandler(InnerForm_Deactivate);
19        this.ClientSize=new Size(200,200);
20        pBox.Dock=DockStyle.Fill;
21        pBox.SizeMode=...StretchImage;
22        this.Controls.Add(pBox);
23        audio=sd.Subscribe(clientId+"/audio") as
24            Collabrary.Subscription;
25        audio.Quench=true;
26        audio.Notified+=new
27            ..._NotifiedEventHandler(audio_Notified);
28        video=sd.Subscribe(clientId+"/video") as
29            Collabrary.Subscription;
30        video.Notified+=new
31            ..._NotifiedEventHandler(video_Notified);
32        name=sd.Subscribe(clientId+"/name") as
33            Collabrary.Subscription;
```

```
26    name.Notified+=new
27        ..._NotifiedEventHandler(name_Notified);
28    this.Text=sd[clientId+"/name"] as string;
29 }

30 void audio_Notified(string Key, string clientId,
31    object val, SubscriptionNotifyStyle reason,
32    object prev) {
33    spkr.Play(val as Collabrary.IWaveform);
34 }

35 void video_Notified(string Key, string clientId,
36    object val, SubscriptionNotifyStyle reason,
37    object prev) {
38    Collabrary.IPhoto Frame=jpeg.Decompress(val as
39    Collabrary.Buffer);
40    pBox.Image=Image.FromHbitmap(Frame.Hbitmap);
41 }

42 void name_Notified(string Key, string clientId,
43    object val, SubscriptionNotifyStyle reason,
44    object prev) {
45    this.Text=val as string;
46 }

47 void InnerForm_Activated(object sender,EventArgs e){
48    audio.Quench=false;
49 }

50 void InnerForm_Deactivate(object sender,
51    EventArgs e){
52    audio.Quench=true;
53 }
54 }

55 class OuterForm : Form {
56    PictureBox pBox=new PictureBox();
```

```

48 Collabrary.Camera camera=new CameraClass();
49 Collabrary.Microphone mic=new MicrophoneClass();
50 Collabrary.JPEG jpeg=new Collabrary.JPEGClass();
51 Collabrary.SharedDictionary sd=new
    ...SharedDictionaryClass();

52 OuterForm() {
53     this.Activated+=new
        EventHandler(OuterForm_Activated);
54     this.Deactivate+=new
        EventHandler(OuterForm_Deactivate);
55     pBox.BorderStyle=BorderStyle.None;
56     pBox.Location=new Point(10,10);
57     pBox.Size=new Size(camera.Width,camera.Height);
58     this.Controls.Add(pBox);
59     this.IsMdiContainer=true;
60     camera.Captured+=new
        ..._CapturedEventHandler(camera_Captured);
61     mic.Captured+=new
        ..._CapturedEventHandler(mic_Captured);
62     Collabrary.CommonUI cui=new
        Collabrary.CommonUIClass();
63     this.Text=cui.PromptForString("Enter your "+
        "name:", "anonymous", "Media Space",
        PromptForStringDialogStyle.SingleLine);
64     string url=sd.PromptForUrl(
        "tcp://localhost:mediaspace");
65     if(""==url) {
66         Application.Exit();
67     }
68     sd.Opened+=new ..._OpenedEventHandler(sd_Opened);
69     sd.Closed+=new ..._ClosedEventHandler(sd_Closed);
70     sd.Entered+=new
        ..._EnteredEventHandler(sd_Entered);
71     sd.Exited+=new ..._ExitedEventHandler(sd_Exited);

72     sd.Quench=true;
73     sd.Open(url);
74 }

75 void sd_Opened(string Url, bool isServer) {
76     sd[sd.Me+"/.transient"]=sd.Me;
77     sd[sd.Me+"/name"]=this.Text;
78     foreach(SDItem e in sd.Instances["*"]...) {
79         sd_Entered(e.Value as string);
80     }
81     camera.FrameRate=5;
82     mic.Recording=true;
83 }

84 void sd_Closed(string Url, bool isServer, bool
    graceful, ConnectionStatus prevStatus,
    string details, out int retries) {
85     mic.Recording=false;
86     camera.FrameRate=0;
87     foreach(SDItem e in sd.Tags...) {
88         (e.Value as InnerForm).Close();
89     }
90     sd.Tags["*"]=null;
91     if(sd.Troubleshoot("Media Space",Url, graceful,
        isServer, prevStatus, details)) {
92         retries=4;
93     } else {
94         retries=0;
95         this.Close();
96     }
97 }

```

Figure 3.3—Complete C# source code for an *n*-way video media space working system.

```

98 void sd_Entered(string clientId) {
99     InnerForm inner=new InnerForm(sd, clientId);
100     inner.MdiParent=this;
101     inner.Show();
102     sd.Tags[clientId]=inner;
103 }

104 void sd_Exited(string clientId) {
105     InnerForm inner=sd.Tags[clientId] as InnerForm;
106     inner.Close();
107     sd.Tags[clientId]=null;
108 }

109 void camera_Captured(Collabrary.IPhoto Frame) {
110     Frame.Distort(0.5F, PhotoDistortStyle.Blur);
111     sd[sd.Me+"/video"]=jpeg.Compress(Frame);
112     pBox.Image=Image.FromHbitmap(Frame.Hbitmap);
113 }

114 void mic_Captured(Collabrary.IWaveform samples) {
115     sd.Signal(sd.Me+"/audio", samples, 100, 0, null);
116 }

117 [STAThread] static void Main() {
118     Application.Run(new OuterForm());
119 }
120 }

```

Figure 3.3—Complete C# source code for an  $n$ -way video media space working system. Some text has been abridged with ellipses for display purposes.

The shared dictionary system automatically deals with **late-comers**: the server provides each client at the time it connects with a completely up-to-date version of the data store. This is similar to what is done in the JAVA SHARED DATA TOOLKIT (Burrige, 2004). This approach is a more lightweight (objective #4) solution to the late-comer problem than, say, the ELVIN (Fitzpatrick et al, 2002), which requires that end-programmer write code to update a late-comer. The `SharedDictionary.Opened` event is raised on the client after it has connected and fully updated its local cache. In the figure, the `sd_Opened` event handler (lines 75~83) performs initialisation tasks to enable audio/video capture and bring the GUI display up-to-date.

When the connection is closed or broken due to a network connectivity problem, the end-programmer can handle the `SharedDictionary.Closed` event and set a flag to have the connection automatically re-established. In the source code, the first part of the `sd_Closed` event handler (lines 84~90) performs cleanup tasks to disable audio/video capture and remove the inner video windows that have been created for the other nodes. The second part (lines

91~97) uses the `SharedDictionary.Troubleshoot` method to display a message to the end-user to help them understand and correct an unexpected network connectivity problem. This convenience method provides a lightweight means for the end-programmer to deal with a common problem: notifying the end-user of connection troubles.

When a client connects to the shared dictionary server, the server informs the other clients already connected to it, and they in turn each raise the `SharedDictionary.Entered` event. An inner video window must be set up for each client as it connects in (lines 98~103). A reference to the video window is stored in the `sd.Tags` collection. This is an unshared data store provided by the `SharedDictionary` as a convenience for programmers to use they wish. When a client disconnects from the server, the other remaining clients are informed and raise their `SharedDictionary.Exited` events. In the source code, the `sd_Exited` event handler (lines 104~108) removes the inner video window from the display that corresponds to the departed client.

### 3.3.2 Organising and storing data in a hierarchical dictionary

As mentioned previously, entries in the shared dictionary are key/value pairs. The keys are simple text strings that look like file system paths. The values may be of virtually any type: the `COLLABRARY` automatically marshals the data i.e., converting it into a form that can be transmitted over a network. Automatic marshalling makes the shared dictionary system more accessible (objective #3: because novice programmers need not concern themselves with it), lightweight (objective #4: because expert programmers need not write any code to take care of it) and flexible (objective #9: because it allows the potential for an unlimited number of end-programmer data types to be shared).

A value is stored at a particular key using a simple assignment syntax: e.g., `sd["/name"]="Michael", sd["/age"]=28, sd["/video"]=camera.Frame`. This API has been designed to be accessible (objective #3: because it is the same syntax as that which is used with the system-supplied hash table class), lightweight (objective #4: because assignment is one of the simplest programming statements) and flexible (objective #9: because the end-programmer decides what keys are used and what values they will hold).

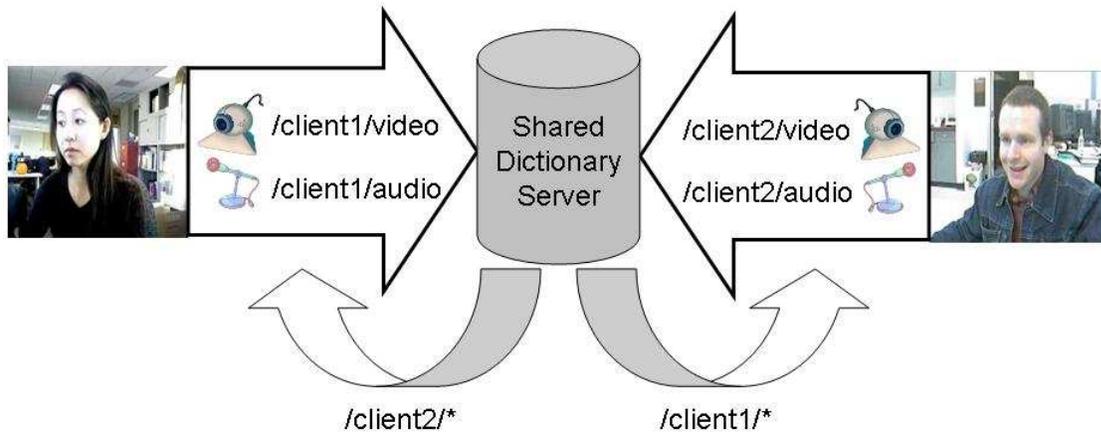


Figure 3.4—In the prototype video media space, clients post audio and video data to a shared dictionary server and receive notification of audio/video posted by others. The shared dictionary maps hierarchically organised string keys to values.

The shared dictionary supports hierarchical organisation of data even though a hash table is (conceptually) not a hierarchical data structure. Keys are made to look like paths in a disk file system, which is a hierarchical data structure. Related keys can be given a common prefix. For example, Figure 3.4 shows the audio and video keys for two participants in the prototype media space. Each is prefixed with a unique identifier for the client: e.g., `/client1/audio`, `/client1/video`, `/client2/audio`, `/client2/video`, and so forth. On line 111 of the source code, uses the `SharedDictionary.Me` property to retrieve its client-id and prefixes it to the `"/video"` substring to generate the complete key used to store its video frames.

### 3.3.3 Subscription notifications and Model-View-Controller architecture

The COLLABRARY shared dictionary system has a notification mechanism whereby the end-programmer can request notification of changes made to the dictionary. The end-programmer obtains a `Collabratory.Subscription` object, specifying a key or pattern of keys to watch, and handles the `Subscription.Notified` event on it. The simple pattern matching language available resembles the “filename globbing” pattern matching language used in UNIX and related disk file systems. (The code in the figure does not need to make use of pattern-based subscriptions.)

In the `InnerForm` constructor (lines 13~28) three subscriptions are used to handle the data coming in from each media space client: one for the node's video, one for its audio, and a third for its name. Note how the client-id of the node has been prefixed to the subscription patterns (lines 20, 23, and 25) and how, on line 27, it is used when retrieving the current value of the name. When a node transmits its audio or video, it sends the data to the server which in turn broadcasts it to all the other clients. When another client receives the data, it updates its local cache and then calls into the `Subscription` objects it has created that have matching keys or patterns. The end-programmer's `Subscription.Notified` event handler is invoked and passed various parameters that describe the change.

Video and audio are streamed by repeatedly storing individual video frames and audio blocks at the same key in the shared dictionary (line 111, for example). As each frame or block is received, the `audio` and `video` subscription notification event handlers (lines 29~31 and 32~35, respectively) render the media to the speaker or display. The whole process is not altogether strikingly different from that used to construct the “mirror” in the previous iteration. The primary difference is the inclusion of an intermediate step of storing data in the shared dictionary.

The ability to organise data hierarchically and receive asynchronous notification of data changes allows the end-programmer to employ the shared dictionary as the “model” within a Model-View-Controller or Presentation-Abstract-Control architecture pattern (Dix et al, 1993). These models are important because they allow the end-programmer to separate the abstract data model from how it is gathered (i.e., the input gathered by the controller) and how that data is displayed (via the view or presentation). This separation is critical in a distributed environment where different clients may have different views or different means of managing user input.

### 3.3.4 Controlling the presence and distribution of keys and values

The `COLLABRARY` shared dictionary system includes features to control how long keys or values stay in the shared dictionary. Normally, when a client puts a value in the shared dictionary, it is sent to all clients and it is stored in the dictionary indefinitely. It can be overwritten (by any client, not just the one that first put it there) by assigning a new value to

the same key. The entry will be removed when a client sets the key's value to `null`. A client receives a copy of all data on the server, regardless of whether it has obtained a subscription for it or otherwise expressed interest in it.

As a bandwidth economisation, the shared dictionary server may silently discard intermediate values. (This is the default behaviour and later in this section I will discuss how the end-programmer can override it.) When an entry in the dictionary is repeatedly updated with rather large blocks of data, it might generate more requests to transmit data than the network can handle. A queue of these updates can accumulate. If there is already an update waiting but not yet sent at the time an update for an entry is to be delivered to a client, the waiting update is discarded and only the newer of the two is actually sent. If this happens repeatedly every time an update is to be sent, heuristics are employed to ensure that an update is routinely sent without being dropped. The updates that are dropped may vary between clients: a client connected to the server by a fast link may receive every update while a client on a slower link may receive only some of the updates. Intermediate values are also discarded similarly by clients when generating requests to the server.

The default persistence and distribution behaviour is good for most purposes, but drawing upon observations given in Chapter 2 it may interfere with finding a trade-off between social transparency and bandwidth. It is known that audio should be given higher priority than video during interactions, thus the default data distribution scheme in the shared dictionary—which gives equal priority to all types of data—is inappropriate for a video media space. It is known that jitter in the audio channel is unacceptable while it is minimally tolerable in the video channel. While intermediate video frames could be dropped if absolutely necessary, audio data should (ideally) never be dropped.

The source code for this iteration of the video media space prototype demonstrates three ways in which the default persistence and data distribution behaviour in the shared dictionary may be changed to make the video media space working prototype more robust in lower-bandwidth network conditions.

First, the end-programmer can override the default intermediate-discard behaviour by using the `SharedDictionary.Signal` method to route data through the shared dictionary (as

though it was being stored) but without actually having the data persist. This method is used for transmitting audio on line 115 of the source code. The `Signal` method includes arguments that allow the programmer to specify the priority: the base priority for data stored using the assignment syntax shown earlier is 0 and outgoing data is sent on a highest-priority-first basis. In the figure, priority of 100 is used to ensure audio is sent ahead of video. Again, heuristics are applied to ensure updates for other keys are periodically sent.

Additional optional `Signal` method arguments specify timeouts for data delivery (i.e., the server discards the data if it cannot be sent before it times out), and a list of clients to which the data is specifically addressed. When this argument is supplied the `Signal` method supports broadcast, selective multicast, and unicast data distribution policies.

Second, the end-programmer can store in the shared dictionary specially named keys that modify how the server processes data in the shared dictionary. For example, on line 65, the client stores its own client-id in a key that ends with `.transient`. The `.transient` metadata key associates a client with a subtree of the dictionary such that the server will automatically remove all of the keys in the subtree when the associated client disconnects. This effectively binds the presence of data in the shared dictionary to the presence of a client. For example, when the client specified in the value for `/people/mike/.transient` disconnects from the server, the server will automatically remove `/people/mike` plus `/people/mike/*` (i.e., any descendent entries in the shared dictionary hierarchy). A variation of this type of behaviour is the `.lifetime` metadata key. It holds a value that specifies the duration of time (in seconds) that a subtree should remain on the server (not shown in this example).

Third, the end-programmer can change the way it receives data such that only data to which it has subscribed will be received. This saves on bandwidth: the server maintains a list of all the subscriptions held by a client and “quenches” or withholds notification of a change made to the data store from a client that has not subscribed to it. In this iteration, only audio subscriptions are quenched. The audio from a particular client is received only when the inner video window for that client is active. On line 72, the `SharedDictionary.Quench` flag is set to `true` to ask the server to turn on server-side subscription quenching for the current client. The client will not receive notification of a change to the data store unless it has a matching subscription. In the `InnerForm_Activated` and `InnerForm_Deactivate` GUI event handlers

(lines 39~41 and 42~45, respectively) the audio subscription for a client is enabled/disabled by toggling the `Subscription.Quench` flag. Setting this flag to `false` temporarily turns off the particular subscription's notifications without affecting other subscriptions. This quenching method is a simplification of the quenching used in the ELVIN notification service (Fitzpatrick et al, 2000).

The data persistence/distribution options discussed in this subsection are optional. If not used, the working prototype will continue to work. Thus, on the one hand, the introduction of these features does not change the accessibility of the toolkit (objective #3). However, a working system prototype built without the modifications described here will "fail" in two ways readily perceived by the media space end-user. First, audio may experience intolerable jitter and dropped packets because of the intermediate value-discard behaviour. Second, a working system prototype may behave erratically (spurious breakages in network connectivity, latency and audio/video de-synchronisation that increases over time) in low-bandwidth networking scenarios such as a home telecommuter connected via a broadband Internet connection.

## 3.4 Multimedia analysis and manipulation

The COLLABRARY is intended to support the rapid prototyping of video media spaces that incorporate novel features for supporting privacy. Until now, the focus has been on showing how the toolkit can be used to build a working video media space prototype that can serve as a platform for the exploration of designs to preserve privacy. In this section, I will describe two ways that the COLLABRARY supports the rapid prototyping of designs that might be useful in preserving privacy.

### 3.4.1 Provide rich, composable operations despite flawed implementations

The COLLABRARY provides a variety of methods to perform rudimentary analysis and manipulation of video. While the COLLABRARY provides comparatively fewer analysis and manipulation methods for audio, the concepts I discuss here still apply. The end-programmer composes these methods to very quickly implement and build upon known video effects. For

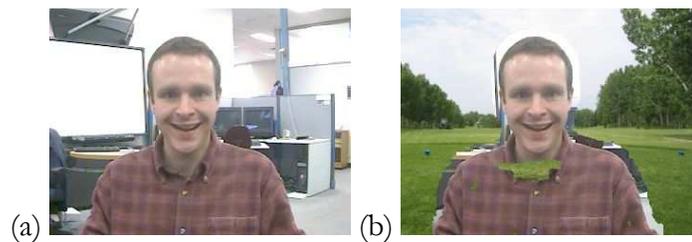


Figure 3.5—Background subtraction can be performed with just a single method call in the Collabrary, but the implementation is not fully robust. Part of the person’s body (a) has been misclassified as background in (b).

example, the source code in Section 3.2 demonstrated how a blur distortion filter may be applied to video by calling the `Photo.Distort` method.

#### 3.4.1.1 Example #1: Background subtraction

Some interesting operations are, in practice, very hard to implement robustly. Algorithms may be entirely unknown or prohibitively computationally expensive. Simplified implementations may not work correctly under all conditions. Background subtraction and replacement is a classic example of a computer vision-based image manipulation technique that is occasionally featured in prototype video media spaces as a means for preserving privacy (e.g., Junstrand, Keijer & Tollmar, 2001). However, it is very difficult to implement this technique for live video captured under naturalistic conditions and consequently it is usually ‘mocked up’ in demonstration videos or high-fidelity prototypes and typically excluded from working system prototypes.

The COLLABRARY includes a rudimentary implementation of background subtraction and replacement. It is demonstrated in Figure 3.5. The image in the figure clearly shows the video image as it appears originally (a) and after the background has been replaced (b). Obvious errors are visible in the result: part of the person’s face and clothing have been mistakenly replaced with the background image, while a “halo” of un-replaced background is visible around his shoulders. The lack of robustness is not the important point here: the important point is that this functionality has been implemented by the end-programmer with a single call to the `Photo.Subtract` method. With almost no effort the end-programmer can test “What would it look like if the background was replaced?” and communicate the idea to others.



Figure 3.6—Blurring can be easily combined with face tracking to produce interesting effects with the Collabrary.

#### 3.4.1.2 Example #2: Face-tracking blur

In Figure 3.6, I take the common-place technique of blurring video to preserve privacy and, inspired by the observation that a person’s face is an important signal of informal awareness cues, come up with the idea of blurring all of an image except a person’s face. This face-tracking-blur design idea can be quickly prototyped with the COLLABRARY. The COLLABRARY provides a `FaceTracker` class to implement the CAMSHIFT face-tracking algorithm (Bradski, 1998). The `Photo.PasteEx` method implements image composition with per-pixel alpha masking and blending. The output of these two methods can be combined by the end-programmer to very simply prototype a blur filter that blurs everywhere except within a bounding box rectangle around the person’s face.

#### 3.4.1.3 Example #3: Visual traces

In Figure 3.7, I show the last example I will use to illustrate the concept that the COLLABRARY provides rich functionality that can be rapidly composed by the end-programmer. This example does not focus on privacy preservation. Instead, it shows an adaptation of Gutwin’s (2002) idea of telepointer traces. Snapshots of the system in action are visually blended together to show a history of the interactions with the system. This concept can be easily applied to video using the `Photo.Paste` method to compose video frames with alpha blending.

In Figure 3.7, I have modified the traces idea in two ways. First, a snapshot is added to the history of recent frames if it differs markedly from the previous snapshot in the history. Thus, rather than using every frame from a fixed window, the modified traces concept I use here shows a history of major changes. Second, a scrolling EKG-like diagram that shows an approximation of the activity level in the video over time appears at the bottom of the display.

```

1 void Traces(Photo currentFrame) {
2 // Compute difference factor.
3 Photo temp=currentFrame.Copy(...) as Photo;
4 temp.Subtract(frames[frames.Length-1], null, null);
5 float psnr=temp.PSNR;
6 // Update composite traces image if major difference.
7 if(Math.Abs(psnr)>PSNR_THRESHOLD) {
8     Array.Copy(frames, 1, frames, 0, frames.Length-1);
9     frames[frames.Length - 1]=currentFrame.Copy() as Photo;
10    for(int i=0; i<frames.Length; i++) {
11        composite.Paste(frames[i], 0, 0, ComputeAlphaForPosition(i));
12    }
13 }
14 currentFrame.Paste(composite, 0, 0,
15     0.4F);
16 // Update history and draw chart.
17 Array.Copy(history, 1, history, 0,
18     history.Length-1);
19 history[history.Length-1]=sign*
20     Math.Min(1F, Math.Abs(psnr/40F));
21 sign*=-1;
22 Photo chart=new PhotoClass();
23 int height=currentFrame.Height/3,
24     halfHeight=height/2-8,
25     baseLine=halfHeight+4,
26     prevY=baseLine;
27 chart.BackgroundColor=chart.MakeColor(255,255,255);
28 chart.Resize(currentFrame.Width, height, false);
29 for(int i=1; i<history.Length; i++) {
30     int curY=baseLine+(int)(halfHeight*history[i]);
31     chart.DrawLine(i-1, prevY, i, curY, chart.MakeColor(255,0,0), 1);
32     prevY=curY;
33 }
34 currentFrame.Paste(chart, 0, currentFrame.Height-chart.Height, 0.6F);
35 }

```



Figure 3.7— C# code to implement Gutwin’s traces for video and an EKG-like activity history display and a representative image. (Contrast increased for print reproduction.)

On lines 4 and 5, I use the `Photo.Subtract` method and `Photo.PSNR` property to obtain a scalar floating point value that approaches negative infinity the more the current video frame differs (on a per-pixel basis using only the raw colour values) from the previous snapshot. The current frame is added to the history if this “difference value” exceeds a limit (lines 7~13 of the figure). A composite of all the video frames in the history is pasted (alpha-blended at 40% opacity) atop the current video frame. In the figure, note how a visitor, wearing a red sweater, is faintly visible.

(using wrapper API)	(using pointer-based API)
<pre> 1 void FleshOut(Photo frame) { 2   Photo hsv=frame.Copy(...) as Photo; 3   hsv.ColorSpace=...HSV; 4   for(int y=0; y&lt;frame.Height; y++) { 5     for(int x=0; x&lt;frame.Width; 6       x++) { 7       if(hsv[x,y,...hsvHue]&lt;20) { 8         frame[x,y,...A11]=0; 9       } } } 10  frame.Distort(0.8F, ...Blur); 11 } </pre>	<pre> 1 unsafe void FleshOut(Photo frame) { 2   IntPtr frameBP, hsvBP; 3   int len; 4   frame.LockMemory(out frameBP, 5     out len); 6   Photo hsv=frame.Copy(...) as Photo; 7   hsv.LockMemory(out hsvBP, 8     out len); 9   len=(((frame.Width*3+3)&gt;&gt;2)&lt;&lt;2); 10  hsv.ColorSpace=...HSV; 11  for(int y=hsv.Height-1;y&gt;=0;y--){ 12    byte* hP=((byte*)hsvBP)+y*len; 13    byte* fP=((byte*)frameBP)+y*len; 14    for(int x=hsv.Width-1;x&gt;=0;x--){ 15      if(hP[0]&lt;20) { 16        fP[0]=0; fP[1]=0; fP[2]=0; 17      } 18      hP+=3; fP+=3; 19    } } 20  frame.Distort(0.8F, ...Blur); 21  hsv.UnlockMemory(); 22  Marshal.ReleaseComObject(hsv); 23  frame.UnlockMemory(); 24 } </pre>



Figure 3.8—Custom video analysis/manipulation technique to remove flesh-coloured pixels.

Lines 19~30 of the figure show how the EKG-like activity chart is produced. The “difference values” are modified arithmetically to generate a sine wave-like sequence of positive and negative values, scaled between -1 and +1. This value approaches zero as there are fewer differences indicating less activity. (That is, the EKG-like display “flatlines” when there is no activity.) This sequence is used to generate the chart that is produced by using `Collabratory.Photo` methods for common raster graphics primitives, like drawing lines, circles, boxes, and text.

### 3.4.2 Direct media access for prototyping custom analyses and manipulations

The pre-packaged analysis and distortion algorithms discussed in the previous subsection make for a good starting point in the development of experimental privacy preservation techniques. For more power, the `COLLABRATORY` has been designed with the idea that custom analysis/distortion methods—those which do not utilise any of the pre-packaged methods—should also be possible. The `Photo` and `Waveform` objects, which act as containers for video

and audio data, provide methods whereby the end-programmer can gain direct read/write access to the multimedia data. This access is provided at a high-level pixels/samples granularity, a medium-level colour values/left-right channels granularity, and a low-level bits and bytes granularity.

In Figure 3.8, I show two implementations of an algorithm that searches a video frame for “flesh coloured” pixels and replaces them with black. This algorithm is greatly simplified: a pixel is deemed “flesh coloured” if its hue falls in a certain range. The first implementation utilises the `Photo` object’s convenient array-like access to the colour values in pixels. The second implementation obtains a pointer to the block of memory that the `Photo` object internally maintains to hold the pixels in the image. While pointer-based access is fastest, it is obviously more complex. It also permits the `COLLABRARY` to be used with external native libraries that the end-programmer may wish to use to implement advanced functionality.

This example illustrates the power that the `COLLABRARY` affords to the end-programmer trying to develop novel approaches to preserving privacy in video media spaces. The `COLLABRARY` makes this level of analysis and manipulation sufficiently lightweight and accessible that it is possible for even novices—programmers who have no computer vision background—to explore their ideas.

### 3.4.3 Summary of multimedia analysis and distortion

The examples in this section illustrate various ways that the `COLLABRARY` provides rich functionality that the end-programmer can easily combine in potentially interesting and useful ways for preserving privacy in a video media space. This approach gives novice programmers access to sophisticated technical algorithms and helps make the `COLLABRARY` very accessible (objective #5).

The `COLLABRARY` API is quite powerful. I have focused on providing powerful generic high- and low-level analysis and manipulation algorithms that can be combined in unlimited ways. This helps the `COLLABRARY` achieve the goal of being flexible (objective #9). This power also makes the `Collabrary` lightweight (objective #6). If the media space was built without the use of a prototyping toolkit (i.e., directly using the low-level operating system APIs) these manipulations would require hundreds of lines of complicated code. Most other rapid

prototyping toolkits do not provide direct access to the multimedia data (objective #9) or—as in the case of the JAVA MEDIA FRAMEWORK—use highly generalised APIs that demand the programmer invest effort writing code that merely converts generalised media types into richer objects that can support the kinds of operations demonstrated in this section.

These basic multimedia manipulations are provided even though the implementations are not always fully robust, e.g., as in the background subtraction example earlier. In other words, the COLLABRARY does not fully meet its reliability (objective #3) and performance (objective #4) objectives. While robustness is important for a final system, these unreliable implementations are still very useful for exploring ideas and communicating them to others. These implementations are acceptable in a working system prototype deployed to resilient users in robust scenarios. The approach I have taken in the COLLABRARY is to provide the functionality anyway: if it is compromised, it at least will serve as a “stub” that can be filled in later by a better implementation if the end-programmer wishes.

### 3.5 Toolkit features not demonstrated

The COLLABRARY toolkit includes scores of features not demonstrated here. Below, for completeness, I briefly enumerate some of this extra functionality. These features are useful for handling both real world needs (e.g., encryption for security purposes) as well as easing mundane operations (e.g., disk files for output).

— Shared dictionary features

- Encryption
- Concurrency control (using a token-passing algorithm)
- Logging: a program can be written to log everything that happens in the shared dictionary. This program uses specialised interfaces to access all shared dictionary network communication at the server.
- Bulk file transfers: e.g., an equivalent to FTP that allows the contents of large files to be uploaded and stored on the server for other clients to download later. Uploads and downloads restart automatically after connection errors.

- Utilities to administer and inspect shared dictionary servers

— Multimedia features

- Audio/video codecs (like MPEG-4)
- File I/O: read/write AVI (audio/video) and WAV (audio only) files with compression
- Audio manipulations: direct access, bandpass filtering, noise generation

### 3.6 Missing toolkit features

The COLLABRARY is missing a few significant features that are useful for the development of multimedia and groupware applications. Below, for completeness, I briefly enumerate the kinds of functionality not yet implemented in the COLLABRARY. While important, these features are far from critical, and their omission does not preclude the successful use of the COLLABRARY. They have not been implemented mostly due to time constraints.

— Multimedia features

- Audio/video stream synchronisation
- Standardised network transport for audio/video, e.g., ITU-G/H/T, RTP/RTSP
- Robust video segmentation algorithms

— General features

- Improved performance, scalability, and reliability
- Platform independence
- Tutorials and assistive documentation

### 3.7 Evaluating the COLLABRARY

Referring to the beginning of this chapter, the COLLABRARY was built to enable the rapid prototyping of different video distortion techniques that I could evaluate in a study (see Chapter 4). The COLLABRARY started out as the `Camera`, `Photo` and `MultimediaFile` (the latter

not being illustrated in this chapter) components: these were the only ones needed to accomplish my immediate goal of building systems reliable enough for a user study. Concurrent to the study, work on the COLLABRARY continued, and it grew to include audio and more video analyses and manipulations. However, at no point during this bottom-up approach to COLLABRARY development was a formal evaluation of the toolkit considered a goal.

The objectives discussed in Section 3.1 loosely guided the development of the COLLABRARY. I have used illustrations of the COLLABRARY in practice to discuss how its API meets (or fails to meet) these objectives. This is, in a sense, a kind of informal evaluation.

A more significant way that the COLLABRARY has been evaluated has been through everyday use by me and others. In the Interactions Lab at the University of Calgary, the toolkit has been a critical technology for numerous research projects and publications.

— Working media space prototypes

- NOTIFICATION COLLAGE (Greenberg & Rounding, 2001)
- COMMUNITY BAR (McEwan & Greenberg, 2005)
- HOME MEDIA SPACE (Neustaedter & Greenberg, 2003)
- VIDEOARMS (Tang, Neustaedter, & Greenberg, 2004)

The NOTIFICATION COLLAGE was used daily for more than three years. The COMMUNITY BAR is under active development and is also presently being used on a daily basis. These working system prototypes are robust enough that they can be deployed to medium sized groups (five to 10 people), which include home-based telecommuters (1, 5, 10 or even 100 km away) and colleagues in remote institutions (several thousands of kilometres away).

— MSc theses which were based on software dependent on COLLABRARY technology

- Neustaedter (2003)
- C. Tang (2003)
- Rounding (2004)
- A. Tang (2005)

— Studies evaluating using systems developed atop the COLLABRARY

- Neustaedter, Greenberg, & Boyle (2005)
- Boyle, Edwards, & Greenberg (2000)

Furthermore, the toolkit has been used to teach groupware programming in a senior undergraduate interaction design class and is used daily in countless small utility applications for automating mundane tasks like producing thumbnails of a large folder of files.

### 3.8 Conclusion

This chapter opened with the following thesis problem and goal statements:

**Thesis Problem #1:** It is hard to rapidly develop video media spaces because the programmatic interfaces for multimedia are complicated and require considerable programmer effort and expertise.

**Thesis Goal #1:** Develop a toolkit to support the rapid prototyping of video media spaces and the distortion filtration method for preserving privacy therein.

**Status of Goal #1:** Completed. In this chapter, I have presented the COLLABRARY. It is a toolkit that supports the rapid prototyping of video media spaces and of distortion effects that might be useful for preserving privacy. It satisfies thesis goal #1 and solves thesis problem #1 by providing programmatic interfaces that can be rapidly used by even novice programmers to build sophisticated video media space prototypes.

In Section 3.1, I outlined a set of nine objectives that I held in my mind as I built the COLLABRARY. These objectives were never formally specified: development on the COLLABRARY proceeded in a bottom-up fashion, growing incrementally to fit my needs.

I illustrated how the COLLABRARY meets these objects by showing it in practice. In Section 3.2, I first showed how the COLLABRARY can be used to capture and render audio and video by showing the source code needed to build a “blurry” mirror. The critical concepts that emerge from this illustration are that the COLLABRARY provides lightweight abstractions of

multimedia hardware and data that can be strung together via an accessible event-oriented pipeline architecture.

In Section 3.3, I showed how the COLLABRARY used to prototype a fully functional  $n$ -way video media space prototype. The critical concepts that emerged from this illustration are that the COLLABRARY provides a shared dictionary which gives programmers a means to flexibly organise and distribute multimedia information and that the shared dictionary's subscription-based notifications can be used to implement the Model-View-Controller architecture for GUI programming.

In Section 3.4, I used the COLLABRARY to rapidly prototype four video analysis and manipulation techniques that could be easily incorporated into the video media space prototype shown earlier. The critical concepts that emerged from this illustration are that the COLLABRARY provides a lightweight API for rich operations that can be composed by end-programmers to produce powerful techniques that analyse and manipulate video (and audio) for preserving privacy—despite the fact that the implementations in the COLLABRARY are not always robust. Also, the COLLABRARY supports the rapid development of custom analysis and manipulation methods since it provides direct read/write access to multimedia data.

In Section 3.5, I enumerated other toolkit features that are implemented in the COLLABRARY but are not illustrated elsewhere in this chapter. Section 3.6 enumerates useful toolkit features that were not implemented due mostly to a lack of time. Finally, in Section 3.7, I talk about the ways that the COLLABRARY has been put into use as a critical technology for numerous research projects and publications. The studies discussed in the next chapter, for example, are both published work that would not have been possible without the COLLABRARY.

## Chapter 4—Evaluating distortion filtration

This chapter discusses an evaluation of the distortion filtration technique as a means of balancing privacy and awareness in a video media space. The blurring used in the examples in Chapter 3 is an easy-to-implement distortion filtration technique. This chapter addresses the second of four thesis problems stated in Chapter 1:

**Thesis Problem #2:** It is widely suspected that distortion filtration may be useful for mitigating privacy issues in video media spaces but its usefulness has not been rigorously evaluated and there is no guidance as to how much filtration is ideal.

**Thesis Goal #2:** Determine if it is possible to use the distortion filtration technique to strike a balance between awareness and privacy in a video media space. If it is possible, determine at which levels a balance can be reached.

In this chapter I present the results of a semi-controlled user study I ran in collaboration with Chris Edwards and Saul Greenberg to achieve this goal. The results of this study have been published as Boyle, Edwards, & Greenberg (2000). I also summarise the results of a second study that Carman Neustaedter ran in collaboration with Saul Greenberg and me. This study was published as Neustaedter, Greenberg & Boyle (2005). The important contribution offered is a critical examination of a widely employed technique to help preserve privacy. The study suggests that the technique is unreliable in all but the most mundane cases of use.

To perform this study I identified useful independent and dependent subjective measures of privacy and awareness and developed a threshold-based analysis procedure for identifying useful ranges to which the distortion techniques may be applied. Many aspects of the experimental method developed and used here could be applied to other examinations of the trade-off between privacy and awareness.

#### 4.1 Why evaluate distortion filtration?

Distortion filtration, in its many forms, is a widely used technique. Television news broadcasts often ‘pixelize’ sensitive areas of the image: for example, a person’s face is replaced by large pixels (squares) to mask identity. Many researchers have investigated techniques for preserving privacy by altering what appears on the video e.g., Crowley, Coutaz, and Bérard (2000); Hudson & Smith (1996); Zhao & Stasko (1998); and Lee, Girgensohn, & Schlueter (1997) among many others. Standard image processing techniques include: full scene pixelization, altering a scene to show only edges, posterizing effects, blurring, Venetian blinds, and so on. Advanced techniques include a shadow-view filter that gives the effect of a ghostly shadow moving about a static scene (Hudson & Smith, 1996) and the use of eigen-filters to analyse a scene and reconstruct images of it in a ‘socially-correct’ form (Crowley, Coutaz & Bérard, 2000).

Although some of these filtering techniques have been deployed in a number of working video media space systems, there has been only one evaluation of how filters reveal or mask awareness information. Zhao & Stasko (1998) examined four filters on brief five-second video clips shown at two sizes (80×60 and 320×240 pixels). Their study asked volunteers to identify which of five of actors were featured in a clip and which of four activities the actors were engaged in. Volunteers were primed with key information ahead of time. They were shown portraits of the actors as well as the kinds of activities to look for. The results for activity recognition show that, with the exception of the shadow-view filter, all the filters tested supported high (>90%) activity recognition. Actor identity was more difficult to recognize, with only the pixelize filter and uniform lens combination able to support moderate (>75%) actor identity recognition levels. From a separate qualitative study that evaluated the

techniques in the context of a video media space application, it was found that actor identity recognition improves with familiarity.

While a good start, Zhao & Stasko's report has several limitations. First, the authors did not study the filters' effects on privacy. Thus it is difficult to tell if filters offer any advantage over an unfiltered scene. Second, the authors only investigated one level of these filters. Although the chosen level was likely reasonable, the authors did not specify how much filtration was applied. Most filters can be adjusted to give varying degrees of fidelity. An example of this is illustrated in Figure 4.2, where two filters (a blur filter and a pixelize filter) have been applied at nine different levels. At low fidelity levels, the filter masks a great deal of information. At high fidelity levels show many of the details in the image. One cannot really pre-judge a particular filter's ability to show awareness information while safeguarding privacy unless it is evaluated over a range of levels. Third, the scenes are not described; the reader is not told where cameras were positioned, how close actors were to it, and so on. Finally, participants were primed with all possible selections ahead of time. They only had to discriminate between a few choices instead of interpreting a scene.

With this study, we complement Zhao & Stasko's research by investigating the effects of a blur filter and a pixelize filter at various levels for their impact on both awareness and privacy.

## 4.2 Methodology

In this section we detail the experimental method used.

### 4.2.1 Hypothesis

The first null hypothesis considers the effect video filtering may have on a person's ability to extract awareness information from a scene, while the second null hypothesis considers the effect filtering may have on a person's perception of privacy afforded by that filter.

**Hypothesis #1:** There is no difference in a viewer's ability to correctly identify particular awareness cues from one of five different QCIF-size video scene sequences that have been altered by either a pixelize or

blur filter, where these filters were applied at ten levels ranging from heavily obscured to unfiltered effects.

**Hypothesis #2:** There is no difference in how a viewer rates these filtered video scenes in terms of how they protect privacy.

## 4.2.2 Materials: Video sequences

4.2.2.1 We created 95 video sequences by applying two filters at nine different filter levels to five different scenes (the 10<sup>th</sup> level is the unfiltered video). These are described below.

### 4.2.2.2 Scenes

We videotaped five video sequences portraying typical media space settings (Figure 4.1). Sequences were shot using a high-quality Canon XL-1 digital video camera. While we could have used a teleconferencing PC camera, its lesser quality could have compromised our study with inferior video; we also expect camera technology to improve over time. No special lighting conducive to filming was used, as we felt this would lead to unrealistic footage. As in many media spaces, some scenes are poorly lit e.g., backlighting, glare, etc. We converted sequences to a frame rate of 24 frames per second at standard QCIF videoconferencing image size (176×144 pixels, 24 bits per pixel), saved each as an AVI file suitable for replay on a PC.

### 4.2.2.3 Video filters

Using the components provided by the COLLABRARY toolkit presented in Chapter 3, we pre-processed each scene at nine levels of the pixelize and blur filters (Figure 4.2). Both filters use standard image processing algorithms, as summarized below.

The **pixelize filter** produces a “mosaic” effect across an image (lower images in each row of Figure 4.2). It works by re-sampling an image (a single video frame) using a coarser grid. The image is divided into rectangles of equal area, where all of the pixels in an area are reset to the mean of their original colours. The pixelize filter is computationally inexpensive and can be applied in real time to video stream.



**Scene 1 (42s):** a coffee room. A seated male actor reads a newspaper. A second male enters, gets coffee, talks to the first person while gesturing over the newspaper, and then leaves. A third female walks by the doorway.



**Scene 2 (25s):** personal office, side view. A male actor, glancing occasionally at his computer screen is eating a snack. He drinks from a can, then reaches over to grab and open a bag of potato chips.



**Scene 3 (31s):** personal office, view from doorway. A single male actor is seen talking on the telephone while looking at his computer screen. The clip ends just after the actor hangs up and starts working on his computer.



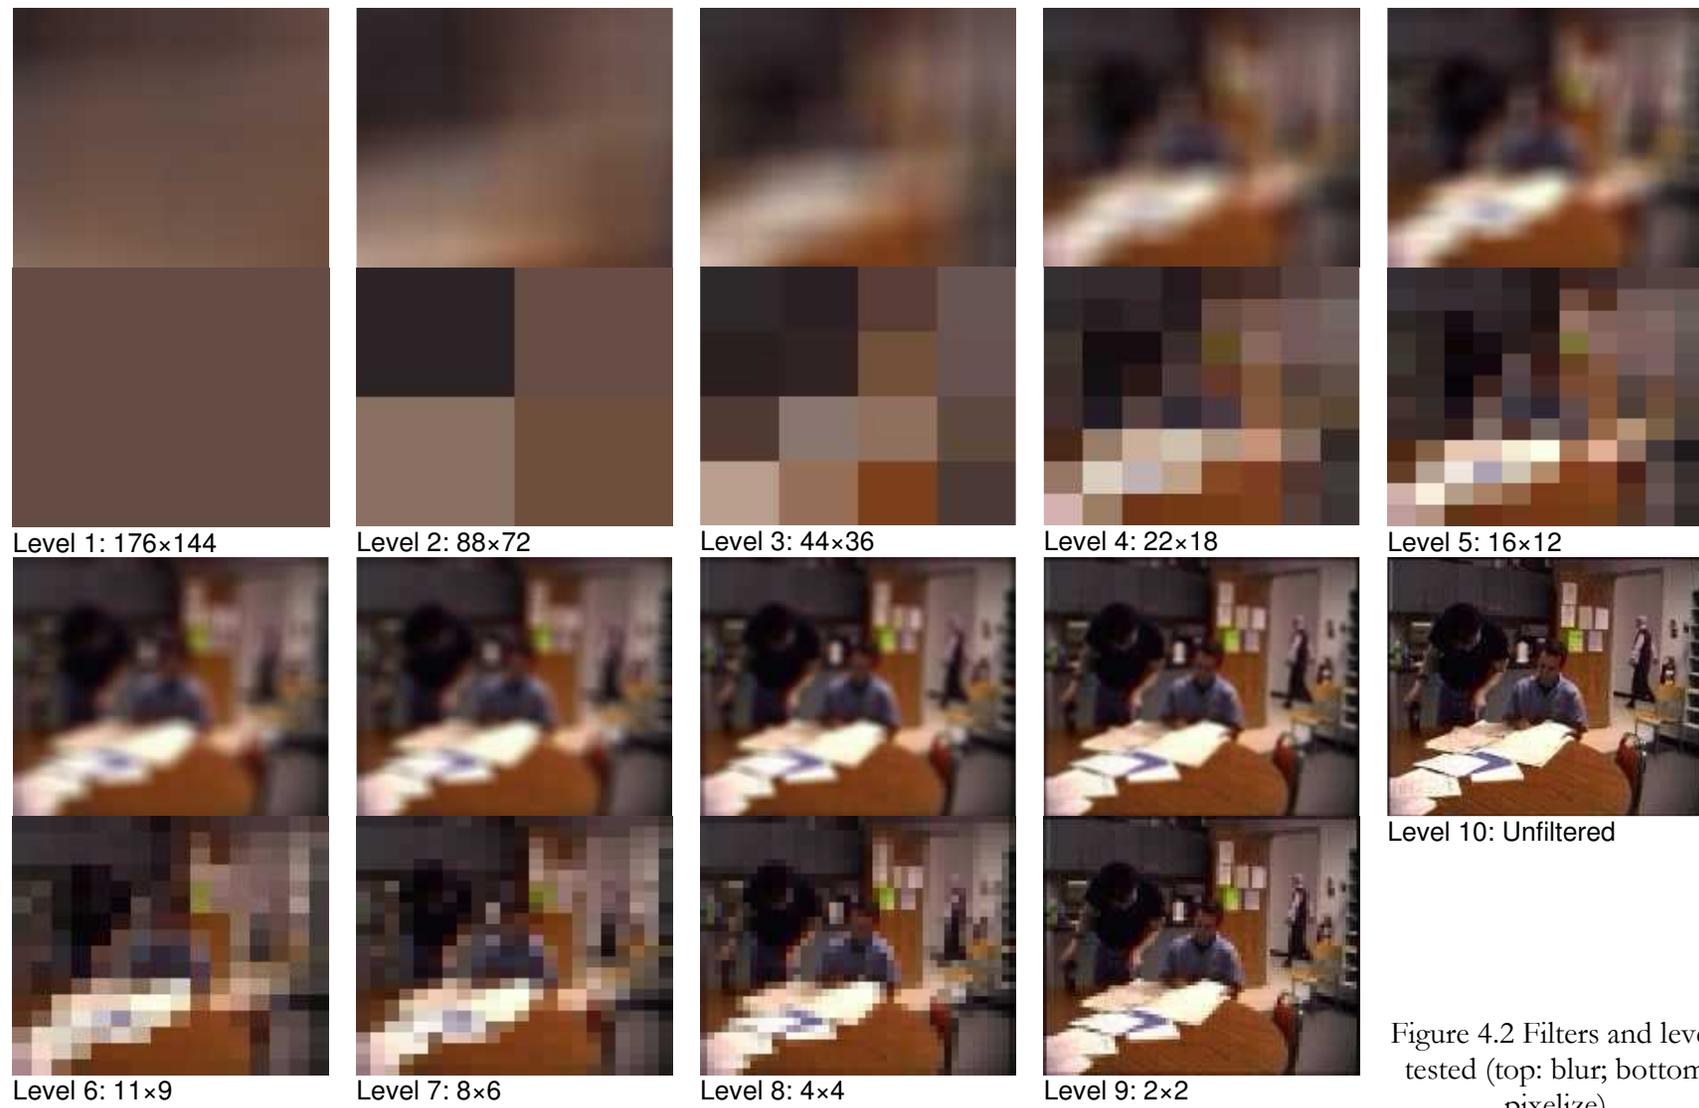
**Scene 4 (86s):** large public office plus counter. A male actor arrives at the counter, and the female actor gets up from her desk to serve him. She retrieves something off screen, and returns. He leaves just as a third male comes to the counter. A fourth male walks by in the background.



**Scene 5 (27s):** over the shoulder, personal workstation. Female actor looks at workstation while typing and moving the mouse. On-screen activities, such as window resizing, are visible.

Figure 4.1—Video scenes used.

The **blur filter** produces a smooth blurring effect across an image (Figure 4.2, upper images in each row). It is a box filter, meaning that a pixel's filtered colour is the mean of the neighbourhood of pixels surrounding it. Unlike the discrete regions seen with the pixelize filter, the image smoothly changes from one region to the next. Our implementation of the blur filter was compute-intensive; similar but computationally inexpensive filters (e.g., Deriche, 1997) could be used instead for real time video manipulation.



These effects can be applied at different *filter levels*. Each level is associated with a “box size” chosen so that they nicely divide into the 176×144 pixel dimensions of the QCIF image. For example, the level 2 box size of 88×72 (see Figure 4.2) divides the 176×144 pixel image into four boxes. For the pixelize filter, this box size corresponds to the dimensions of each “macro-pixel” in the filtered image. For the blur filter, the box size corresponds to the neighbourhood of pixels used to determine a pixel’s value.

We applied the two filters at nine different filter levels to each of the five scenes (Figure 4.2 illustrates one frame from one scene using both filters at all levels). Levels 1-9 range from heavily to lightly filtered. Level 10 is the unfiltered sequence. This gave 95 sequences: 5 (video scenes) × 2 (filter types) × 9 (filter levels) + 5 (unfiltered video scenes).

### 4.2.3 Materials: Questionnaires

We designed three questionnaires to gather experimental data.

*Pre-test questionnaires* captured standard data about each volunteer, including age, gender, and computer and groupware experience. The questionnaire also asked the participants about:

- their willingness to give out personal information over the Internet (e.g., date of birth, credit card number);
- their perception on whether their privacy is protected when using a computer;
- their willingness to let personal acquaintances view their video image on another computer;
- their willingness to let strangers view their video image on another computer, given that reciprocity would be enforced.

*During-test questionnaires* captured how a person perceived a single video sequence using a particular filter at various filter levels. We developed a special web-based system that automatically displayed the appropriate video clip, where participants could fill in the questionnaire by selecting a combination of radio buttons, sliders, and by typing into text fields. Figure 4.3 illustrates a screen snapshot. The questionnaire asked people to identify (when feasible) scene features listed in Table 4.1, for each of these items, the questionnaire also asks people to rate their confidence of their assessment.

- Number of actors present
- Whether each actor is seated or standing
- Whether each actor is moving or still
- Gender of each actor
- Objects in the scene and their location within it
- General activity of each actor
- How busy or idle each actor appears
- 'Tone' of the activity i.e., from casual to serious
- How approachable or withdrawn each actor appears

Table 4.1 Topic of questions asked in the during-test questionnaire.

The software also included a special 'final' question that asked them to rate the privacy-preserving potential of a still snapshot taken from the middle of a video sequence and filtered at different levels i.e., from unprotected to protected (Figure 4.4). Participants could replay the filtered sequence by clicking a given snapshot.

*Post-test questionnaires* captured the volunteer's perceptions of the filter they tested.

— likes and dislikes;

— the situations where they would enable the distortion and what level(s) they would set it to; and

— if they would use an open video link if it was filtered.

#### 4.2.4 Experimental design

**Variables.** The independent variables are scene type (5)  $\times$  filter type (2)  $\times$  filter levels (10). The dependent variables, captured via the questionnaires, were people's ability to identify correctly particular awareness cues, their confidence in their identification, and their perceptions of how privacy is maintained. Particulars were listed in the questionnaires section.

**Participants** were recruited from junior-level university courses, mostly from computer science with some from other disciplines. They comprised ten male and ten female volunteers between the ages of 17 to 32 years. All participants had good or corrected eyesight

**Subject assignment.** We used a mixed design. Scenes and filter type were between subject: each volunteer was randomly assigned to one of the five scenes, and to either the pixelize or blur filter. Filter levels were within subject: each volunteer saw that scene ten different ways: the nine filtered levels and the tenth unfiltered view.

The screenshot shows a web browser window with the address `http://www.sern.ualgary.ca/grouplab/video2000`. The main content area is a questionnaire interface for video analysis. It features a video player on the left showing a scene, and a table of objects and locations on the right. Each object and location has a corresponding confidence slider ranging from 'Unsure' to 'Confident'. Below the table, there are several sections for identifying people in the scene, including radio buttons for 'How many people are present?' and 'Pick a person. Call this person #1. Is this person...', and sliders for attributes like 'Idle vs. Busy', 'Casual vs. Serious', and 'Withdrawn vs. Approachable'. At the bottom, there are dropdown menus for 'Activity' and 'Location'.

Object	Location	Confidence
Table	lower left	Unsure ————— Confident
Counter	along rear of room	Unsure ————— Confident
Door	upper right	Unsure ————— Confident
Papers	on table	Unsure ————— Confident
		Unsure ————— Confident

Level: 4/10

How many people are present?  
 0  1  2  3  4    Unsure ————— Confident

**Pick a person. Call this person #1. Is this person...**

Don't know  Moving  Still    Unsure ————— Confident

Don't know  Standing  Seated    Unsure ————— Confident

---

Don't know    Idle ————— Busy    Unsure ————— Confident

Don't know    Casual ————— Serious    Unsure ————— Confident

Don't know    Withdrawn ————— Approachable

---

Activity:     Unsure ————— Confident

Location:     Unsure ————— Confident

Figure 4.3 Screenshot of in-test questionnaire software.

Filter order was fixed, always beginning with a heavily filtered scene and ending with an unfiltered scene. In effect, this within-subject assignment created a ‘threshold detection’ study, where people could report when they first saw certain types of information in the scene as it became clearer across each filter step.

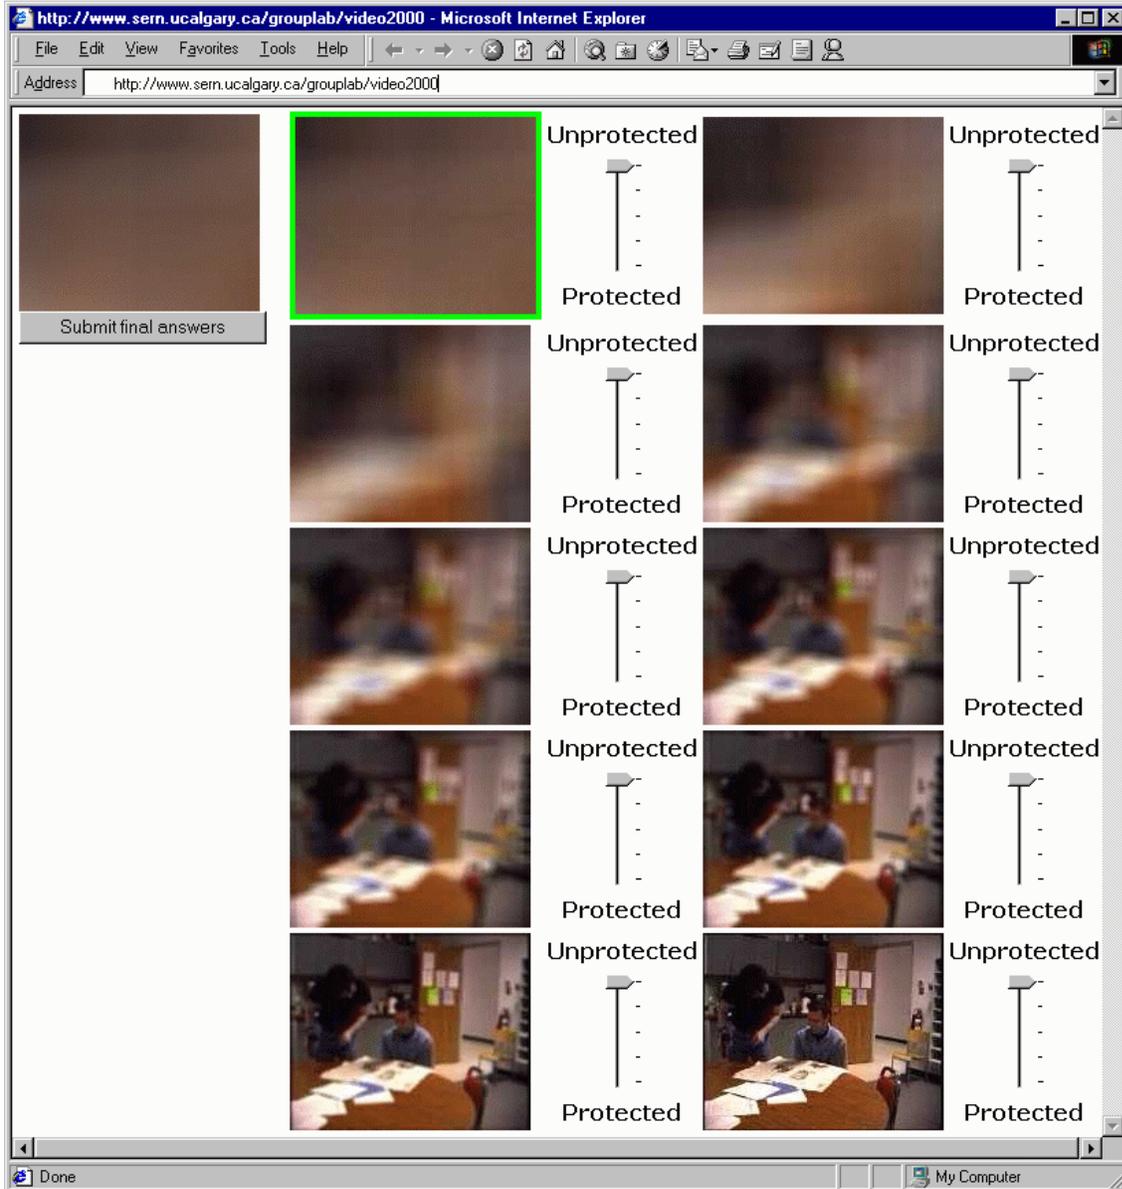


Figure 4.4 Screenshot of in-test software showing last question.

#### 4.2.5 Procedure

Testing was largely self-directed and self-paced. Each participant was first assigned to a particular scene and filter. After completing a consent form and pre-test questionnaire, the participant started our software on a 19" 1280×1024 display. The software loaded the video scene filtered at level 1 (heavily filtered) and the 'during-test' questionnaire form for that scene (similar to Figure 4.3). The participant then viewed this scene as many times as desired: the

video clip operated on a loop, and the scene faded to black between loops. The participant then tried to identify features in the scene and their confidence in their reply by filling in the during-test questionnaire. When satisfied with their answers, the participant proceeded to the next filter level and questionnaire form. Answers provided at the current level were stored for data analysis and automatically copied to the next level's response form. Thus, participants needed only to modify their answers as they identified further features at each new level.

Once the participant completed all ten levels, the software asked the subject to rate the privacy protecting potential of each filter level. Similar to what is seen in Figure 4.4, this form displayed a single frame extracted from the middle of each scene, but shown at the ten filter levels. Participants were asked to imagine themselves as the main actor in the scene, and then to rate how they felt their privacy would have been protected had someone else been observing them. While we deliberately left the identity of the observer somewhat vague, only a few participants requested clarification: these were told that the observer was a peer e.g., a co-worker.

## 4.3 Results

In this section, we describe our observations from the study.

### 4.3.1 Pre-test questionnaire

In the pre-test questionnaire, 12 of the 20 participants answered 'yes' when asked if their privacy was protected when using a computer, and if they would be willing to provide personal information over the Internet. When asked if they would be comfortable if someone they knew could view them through an always-on computer-based video connection, 13 out of 20 responded 'yes.' Similarly, when participants were asked if they would be comfortable if a stranger could view them over a reciprocal video channel 12 responded 'yes' to the question. The participants who responded negatively to these questions raised several concerns. Some said their answer would depend on the situational context i.e., whether they were in a business, academic or home setting. For video links, they were concerned about how well they knew the individual on the other side. Some displayed a general unwillingness to let others know what they were doing. Others stated a sense of discomfort with the idea of always-on video. What

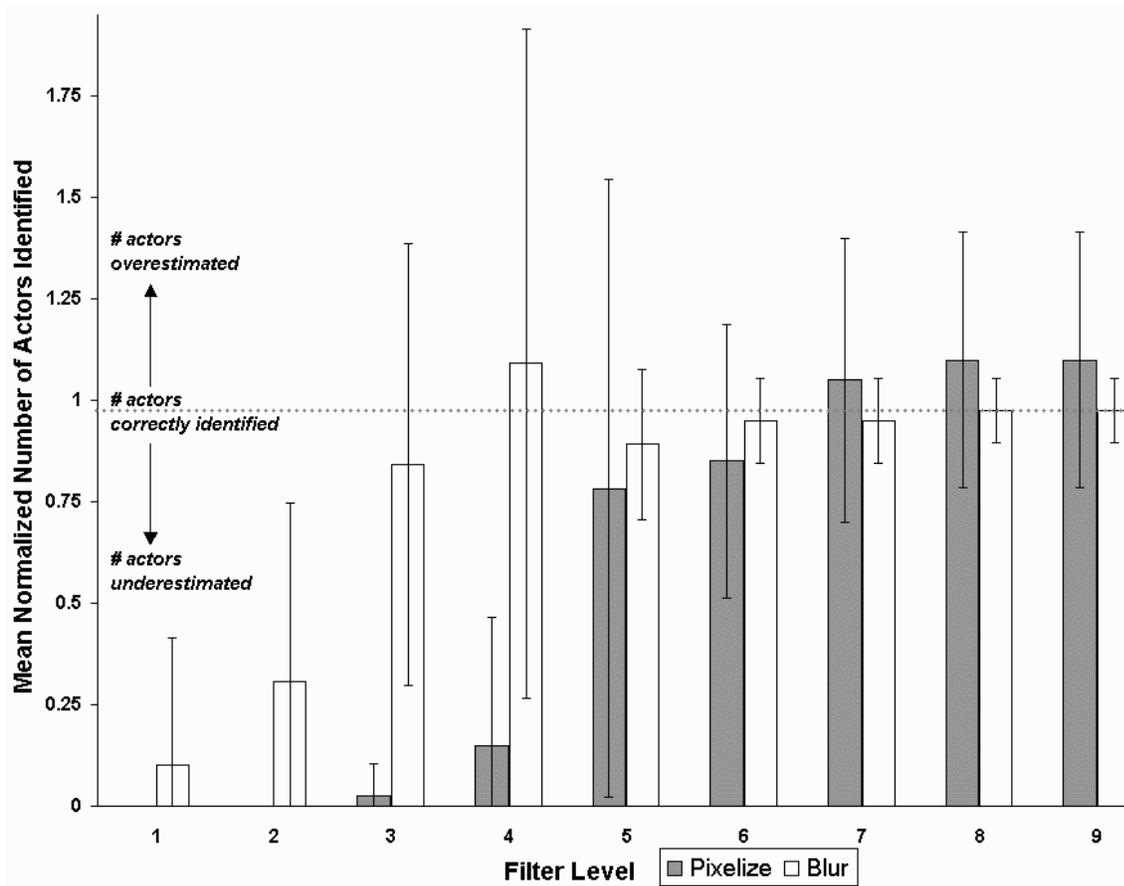


Figure 4.5 Normalised number of actors identified *vs.* filter level.

was somewhat surprising is that these answers were spread across participants i.e., while it seems that about 12 participants always answered positively, it was not the same 12 for each question.

### 4.3.2 Identifying the number of actors in a scene

We wanted to know how well participants could correctly identify the number of actors in a particular video scene under different filters and levels. Since the number of actors featured varied with the scenes, we transformed our result into a normalised metric that would facilitate comparison across scenes. Using this metric, values equal to 1.0 mean that the subject correctly identified the number of actors present. Values less than 1.0 meant that they identified fewer actors than actually present, while values greater than 1.0 mean they overestimated the actors.

Figure 4.5 plots the mean normalised number of actors observed at each level for the two filters tested across all scenes; the error bars show the standard deviation. Although there was much variation, we see that participants using the blur filter were more willing to take a guess at the number of actors early on. Participants could also assess the number of actors in the scene with accuracy at 0.8 or above as early as level 3 with the blur filter, but not until level 5 with the pixelize filter. The amount of variation drops considerably with the blur filter by level 5.

We then analysed the participants' confidence of their guesses at each level of a filter. Not surprisingly, a person's confidence increases along with the filter level i.e., as the clarity of the image improves. There was little difference between the reported confidence and filter type; people seemed moderately confident in their answers by level 5.

We then looked at these results on a per-scene basis. In particular, scenes 1 and 4 contained background actors that were in view for only a brief period, and we wanted to know if participants' guesses differed in these scenes. These scenes proved to be the ones that contributed most to under-estimation errors. However, most participants correctly spotted the background actor by level 6 with both filters.

### 4.3.3 Identifying posture

We wanted to know if participants could identify the posture of the principle actor in a scene i.e., whether they were seated or standing, and whether they were moving around or stayed fairly still. We analysed the data and determined, for each subject, the level (threshold) in a given scene where they correctly identified these posture attributes. Results are plotted in Figure 4.6 for each scene.

As seen in this figure, participants correctly determine posture early on with the blur filter: levels 2-3 for movement, and levels 3-4 for their seated/standing position. With the pixelize filter, participants do not do this until later on: about level 5 for both movement and seated/standing position. In both cases, most everyone correctly states the posture of the main actor featured by level 6.

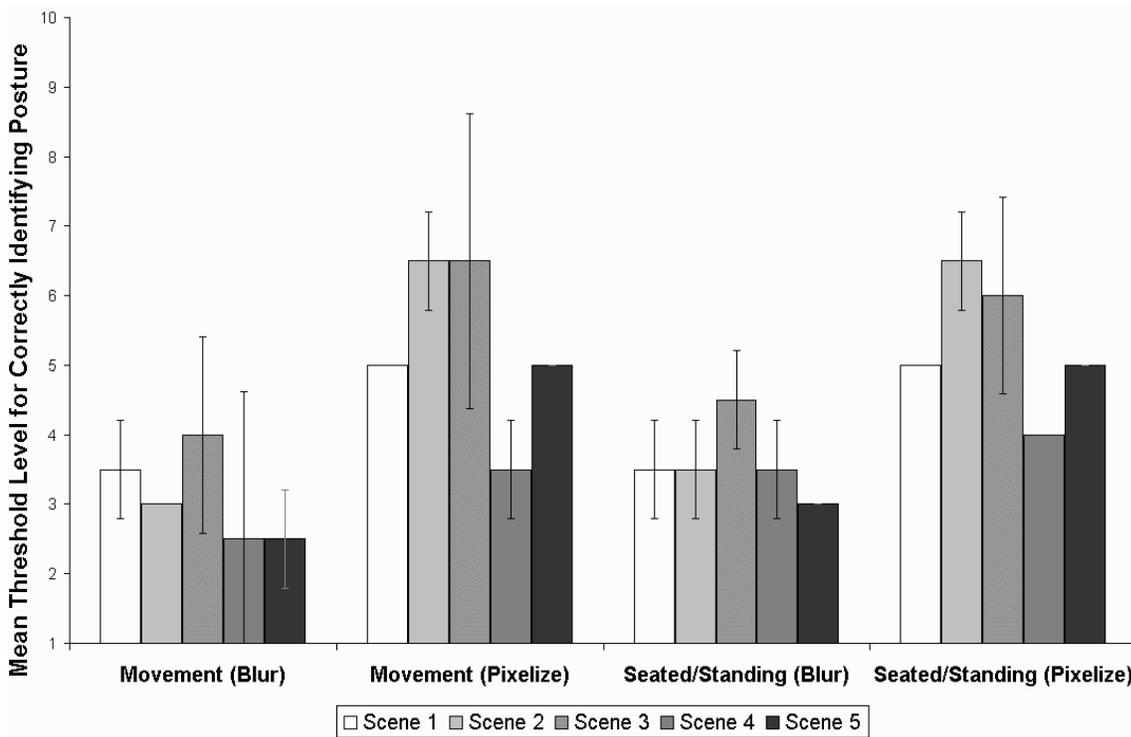


Figure 4.6 Mean threshold level where participants correctly identified actor posture.

With both filters, people seemed to be able to identify posture at the same time they correctly identified the number of people in a scene. There is a strong correlation (0.97) between determining someone's seated or standing position with their ability to identify the number of actors in the scene.

#### 4.3.4 Identifying gender

We wanted to know how well participants could assess the gender of the principle actor in each scene. This proved perhaps the most problematic category, in part because this is difficult to determine even at full fidelity because of the small size of actors portrayed and the poor lighting in some scenes. While we do not show our analysis graphs here, we found that participants are unable (or unwilling) to assess gender for the pixelize filter until levels 6 and 7, and even then only 30-40% did so correctly.

By level 8, however, all participants are able to correctly assess gender. By contrast, participants appear more willing to assess gender earlier when viewing the scene under the blur

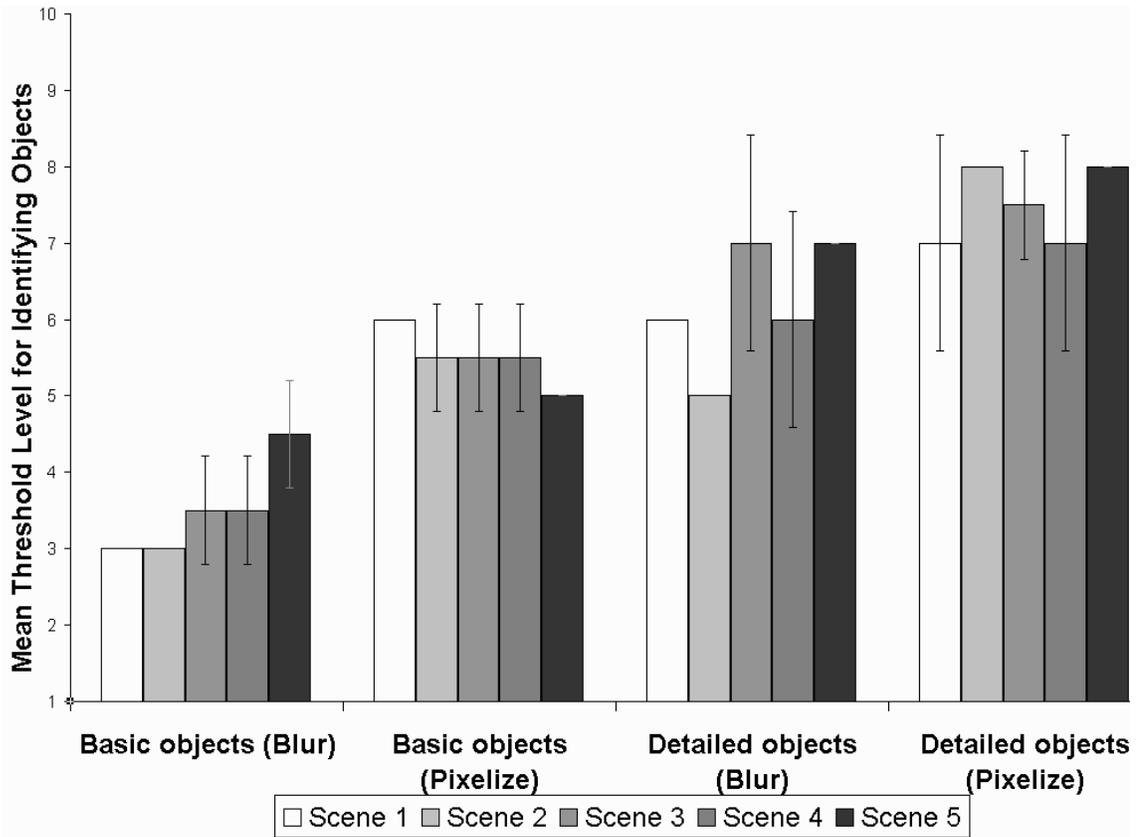


Figure 4.7 Mean threshold level where participants could identify objects.

filter, although they are often wrong. At levels 6-7, 60-65% of the participants correctly assess gender, and most everyone gets it by level 8.

#### 4.3.5 Identifying objects in a scene

As participants viewed each scene across the different filter levels, the images would progressively reveal more information. We asked participants to try to identify any objects visible within the scene as soon as they could. We then analysed the raw data, where we looked at each participant's answer and decided at what level (i.e., the threshold) the participant had roughly and correctly identified some of the basic objects in the scene, and at what level he or she had correctly identified some of the more detailed objects.

The results are plotted per scene in Figure 4.7. While there are differences between scenes, most participants could roughly identify several basic scene objects by level 3 using the blur filter, but not until about level 6 with the pixelize filter. Similarly, participants began

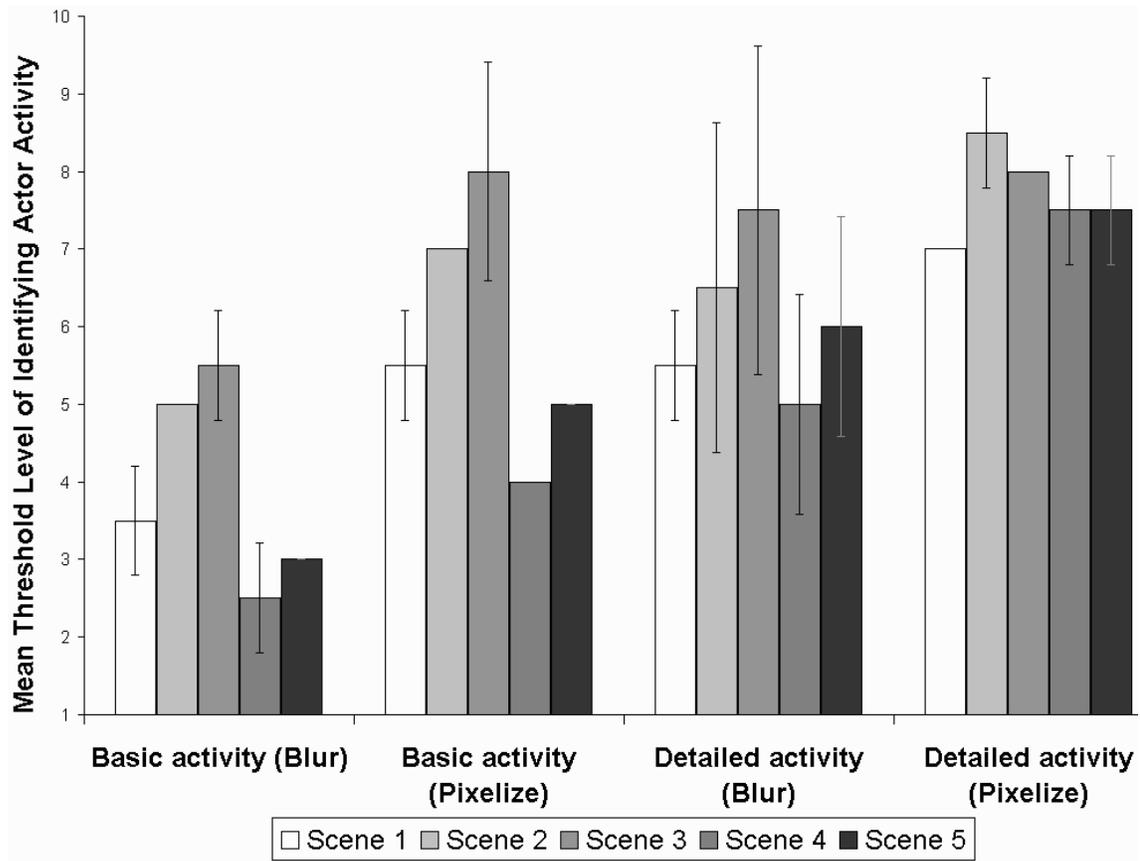


Figure 4.8 Mean threshold level where participants could identify activities.

identifying detailed objects by level 6 with the blur filter, but not until levels 7-8 with the pixelize filter. From our own personal observations and further analysis of participants using the pixelize filter, we noticed there was a very small gap between when people first identified the basic objects in a scene and when they were able to identify the objects in detail; that is, there was only about a one level difference.

#### 4.3.6 Identifying actor activity in a scene.

Similar to how we analysed object identification, we analysed activity identification by deciding at what level participants correctly identified the basic and then the detailed activities of actors. Results are plotted in Figure 4.8. Again we see that people identify basic activities using the blur filter earlier (around levels 3-4) than when they use the pixelize filter (levels 5-7). Similarly, they see detail by around levels 5-6 with the blur and levels 7-8 with the pixelize filter.

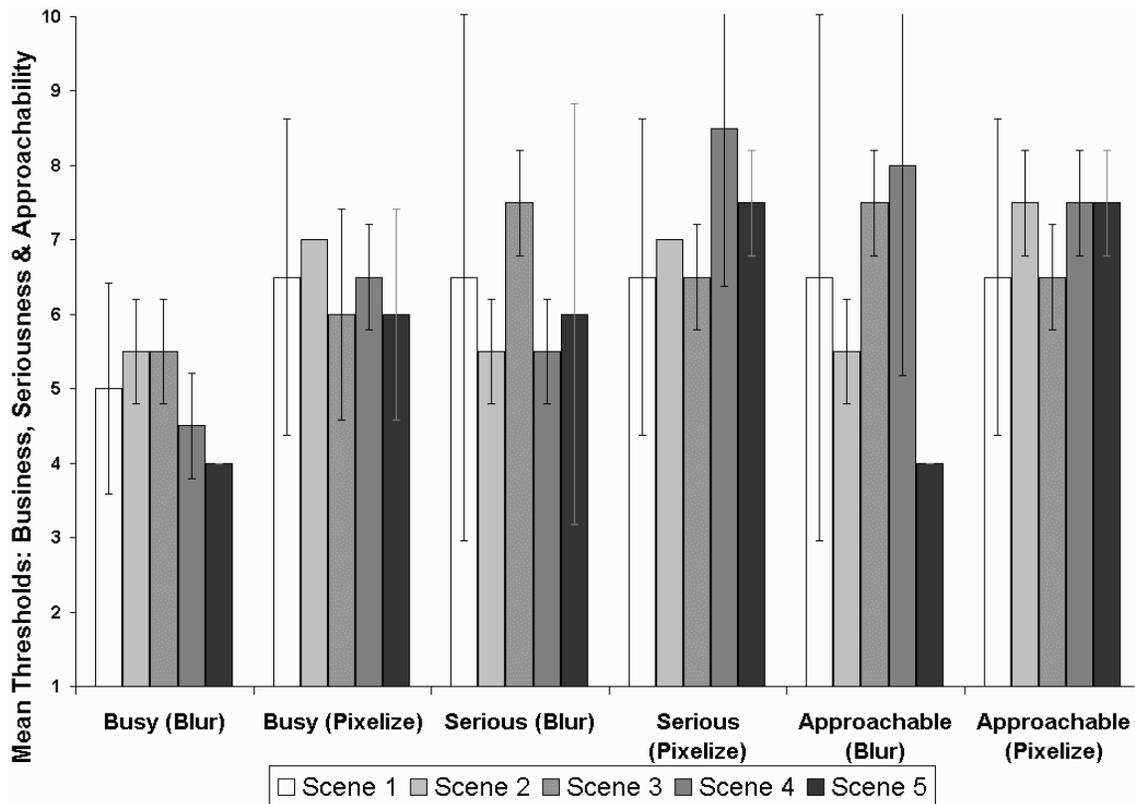


Figure 4.9 Mean threshold at which participants could confidently indicate busyness, seriousness, and approachability.

#### 4.3.7 Identifying busyness, seriousness and approachability in a scene

While there are many ways for people to determine availability, we feel that a person's estimates of an actor's busyness, seriousness, and approachability are reasonable indicators. Yet a person's determination of these availability metrics is highly subjective: even given perfect video fidelity, people may make quite differing judgments. Consequently, we analysed when participants were willing to make a judgment of availability without considering whether this judgment was correct.

Figure 4.9 plots the thresholds that participants appeared willing to commit themselves to a decision of busyness, seriousness, and approachability. Results are highly variable both between participants and across scenes. Still, people seem to declare busyness at around level 5 (blur filter), and level 6 (pixelize filter). Seriousness and approachability seems to

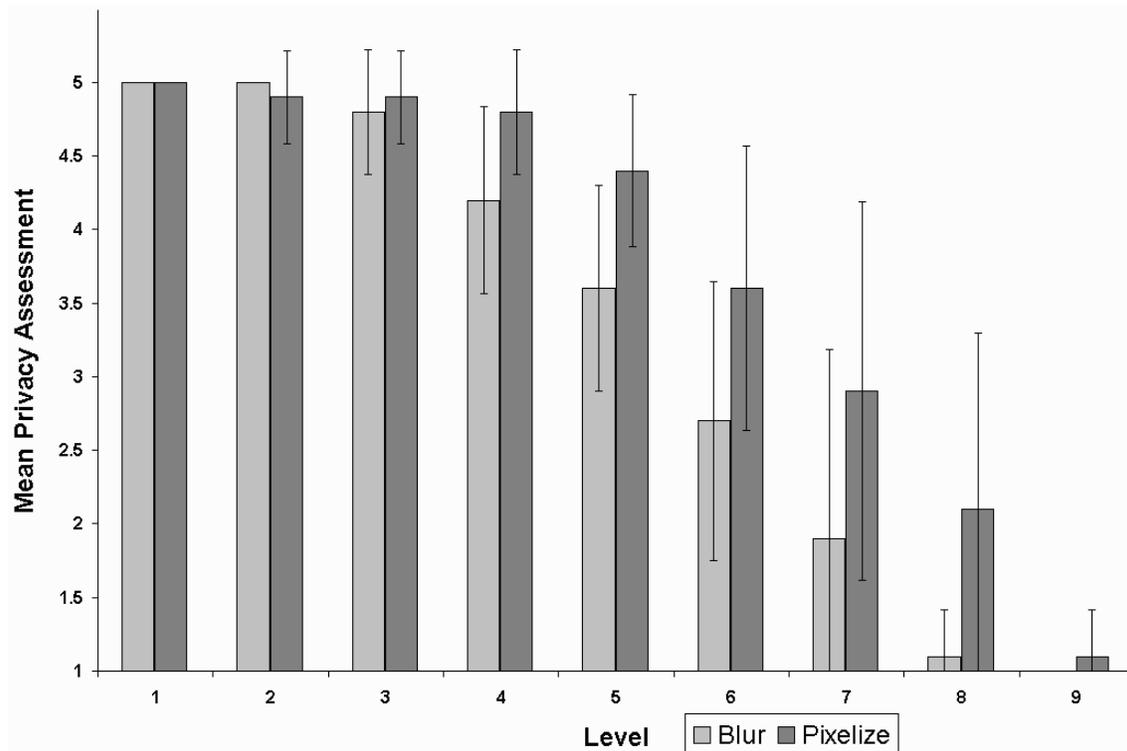


Figure 4.10 Mean privacy rating by level: 1 is unprotected, 5 is highly protected.

demand more fidelity, perhaps to make out actors faces and expressions. People offer judgments by about levels 5-6 (blur filter) and levels 6-7 (pixelize filter).

These judgments are stable: once made, participants rarely changed them even as fidelity increased. As would be expected, confidence in judgments increased with fidelity. For example, while people were only moderately confident of their first attempt to assess busyness, this quickly increased. In particular, more than half the observers were very confident about the accuracy of their assessments by level 5 with the blur filter, but only somewhat confident by level 7 with the pixelize filter.

#### 4.3.8 Rating privacy

We had asked people to imagine themselves to be the main actor in the scene, and then to rate how well a filter at each filter level would protect their privacy. Ratings went from 1 (unprotected) to 5 (protected). In Figure 4.10, we plot these privacy assessment by levels, averaged across all scene types. We see that both filtration techniques do conserve the privacy

factor at more opaque filter levels. Including the standard deviation, participants give ratings of  $\geq 3$  (which means privacy is moderately to fully protected) between levels 1-5 with the blur filter, and between levels 1-6 with the pixelize filter.

We also noticed that the scene type somewhat affected participants' privacy assessments. In general, people were more relaxed about privacy in Scene 1 (the public coffee room) but more concerned about Scene 2 (someone eating a snack) and Scene 4 (viewing a public serving area from afar, which usually implies that some people will not be aware of the camera).

#### 4.3.9 Post-Test questionnaire

Post-test questionnaire items asked participants various questions regarding the filtering technique they tested. When asked what they liked and disliked about the filtration method they used, participants generally gave more 'likes' comments to the blur filter as compared to the pixelize filter. They liked the way the blur filter concealed identity, they liked how they could determine movement while still masking details, and they felt it was visually 'smooth.' They also felt the blur filter had potential to regulate privacy. While there were similar positive comments about the pixelize filter, they disliked that it was hard to see who was there and that it was often difficult to tell what kind of scene/environment was captured by the camera.

Another question asked whether people would leave a video link on themselves if it were filtered. Responses differed depending upon which filter the person had used. Six out of 10 replied that they thought a pixelize filter would be sufficient, while nine out of 10 thought that a blur filter would be effective enough to leave the video always on.

### 4.4 Discussion

Given these results, we can reject the two null hypotheses. Of course, this comes as no surprise. We expect that people can clearly identify more awareness cues as fidelity increases across filter levels, just as we would expect protection of privacy to decrease. A more interesting question is to consider if there is a filter type/level combination that provides basic awareness and a reasonable safeguard to privacy. We answer this by considering all our results together, as summarized in Table 4.2.

	<i>Blur Pixel</i>	
<b>Number of people</b>	3	5
<b>Posture:</b>		
movement	2-3	4-5
seated/standing	3-4	5
<b>Gender</b>	6-7	~7
<b>Objects:</b>		
basic	3	6
detailed	6	7-8
<b>Actor activity:</b>		
basic	3-4	5-6
detailed	5-7	7-8
<b>Availability:</b>		
busyness	5	6
seriousness	5-6	6-7
approachability	5-6	6-7
<b>Privacy protected to</b>	5	6

Table 4.2 Thresholds for identifying awareness cues and for preserving privacy.

Privacy is a key issue when judging the filtering techniques used in the present study. While privacy is best protected when filtration effects are strong, this comes at the cost of a person's ability to identify information that could be crucial for accurately determining availability. With both filters, however, there appears to be a filtration level that provides safeguards to privacy while still providing basic awareness information.

As indicated in Table 4.2, this balance of awareness and privacy occurs at level 5 with the blur filter. Participants chose this level as the highest level that still provides some safeguard to privacy. Yet, we saw that participants could assess basic awareness information quite early on (number of people, posture, basic scene objects, basic actor activity) and finer attributes at the threshold (the basic availability parameters). We also noticed that participants were generally willing to speculate on key awareness cues as early as level 3, even when this came at the expense of accuracy. At level 5, participants had more difficulty assessing detailed information that could contribute to privacy violations i.e., they could not describe scene objects in detail, nor could they assess gender, and they were just on the edge of identifying actor activity details.

The pixelize filter provides a more precarious balance between awareness and privacy, this time at level 6. What differs is that the level at which privacy safeguards are acceptable is very close to the level at which people are just identifying some basic awareness cues. In particular, the threshold for identifying most awareness information is at levels 5-6, and

participants appeared more reluctant to hazard guesses earlier on. In practice, this probably means that users of the pixelize filter will sometimes find it difficult to accurately identify certain cues i.e., when events in the scene happen further away from the camera; when lighting is poorer, and so on. That is, there is more danger that the pixelize filter at level 6 will operate poorly in some settings because awareness cues are difficult to extract from the scene. On the flip side of the coin, another danger is that it is very close to level 7, which participants said (in the post-test questionnaire) pretty well displayed almost all there was to reveal in the scene. That is, when it is possible to tell anything from a pixelize scene, it soon becomes possible to tell most everything. We checked this by examining our data closely: indeed, the pixelize filter tended to be characterised by large jumps across these threshold levels. When taken together, all these points suggest that, although the pixelize filter is widely used in existing prototypical privacy-preserving video media space applications, it may in fact be a poor choice of filter.

The better ratings of the blur filter over the pixelize filter are echoed in the post-test questionnaire. Perhaps the most telling response was that almost all of the blur filter participants felt that they would use an always-on video link if the blur filter were part of it. While pixelize filter users were also positive, their response was not as overwhelming.

All these results come with large caveats. First, the study is small: there are only 20 demographically similar participants. Second, there is high variability in participants' responses across scenes, across filters, and particularly across the levels tested. We believe that no single level of filtration can guarantee privacy safeguards in all cases for all people. Third, the size of objects within a scene can affect how people view it. Our scenes were mostly medium-range to far shots. If people position their cameras so that certain objects appear large, they will likely be more identifiable. Fourth, people's willingness to expose themselves to others will depend greatly on where their camera is situated. For example, people will be more sensitive to privacy concerns if they are being viewed in their casual home office (where other members of the family may enter in various states of undress) than in a public work office. Fifth, we recognize that our measures of people and object identification for awareness and subjective ratings for privacy violations are crude ones. People's perception of where projecting awareness begins to violate privacy will depend on many other things, such as their relationship with the person on the other side of the video link. While we believe that many

people will be receptive to video links with intimate collaborators—close colleagues, good friends, etc.—only a minority of people would be willing to expose themselves to more distant people (e.g., supervisors) and to the world at large (e.g., where the video can be viewed by anyone on the world wide web).

All this implies that, while a balance between privacy and awareness may be possible, it will be a precarious balance at best. Filtering alone may be adequate for some situations, but certainly not all. Our participants were well aware of this fact, and many underscored the need for ‘blocked’ modes, where a person could completely cut off the video.

## 4.5 Neustaedter’s study

The principle weakness in the evaluation described above is that the scenes used were mundane. They did not depict risky scenarios in which a person would very much want to have privacy preserved. Although they captured naturalistic office scenes and activities, they did not include, for example, home telecommuting scenarios. People appear in the home in various states of undress. Common activities in the home, e.g., a passionate kiss between spouses, might be uncommon or even inappropriate in an office.

Carman Neustaedter, Saul Greenberg, and I conducted a second study (Neustaedter, Greenberg & Boyle, 2005) to address this caveat to the first study results. This study examined the blur filter only, and used a modified experimental method that featured high risk video scenes captured in a home setting (Figure 4.11). Neustaedter was the lead researcher and reported study details in his thesis as well as the joint publication. Consequently, this section were only summarise the study goals, methods, and results.

### 4.5.1 Materials

Figure 4.11 shows representative still images taken from the five video scenes used in the second study, with a brief description of each and an initial projection of the risk that study participants will associate with the scene.

**Working at a computer:** The telecommuter is working at a computer while wearing clothes appropriate for both home and the office (Low risk).

**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 (Moderate risk).

**Working with no shirt on:** The telecommuter is working at a computer with no shirt on (Moderate risk).

**Kissing a partner:** The telecommuter is working at a computer when his/her partner enters the room, kisses the telecommuter intimately, and leads him/her out of the room (Moderate risk).

**Changing clothes:** The telecommuter enters the room in a robe, is shown completely naked, and then puts on underwear (High risk).



Figure 4.11—Scenes used in Neustaedter's second study were of much higher risk than those used in the first study.

Each scene was shot twice, once with a paid female actor (nicknamed Linda) and once with a paid male actor (nicknamed Larry). The actors were deliberately chosen to appear as middle-aged working professionals doing otherwise “normal” home activities. In the Change Clothes scene, the actors were shown completely nude. As in the first study, participants were asked to imagine that the actor depicted was a colleague who would not want to be seen in a compromising situation. NTSC DV-quality video (720×480 pixels, 30 fps) was shown to participants. This is considerably better quality than was shown to participants in the first study. The ten blur levels used roughly correspond to those used in the first study.

### 4.5.2 Method

As in the first study, pre-test, during-test, and post-test questionnaires were administered to participants in the second study.

The during-test protocol was essentially the same as the first study: each video scene was shown, beginning at the lowest level of fidelity. At each fidelity level, the participants were asked to describe what was visible, deduce important informal awareness cues in the scene, and rate how confident they were in their guesses. There were two important differences between the initial and follow-on studies' during-test protocol. First, as each filter level was shown in the second study, participants were asked to rate how threatening it is to the actor and how threatening it is to the actor's family members. Second, when all levels for a scene were shown, participants in the second study were asked to choose a blur level they felt would make the scene appropriate for a colleague of the actor to view.

The post-test protocol differed between the two studies. In the second study, participants were asked to rank the scenes according to risk by positioning representative still images along a "line of privacy risk."

### 4.5.3 Results

The during-test questionnaire responses dealing with the perceptibility of awareness cues were analysed to identify thresholds for understanding. These results showed that awareness cues could be discerned reliably beginning around levels 3~5. There were significant differences in the minimum fidelity level depending on scene and the kind of awareness information looked at. These results are similar to those of the first study.

Most of the important differences between the first study and the second study deal with privacy, rather than awareness. First, participants' assessments of the risk in each scene roughly matched our expectations (stated in Figure 4.11). There were statistically significant results to show that blurring affect the risk assessment for a given scene. Open-ended questions suggest that the actor's appearance or activity were the things in the scene that made it threatening to the actor.

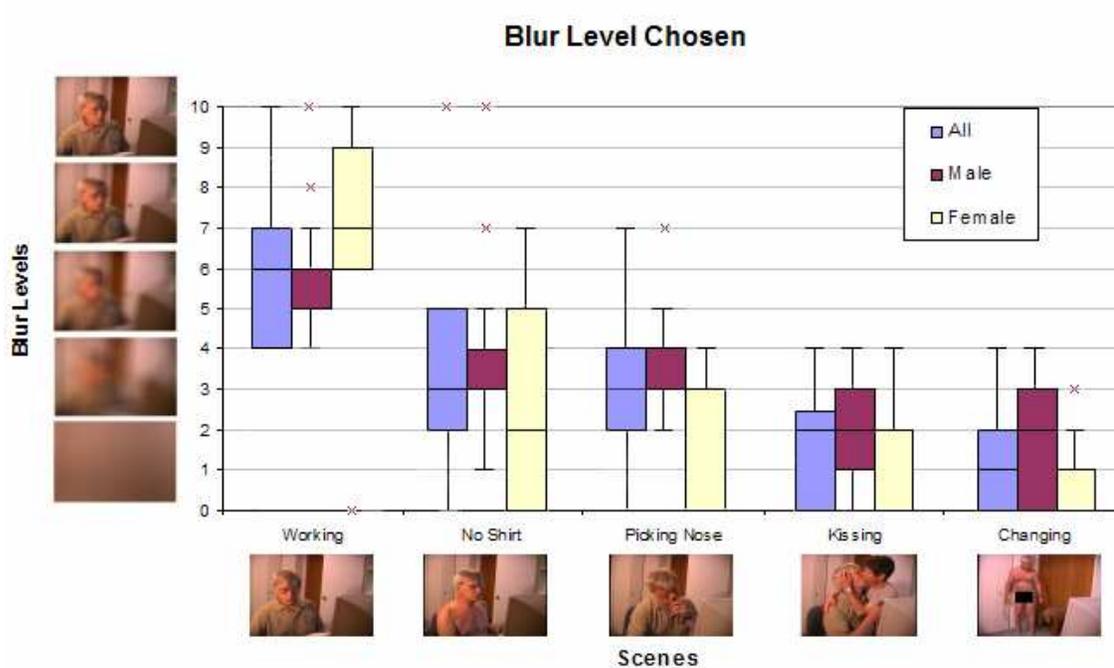


Figure 4.12—The median and range of blur levels chosen by participants for each scene. Blur level 0 represents choosing to turn the camera off. Figure reproduced from Neustaedter, Greenberg & Boyle (2005).

The main improvement in the second study is the use of scenes depicting higher risk scenarios than those of the first study. Consequently, the main difference between the second study and the first study lays in blur levels chosen by participants to make the scenes acceptable for viewing by office colleagues. Figure 4.12 shows a plot of the blur levels chosen by participants in the second study. In particular, the mean blur levels for the highest risk scenes (Kissing and Changing) are very much lower than the minimum needed for awareness and in many cases it was found that users wanted to turn the camera off entirely (blur level 0 in the figure).

## 4.6 Conclusions

Returning to the thesis problems and goals stated in Chapter 1:

**Thesis Problem #2:** It is widely suspected that distortion filtration may be useful for mitigating privacy issues in video media spaces but its usefulness

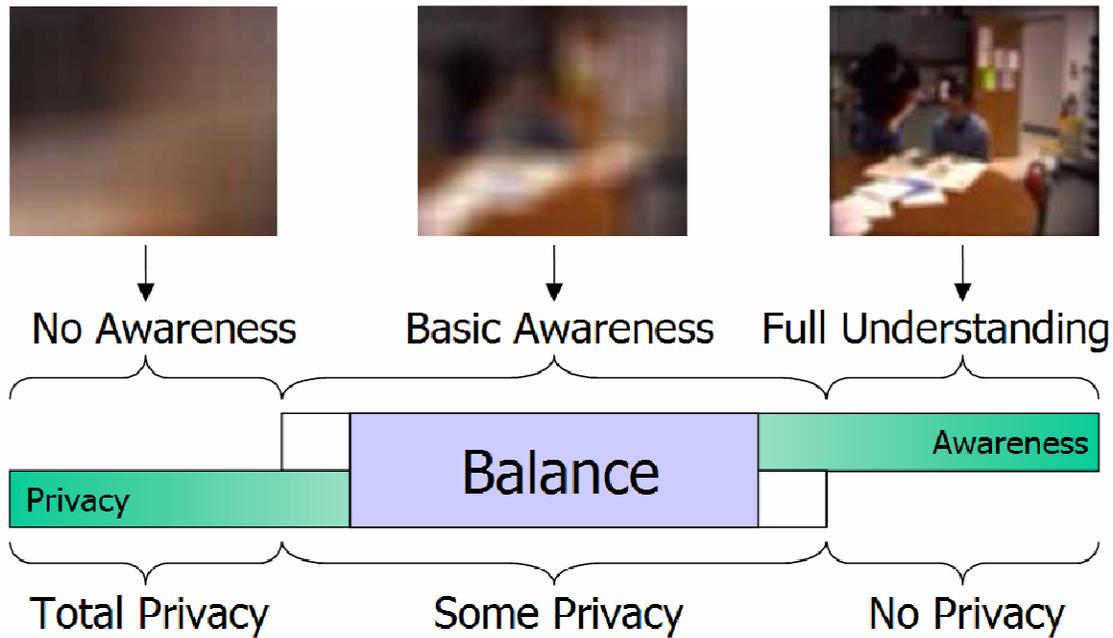


Figure 4.13—The first study found that in mundane office scenes that have low privacy risk, privacy and awareness can be balanced with the blur filter around level 5.

has not been rigorously evaluated and there is no guidance as to how much filtration is ideal.

**Thesis Goal #2:**

Determine if it is possible to use the distortion filtration technique to strike a balance between awareness and privacy in a video media space. If it is possible, determine at which levels a balance can be reached.

**Status of Goal #2:**

Completed. This chapter described the results of the two semi-controlled laboratory user studies that evaluated the blur and pixelize filters for balancing awareness and privacy in a video media space. The results of these studies suggest that while an adequate balance between awareness and privacy may be found for benign situations, this balance will not be found in many risky scenarios using the distortion filtration technique alone.

Our results show that the filter and the level it is operated at do have an impact on privacy and awareness. It is possible to filter a scene so that some aspects are discernable, but

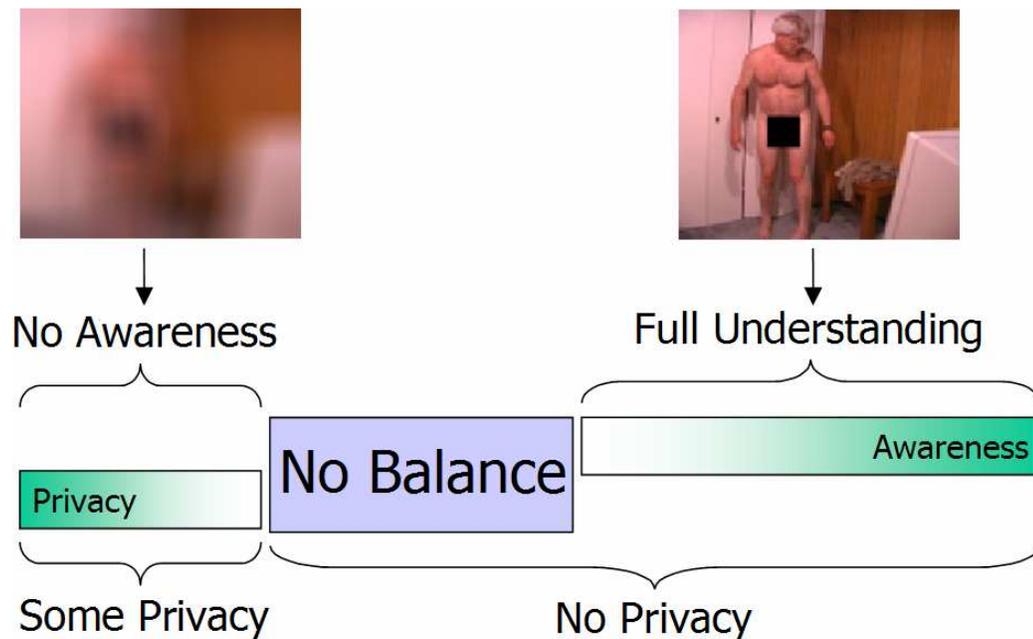


Figure 4.14—The second study found that privacy and awareness cannot be adequately balanced in risky scenarios routinely expected for home telecommuters.

others are not. In both studies, we found that awareness cues could be discerned with the blur filter starting around levels 3~5 (middle image in Figure 4.13) but could not be discerned at levels 1 and 2 (leftmost images in Figure 4.13 and 4.14). We also found that privacy was not preserved at higher fidelity levels (rightmost images in Figure 4.13 and 4.14)..

Where the two studies differ is in the region of overlap found. In the first study, we found that the blur filter adequately preserved privacy in mundane scenes at levels 3~5. In the second study, these levels did not preserve privacy in risky scenes. Furthermore, in the second study we found that filtration down to the lowest levels of fidelity may not even be enough: when the risk is high many people will want the camera to be turned off altogether.

Thus, **there is no general-purpose range of blur filter levels which balance awareness and privacy in high-risk scenarios** i.e., provide awareness cues yet make the scene acceptable for viewing. Interpreting these results, it is possible that distortion filtration may have a place in an approach to supporting privacy that incorporates a broad spectrum of other techniques, but in isolation the technique will be of limited success

This result calls into question the near-universal application of this technique. It is now questionable if using this technique in television news broadcasts suffices to protect the identity of, say, juvenile criminal offenders (for whom there are often legislated protections). Blurring is, of course, just one technique. Perhaps other promising filters may be preferable to blurring e.g., the eigenfilters of Crowley, Coutaz & Bérard (2000). The methodology proposed here might be repeatable and therefore useful for evaluating these other filters. The results of this study do not deny the possibility of technology-based support for privacy preservation in a video media space, but they do indicate that we need to examine more closely the individual and group behaviours that support the broader context of privacy as people experience it. The understanding is necessary so that we may understand how the interface to a technique like distortion filtration can be designed to support privacy as on-going practice.

One approach, explored by Neustaedter in his HOME MEDIA SPACE, is to identify a handful of simplified heuristics for applying the technique in ways that fit the current context. His system uses simple environmental sensors that serve as inputs for automated heuristics that adapt the level of distortion according to the risk faced by media space participants, even turning the camera off and physically rotating the camera to point it at a wall when the risk is high (Neustaedter & Greenberg, 2003). Another approach is to develop user interface components which allow people to dynamically adjust the level of filtration according to their preference in a very lightweight manner. Similar work is being done on this topic outside the video media space domain such as the Place Bar widget for selectively revealing location information to location-enhanced web services at various fidelities (Hong et al, 2003).

This chapter also brings to a close the first act of this thesis, which examines low-level technical factors related to the design of video media spaces. This examination covered video media spaces as a technology to support distributed collaboration by providing informal awareness cues leading into spontaneous casual interactions; the COLLABRARY as a technology to support the rapid development of working video media space systems; and, the distortion filtration technique as a technology to mitigate privacy problems in video media spaces. Although human factors were considered in each chapter, the focus has remained on the important technical factors. The next act takes a major turn by focusing on the important human factors related to privacy first while considering technical factors second.

## PREFACE TO ACT II

### HIGH-LEVEL SOCIAL FACTORS RELATED TO PRIVACY-PRESERVING VIDEO MEDIA SPACE DESIGN.

The three chapters bundled in this second act of my thesis present work I did to understand the complex nature of privacy.

I reached a point in my attempt to design a privacy preserving video media space where it became clear that I did not understand what privacy is within this context. That is, I came to realise that in order to design a privacy preserving video media space I needed a better understanding of the individual and social practices that make up privacy as people experience it on an on-going basis: the broader context into which the privacy preserving techniques of a video media space will fit. Thus, I put aside the bottom-up approach of Act I, and immersed myself in theoretical descriptions of privacy I found in psychology, sociology, anthropology, law, architecture, and computer science. What I found were incredibly rich theories of privacy that, for the most part, had absolutely nothing to do with computers.

I found myself at every turn re-expressing the highly abstract and generalised theories of privacy I encountered and applying them to the concrete problems reported in video media space literature in order to make sense of them. In doing so, I became aware of the important role that language has on shaping our understanding of privacy and I came to unify in my mind the vocabularies of each discipline from which I drew my background. I borrowed selectively—including only what I felt was needed with no commitment to the rest—and freely intermixed my own observations and analysis to resolve incompatibilities and render a new theory of privacy that satisfactorily accounted for what had written about privacy in video media spaces. In my opinion, these are the most important chapters in this thesis.

This act has been published as Boyle & Greenberg (2005). It builds a vocabulary of terms that permit unambiguous and holistic description of privacy in the context of video media space design and use. It lays the conceptual foundation for Act III, which presents a method to use these concepts and vocabulary terms to systematically analyse and describe the

relationships between privacy and design in a video media space, illuminating hidden assumptions or omissions in the notions of privacy embodied in the design.

A fundamental premise in this thesis is that privacy is intuitively understood by everyone, and I often preface presentations of my privacy theory with statements to the effect that “everyone is an expert in privacy.” I follow this up with the claim that even though everyone intuitively understands privacy, it is a difficult thing for people to think analytically and talk clearly about. I have chosen to attack this problem by building a vocabulary for describing privacy in a comprehensive and unambiguous manner. This vocabulary is the main contribution of Act II.

I am hardly the first researcher to talk about privacy. Naturally, I have grounded my theory atop a broad base formed out of others’ theoretical descriptions of privacy, i.e., Altman, Bellotti, Palen & Dourish, and Schwartz. However, I stress that the vocabulary I have produced is much more than a mere literature review. As we will see, the theory presented in the chapters that follow is original and goes beyond mere synthesis of others’ ideas in several important ways.

I present not only the concepts already discussed in privacy in video media spaces, but also concepts from other disciplines that have never been used with video media spaces before: e.g., refuge; territoriality; role conflict.

I give existing concepts new interpretations that are more powerful and more precise than the accepted O.E.D. definitions: e.g., control as a normalised, cooperative phenomenon; violation as impairment of control; risk as the integration of the probability and severity of harms arising from a violation.

I introduce entirely new concepts that unify smaller constituent ideas: e.g., solitude as attention control; confidentiality as fidelity control; autonomy as self-presentation control.

I bring a completely original organisation of the concepts: e.g., privacy is a synergistic union of solitude, confidentiality, and autonomy in a dynamic, normalised, situational dialectic; computer security, territoriality, cryptography, disinformation, distortion filtration, access control, ambiguity, plausible deniability and genres of disclosure all exist in a single unified space of different means people exploit to regulate confidentiality.

I apply these concepts to video media spaces to yield original insights into the design problem: e.g., risk/reward disparity; apprehension as a top-level design problem.

Finally, the work in this act and in Act III does not override or negate the work presented in Act I. In particular, it does not change the results of the studies described in Chapter 4 or even call into doubt those results. Rather, the following chapters:

help us make sense of the results;

point the way towards new hypotheses to test, new methodologies to employ, and new toolkit features to implement;

change the questions we ask about video media spaces and privacy; and,

reveal important assumptions and omissions in the conception of privacy that constrains the scope of the earlier work.

In doing so, these Acts extend and enrich the earlier work, by taking it in a much needed different direction.

## Chapter 5—Language of privacy in CSCW

Privacy is a multifaceted thing, connected with much of daily life. Perhaps because of the many varied aspects of privacy, it is notoriously difficult to discuss. Each word in the vocabulary that researchers use to talk about privacy is as multifaceted as the thing itself. Perhaps because of this, privacy has been given considerable diverse treatment by hundreds of authors in scientific, engineering, and humanities literature (Brierley-Newell, 1995).

Out of this diversity, however, arises confusion. Different authors may use the same word to describe different concepts or phenomena, or the same author may use different words to describe the same concept/phenomenon without relating the words to one another. Interdisciplinary discussion of privacy is made complicated by obvious differences among the stereotypical conceptions of privacy in different domains. Lawyers stereotypically equate privacy with autonomy (being let alone). Psychologists stereotypically equate privacy with solitude (being apart from others). Technologists, economists, architects and others stereotypically equate privacy with confidentiality (keeping secrets).

To deal with the basic problem of deciding which words to use to talk about privacy in the context of video media spaces, I proposed this thesis problem/goal pair in Chapter 1:

**Thesis Problem #3:** There is no comprehensive vocabulary of privacy terms—one that integrates conceptions and theories of privacy from many disciplines—to support unambiguous description of how privacy is affected by video media space design and use.

Thesis Goal #3: Integrate privacy theories and observations from many disciplines of scientific inquiry to produce a vocabulary for describing privacy and a video media space's effect on it in an unambiguous and comprehensive manner, accounting for at least the privacy issues reported in previous literature.

This act consists of three chapters that presents a vocabulary I have assembled to discuss the relationship between privacy and the design of technology. The vocabulary integrates empirical knowledge of privacy uncovered by researchers in design-oriented disciplines like CSCW, law, and architecture with theoretical frameworks for understanding privacy developed in social sciences like environmental psychology, anthropology, and behavioural psychology.

I have split my description of the vocabulary into three chapters to make it more “digestible.” In truth, the discussion of one carries naturally to the discussion of the next. In this first chapter, I situate the problem by reviewing the notions of privacy that are richly embedded in the CSCW field, from which work on video media spaces emerges.

In the text, vocabulary terms are set in bold when first introduced and discussed. We have extracted many of these keywords and organised them into a shallow hierarchies seen in Table 5.1. This list will be repeated again in Chapter 8, and appears in Appendix B with one-line descriptions of each. This list is an abridged summary guide to the vocabulary: some vocabulary terms set in bold in the text do not appear in the list. Please see Chapter 11 for a description of the list generation.

Readers are encouraged to scan the vocabulary list beforehand, so as to get a sense of variety of topics that will be covered and use Appendix B as a tear-out reading aide. As each term is discussed, we also give the number of it in the vocabulary. We relate the terms to observations drawn from video media space design practice and use. While this chapter does not present specific solutions to privacy problems in video media spaces, it does satisfy our goal of creating a vocabulary that will permit CSCW researchers to discuss privacy issues in video media space design in a holistic yet unambiguous way.

Table 5.1—Vocabulary that embodies the descriptive theory of Act II.

**1. SOLITUDE**

- a) Physical Dimensions
  - i) Interpersonal Distance
    - (1) isolation to crowding
  - ii) Attention
    - (1) focus to periphery
- b) Psychological Dimensions
  - i) Interaction to Withdrawal
    - (1) anonymity and reserve to intimacy
  - ii) Escape
    - (1) refuge
    - (2) fantasy
- c) Presentation Dimensions
  - i) High-level Awareness
    - (1) availability
    - (2) accessibility
  - ii) Distraction
    - (1) relevance
    - (2) salience

**2. CONFIDENTIALITY**

- a) Information Channels
  - i) Medium
    - (1) aural
    - (2) visual
    - (3) numeric
    - (4) textual
  - ii) Processing
    - (1) sampling
    - (2) interpolation
    - (3) aggregation
    - (4) inference
  - iii) Topic
    - (1) information about the self
    - (2) personally identifying information
    - (3) activities
    - (4) whereabouts
    - (5) encounters
    - (6) utterances
    - (7) actions
    - (8) relationships
- b) Information Characteristics
  - i) Basic Characteristics
    - (1) sensitivity
    - (2) persistence
    - (3) transitivity
  - ii) Fidelity
    - (1) precision
    - (2) accuracy

- (3) misinformation
- (4) disinformation
- iii) Certainty
  - (1) plausible deniability
  - (2) ambiguity
- c) Information Operations
  - i) Basic Operations
    - (1) capture
    - (2) archival
    - (3) edit
  - ii) Use
    - (1) accountability
    - (2) misappropriation
    - (3) misuse
  - iii) Scrutiny
    - (1) surreptitious surveillance
    - (2) analysis

**3. AUTONOMY**

- a) Social Constructions of the Self
  - i) Front
    - (1) identity
    - (2) digital persona
    - (3) appearance
    - (4) impression
    - (5) personal space
  - ii) Back
    - (1) flaws
    - (2) deviance\*
    - (3) idealisations
  - iii) Signifiers\*
    - (1) territory
    - (2) props
    - (3) costumes
  - iv) Harms
    - (1) aesthetic
    - (2) strategic
- b) Social Environment
  - i) Social relationships
    - (1) roles
    - (2) power
    - (3) obligations
    - (4) status divisions
    - (5) trust
  - ii) Norms
    - (1) expectations
    - (2) preferences
    - (3) social acceptability
    - (4) conformance
    - (5) deviance
    - (6) place

#### 4. MECHANICS OF PRIVACY

- a) Boundaries
  - i) disclosure
  - ii) temporal
  - iii) spatial
  - iv) identity
- b) Process Characteristics
  - i) dialectic
  - ii) dynamic
  - iii) regulation
  - iv) cooperation
- c) Violations
  - i) risk
  - ii) possibility
  - iii) probability
  - iv) severity
  - v) threat
- d) Behavioural and Cognitive Phenomena
  - i) self-appropriation
  - ii) genres of disclosure
  - iii) policing
  - iv) reprimand
  - v) reward
  - vi) risk/reward trade-off
  - vii) disclosure boundary tension
  - viii) disinformation\*
  - ix) reserve\*
  - x) Signifiers\*
    - (1) implicit
    - (2) explicit
- e) Environmental Support
  - i) situated action
  - ii) reflexive interpretability of action
  - iii) constraints
  - iv) transitions
  - v) choice
  - vi) reciprocity
  - vii) liberty
  - viii) refuge\*
  - ix) Embodiments
    - (1) rich to impoverished
  - x) Cues
    - (1) feedback

(2) feed-through

#### 5. COMPUTERS AND PRIVACY

- a) Support Methods
  - i) computer security
  - ii) cryptography
  - iii) pseudonymity
  - iv) Access Control
    - (1) authentication
    - (2) authorisation
  - v) Content Control
    - (1) distortion filtration
    - (2) publication filtration
  - vi) Reliability
    - (1) data integrity
    - (2) process integrity
    - (3) stability
- b) Problems
  - i) inadvertent privacy infractions
  - ii) apprehension
  - iii) resentment
  - iv) the four 'D's': decontextualisation, disembodiment, dissociation, desituated action
  - v) role conflict
  - vi) Deliberate abuse
    - (1) misappropriation
    - (2) misuse
    - (3) identity theft
    - (4) impersonation
- c) User Interface Issues
  - i) degrees of temporal/spatial freedom for information access
  - ii) risk/reward disparity
  - iii) Feedback and Control
    - (1) believability
    - (2) socially natural qualities
    - (3) utility of privacy countermeasures
  - iv) Effort
    - (1) cognitive
    - (2) physical
  - v) Control Granularity
    - (1) fine- to coarse-grained

Table 5.1 Vocabulary that embodies the descriptive theory of Act II.

### 5.1 Video media spaces: A crucible for studying privacy

Undoubtedly, privacy is a concern for technologists. Some of the ways that technology affects privacy are deemed undesirable. Ethical, political and economic forces compel research on

methods for designing, building and deploying systems that benefit individuals and society without eroding privacy.

Specifically, privacy is important to human-computer interaction design. HCI and CSCW researchers have contributed abundant empirical findings relating privacy to technology design. This is especially the case with research regarding video media spaces (Bellotti, 1998). Video media spaces (VMS) connect small groups of distance-separated collaborators with always-on or always-available video channels. Via these video channels, people gain informal awareness of others' presence and their activities. This awareness permits fine-grained coordination of frequent, light-weight casual interactions. As discussed in Chapter 2, a variety of VMS designs have emerged.

- Snapshot-only video portholes that show occasionally-updated small images of what is happening at other sites e.g., Dourish & Bly (1992) and Lee et al (1997).
- Intermittently open links between personal offices, where people can selectively establish brief or long connections into other spaces, and where they can create the equivalent of an open 'videophone' call e.g., Olsen & Bly (1991), Mantei et al (1991), Gaver et al (1992); and Tang et al (1994).
- Persistently open links between common areas (e.g, cafeterias, lounges) where the video feed from an always-on camera is continuously displayed at distant sites e.g., Fish et al (1990) and Jancke et al (2001).
- 'Video-as-data' uses, where video provides access to a shared visual workspace about which local and remote users can micro-coordinate their individual activities and group interactivities e.g., Nardi et al (1997). Unlike the other conditions, 'video-as-data' configurations use video to transmit workspace awareness cues instead of affective conversation and informal awareness cues.

While video media spaces are a promising way to increase group interaction, they are perceived by users and non-users alike to be privacy invasive and privacy insensitive e.g., Gaver et al (1992), Bellotti & Sellen (1993), Lee et al (1997). They permit privacy violations that range from subtle to obvious and from inconsequential to intolerable. Early media spaces users were typically enthusiastic about the technology yet highly aware of its potential for

sociological and psychological impact. This combination of participants and problems makes video media spaces an excellent crucible for examining the privacy-design link. For example, it is the application area in which Bellotti applies her framework for privacy in CSCW and CMC (Bellotti, 1998).

## 5.2 Approaches to privacy research

Researchers in CSCW generally assume that privacy problems caused by technology arise because of the way systems are designed, implemented, and deployed. For example, Grudin suggests that the underlying drive to increase human efficiency through technology—specifically context-aware systems—leads to design decisions that conflict with privacy (Grudin, 2001). This argument applies equally to video media spaces.

Although there is now a reasonable body of literature that discusses the design problems found in video media spaces, the emphasis thus far has been on generalising about the symptoms observed and then proposing specific countermeasures—point solutions—to offset specific symptoms. Although there has been excellent empirical discussion of the human and technical factors that prompt privacy problems e.g., Bellotti (1998), not all factors are discussed nor are these factors related to one another in a cohesive fashion nor do they completely account for all problems observed. Techno-centric bottom-up approaches do not readily yield insight into how to diagnose privacy problems and predict when they will occur, or provide an intellectual foundation from which to generate new kinds of solutions. Grudin compares ‘bottom-up’ versus ‘top-down’ methods for exploring privacy-design issues (Grudin, 2001). He suggests that while bottom-up approaches readily address technical issues, they demand trial and error to address social issues and are thus too slow and unethical to use for problems like privacy.

Recently, several researchers have begun top-down examinations of privacy-and-design that integrate CSCW findings with theories developed in sociology and psychology. Palen & Dourish (2003) motivate their work on privacy by pointing out that a lack of conceptual frameworks stifles analytical reasoning about privacy and design. Grudin (2001) cautions though that ‘top-down’ approaches suffer from validity and completeness concerns, are complex and time-intensive to produce, and the results can be difficult for designers to

consume incrementally. Theory-based approaches are intended to inform design, yet top-level theoretical abstractions are too abstract to be directly applied to concrete design issues.

These top-down deconstructions of privacy can proceed seemingly *ad infinitum* and it can be difficult to navigate the transition from theory back into design. Also, findings from specific instances of design (empiricism) symbiotically serve to inform theory-making. Hence, in the field today a variety of approaches are being taken to explore privacy as it relates to technology design. In video media spaces and other system areas the HCI research community is steadily progressing towards a comprehensive understanding of the privacy-design link.

### 5.3 Overview of CSCW perspectives on privacy

The CSCW perspective of privacy is rooted strongly in the thoughtful analysis of the impact of technology and its design. This perspective arose from a milieu of self-experimentation: early researchers in video media space design built prototype systems and then used them for extended periods e.g., Mantei et al (1991). By building the technologies the researchers identified and overcame important technological road-blocks, but by living with the technology they came to experience first-hand the privacy consequences of various design decisions and the symptoms of underlying problems. By carefully reflecting on their experiences, these researchers came to intimately understand the relevant technological and individual and social human factors. In this section we build upon Bellotti's (1998) dichotomy of problems, which consists of the following.

- Deliberate privacy abuses are possible.
- Inadvertent privacy violations are possible.

To these, we add a third problem.

- Users and non-users feel apprehensive about the technology.

Although apprehension as a problem theme has received little direct attention we include it here because there is a large body of related research that sheds important insight onto it.

## 5.4 Deliberate privacy abuses: Issues of control

A fundamental premise of much privacy research is that privacy is a thing that can be intentionally controlled (to a limited extent) by groups and individuals. This control is afforded by environmental constraints to interactivity. Technology confounds privacy control by lifting or changing these constraints (Palen & Dourish, 2003; Grudin, 2001). There is an implicit assumption that there are some times when some people—who may or may not be part of the VMS community—go out of their way to violate others' privacy.

Thus, even though video media space users might never willingly violate their peers' privacy the system affords the potential for such **deliberate abuses** (5.b.vi). Worse, media spaces are not adequately designed to safeguard against malicious use arising from unauthorised access. Thus, they afford the potential for undiagnosed abuse by outsiders. One example is surreptitious surveillance: e.g., a thief—or worse, a violent sex offender—intercepts a VMS video stream on the Internet so as to monitor the presence and activities of others as he plots the perfect time to commit his crime.

### 5.4.1 Methods for controlling media space access

One way to solve deliberate privacy abuses is with **access control** (5.a.iv), which puts into place computer security and cryptographic measures to deny unauthorised individuals access to sensitive information (Smith et al, 1995). While access control is common on virtually all computers, those wishing to restrict access have faced a constant and unrelenting battle with those wishing to crack systems.

Another way to solve deliberate privacy abuses is to simply remove sensitive information from the media space so there is nothing of worth for others to access and to reduce the harm that may result if access control measures are defeated. We call this technique **content control** (5.a.v). It is hard to put this technique into practice in a VMS because the purpose of a media space is to reveal (Gaver et al, 1992). There is a fundamental trade-off between privacy and the utility of VMS for awareness: for one person in the media space to have richer awareness, others must have necessarily less privacy (Hudson & Smith, 1996).

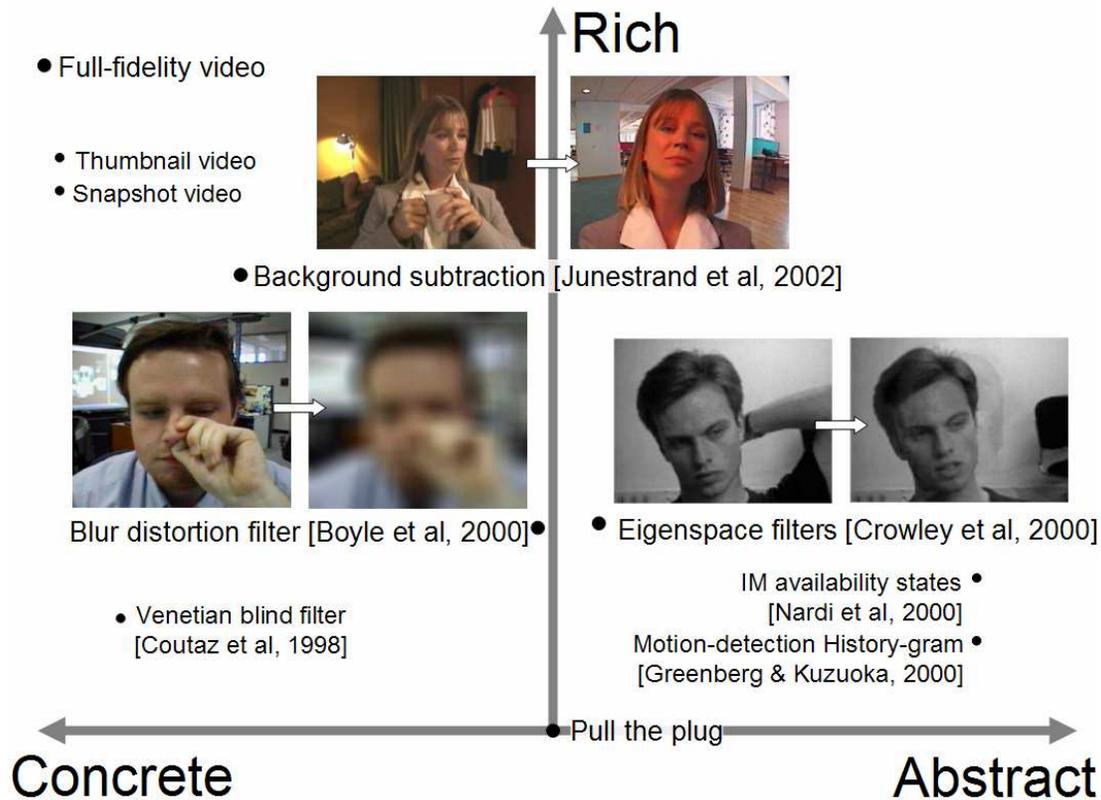


Figure 5.1—A design space showing some previously explored techniques for preserving privacy in video media spaces.

Figure 5.1 shows several techniques for preserving privacy in video media spaces based on content control. Distortion filters such as the blur filter in Figure 5.1 mask sensitive details in video while still providing a low-fidelity overview useful for awareness (Zhao & Stasko, 1998; Boyle, Edwards & Greenberg, 2000). Publication filters such as the background subtraction filter in the figure remove details considered unimportant for awareness information (Coutaz et al, 1998; Junestrand, Keijer & Tollmar, 2001). Finally, potentially privacy-threatening details can be abstracted away from the video altogether such as in instant messenger status icons and in the eigenspace filter in the figure (Crowley et al, 2000).

#### 5.4.2 Control: User interface issues and trade-offs

Both the above approaches involve control over what information is in the media space and who gets to see it. It is hard to design a video media space that provides fine-grained control in a lightweight manner, yet both are vital to preserving privacy (Bellotti, 1998). **Fine-grained control** (5.c.v) can be adjusted on a person-by-person, instance-by-instance basis.

**Lightweight control** (5.c.iv) needs little cognitive or physical effort. In the physical environment, strategies for controlling information access are both lightweight and fine-grained. For example, a person holding a notepad close to his chest prevents all others from seeing it. Yet with a subtle twist he can open it up for the person immediately next to him to see while still keep it mostly concealed from all others (Luff & Heath, 1998). This kind of privacy regulation demands very little cognitive or behavioural effort from the people involved and is usually an implicit activity realised as a natural consequence of the other activities.

There are few fine-grained yet lightweight strategies for controlling a video media space. Unplugging the camera is a lightweight and undeniably effective means for blocking access to all, but it is not very fine-grained. Consider a female worker who wants to offer full-fidelity video to colleagues from both her work and home offices. She wants only some work colleagues to see her at her work location. She also wants another set of (possibly overlapping) colleagues and friends to see her at home, but only when she does not have anyone else in the home office and only during normal working hours (although occasionally seeing her in the early evening is fine). This level of fine-grained control is usually unavailable in the media space. Even if it were, the typical user interface—complex panels of GUI widgets or if-then-else scripted access rules—make configuring the system very heavyweight. As a result, people often do not make changes when appropriate, and often end up configuring the system to grant all others either full access at any time, or no access whatsoever. Unfortunately, these behaviours thwart the security of a system and open it up to deliberate privacy abuses.

Heavyweight and coarse-grained privacy control interfaces prompt an “all or nothing” trade-off. Some users will err on the side of ‘nothing’ and reject the system. While these users avoid the privacy problems, they miss out on the benefits afforded by the system. Other users will err on the side of ‘all’ and forgo best practices of use for convenience, but they must endure the privacy problems that arise.

Control user interfaces must also be **believable** (5.c.iii.1): be readily understood and effect meaningful change in a predictable manner. For example, giving participants the chance to turn the camera around and point its lens out a window affords believable control. As with any direct manipulation UI, the result of the change is also immediately apparent. There is no

disassociation of action and result, as would be the case if the camera was controlled through, say, a command-line or graphical interface.

Control must also be easily interpreted by others. **Dissociation** (5.b.iv), where one's actions become logically separated from one's identity, makes it very difficult for VMS participants to determine who is accessing information about them even though they may be able to tell that it is being accessed (Bellotti, 1998). Dissociation makes deliberate privacy abuses possible because information can be accessed in an unchecked, untraceable, and anonymous manner (Langheinrich, 2001). People have poor strategies for dealing with dissociation because it rarely occurs in the physical environment: one's body, as it is performing an action or gaining access, communicates a wealth of identifying information, coupling action to identity. The predictability of physically mediated access permits institutionalised control and so some deliberate privacy abuses are permitted because the social infrastructure needed to prevent them cannot keep pace with technological advancement (Langheinrich, 2001).

## 5.5 Inadvertent privacy violations

A fundamental premise of the cognitive sciences is that people are mostly rational (Simon, 1996). Rational people will usually protect their own privacy and respect the privacy of others. Undoubtedly, not all privacy violations are deliberate nor are all opportunities for deliberate privacy abuses capitalised upon. Accidental violations are known to happen from time to time. **Inadvertent privacy infractions** (5.b.i) are believed to occur because media space designs fit poorly with individual human and social factors thereby causing breakdowns in normal social practice (Bellotti, 1998). Specifically, privacy regulation is **situated action** (4.e.i) (Suchman, 1987). Environmental constraints for interactivity keep interactions situated in a temporally and spatially localised **context**. Technology changes these constraints, causing actions and interactions to be desituated and **decontextualised** (5.b.iv) (Grudin, 2001). Inadvertent privacy violations occur because people are no longer operating in clearly situated contexts (Palen & Dourish, 2003).

### 5.5.1 Disembodiment confounds self-appropriation

**Self-appropriation** (4.d.i) is a regulatory process where people modify their behaviour and appearance according to social norms and expectations (Bellotti, 1998). Self-appropriation depends on cues for behaviour sense from the environment, such as place and the people in it. For example when a person is at work, he acts, dresses, and speaks to match others' expectations of professionalism. This will differ markedly from how he appropriates himself on the basketball court. As people move between contexts—the office, the bathroom, the hallway, the basketball court, the home—they modify their expectations for social behaviour (**norms**) and adapt their behaviour accordingly.

The impoverished nature of a video media space means that people often do not appropriate themselves correctly for viewing by distant colleagues. **Disembodiment** (5.b.iv)—where a user becomes cut off from the (multiple) contexts of those people viewing him—confounds self-appropriation and leads to inadvertent privacy violations (Bellotti, 1998). Although disembodiment is endemic to computer mediated communication such as in collaborative virtual environments e.g., Benford et al (1995) it has particular implications for privacy.

The **presence** or **absence** of others in the media space is an extremely important cue for self-appropriation, but it too is confounded by disembodiment, specifically the disembodiment of others. A person is entirely dependent on the VMS to feed context cues back to her in order to determine how she should behave. Also, a person is entirely dependent on her embodiment in the VMS to signal to others how they should treat her privacy. Nardi et al (1997) suggest that disembodiment negatively affects the interpretability of others' speech and actions, increasing chances for mis-coordination and miscommunication. This makes it difficult for media space users to understand the privacy wishes of others, leading to inadvertent violations. More broadly, Palen & Dourish (2003) suggest that rich embodiments support performances tailored for the local context (i.e., situated action) through “reflexive interpretability of action” and “recipient(-specific) design” of communications.

### 5.5.2 Presence in multiple places forces appropriation in multiple contexts

Place—its architecture and use (Harrison & Dourish, 1996)—is an important feedback cue for self-appropriation. Places differ with respect to privacy expectations e.g., kitchen versus boardroom. A media space participant always concurrently operates in at least two places: the unmediated one and one or more mediated ones. How one should appropriate herself may differ among these places and so too may cues for self-appropriation. While the unmediated environment—the walls that enclose the room a participant occupies and the lack of visible presence of others in that room—suggest that the participant’s privacy is assured to some extent, this assurance could be completely violated by the virtual environment.

Usually there are physical transitions when one moves between two places supporting distinct privacy cultures: a partition, a doorway, and even distance itself (Altman, 1975; Palen & Dourish, 2003). This transition is a feedback cue for self-appropriation. The time needed to navigate the transition affords opportunity to assess the resulting change in expectations and make changes in appearance and behaviour as appropriate.

Media spaces, though, join places with differing privacy cultures without such smoothing transitions, permitting weird intersections of privacy expectations (Bellotti, 1998; Palen & Dourish, 2003). Video media spaces prompt inadvertent privacy violations because they offer a juxtaposition of places that does not occur easily in real life. Without the transition, people are unaware of the juxtaposition and its impact on self-appropriation.

### 5.5.3 Feedback: User interface issues and trade-offs

The design of feedback channels to support self-appropriation is fraught with technical factors that permit inadvertent privacy violations. It is hard to balance VMS feedback salience and distraction (Gaver et al, 1992; Hudson & Smith, 1996; Bellotti, 1998). If the cues are not saliently presented they will go unnoticed, fostering disembodiment and poor self-appropriation. If they are too distracting there is the risk that the VMS user will either disable the feedback channel or disable the VMS altogether.

It is also hard to design VMS feedback cues for self-appropriation that integrate well with social protocol for conversation initiation. In the physical environment, feedback cues are

given socially natural forms, placements, and meanings. For example, a person in his office can hear, emanating from the corridor, the footsteps of a colleague approaching him to strike up a conversation. This audible cue signals the onset of interactivity (who, when, and where) and there is a rich, socially-based (and often unconscious) protocol for initiating conversations built around this doorway approach. Providing a media space user interface to support this protocol is full of subtle problems. For example, Buxton's DoorCam situates the VMS camera and display near the user's office doorway to provide a more natural placement, but this placement is natural only for the initiation of conversation, after which conversation to be continued is ushered inside the room (Buxton, 1997).

Bellotti presents a framework for analysing deliberate and inadvertent privacy problems in systems and evaluating solutions (Bellotti, 1998). Her framework consists of topic areas for formulating questions about the feedback and control a system affords over information in it and topic areas for evaluating the feedback and control user interface. Bellotti's framework includes **intention** (2.c.ii) for access and minimal needed disclosure as feedback cues that are important to evaluating privacy options. In unmediated settings, intention may be revealed implicitly as a consequence of an attempt to access (prior to access is made) or through explicit (e.g., verbal) communication of it. In either case, the communication process is kept extremely lightweight. It is not lightweight in media spaces. Disembodiment and disassociation confound the implicit signalling of intentionality before access is made. Even if there are audio or text channels, getting everyone into a state where they can use them is not lightweight. Beyond cumbersome user interfaces, networking delays during the initiation of conversation denies quick and graceful transition into it (Tang et al, 1996).

Bellotti's framework focuses on the practice of system design as illustrated through case studies. In this way, her work informs design. Her approach is meaningfully different from that of providing principles regarding the ethical treatment of confidential information (Hochheiser, 2002). Such principles inform the practice of system use. Bellotti is quick to point out that the value placed on privacy fluctuates with social events and that the rapid march of technological advancement causes guidelines concerning design and principles concerning use to "show their age" rapidly.

## 5.6 Apprehension

Privacy violations can be **aesthetic** (3.a.iv.1) —affecting appearances and impressions— or **strategic** (3.a.iv.2) —affecting the execution of plans— (Samarajiva, 1998). In social environments, aesthetic privacy violations can have consequences of a strategic nature. Humans, as social creatures, fear and resent both kinds of violations. Non-users are often so suspicious of the media space that they go out of their way to sabotage the system (Jancke et al, 2001). Even users themselves are often wary about the system’s handling of their privacy (Tang et al, 1994). Thus, in addition to specific deliberate or inadvertent privacy threats, prior analysis of video media space privacy indicates that **apprehension** (5.b.ii) itself is a significant problem. Specifically, participants are apprehensive about making bad **impressions** (3.a.4) in the media space and the aesthetic or strategic consequences of them.

### 5.6.1 Surveillance confounds impression management

A fundamental premise of privacy research in VMS design is that people do not want to look bad in front of others—especially peers—yet they from time to time do and say things that may make them look bad. When we speak of ‘looking bad,’ we mean many things. For example, they may be concerned about being seen with inappropriate or untidy dress (e.g., seen in an office media space changing clothes after jogging during lunch) or behaving in ways that others might judge unacceptable (e.g., seen in a home office media space spanking a disobedient child).

Users are apprehensive about making mistakes that make them look bad in the media space (Tang et al, 1994). Since video media spaces permit detailed, surreptitious surveillance at any time, users must monitor their appearance, behaviour, and speech at all times (Lee et al, 1997). Coping with surveillance requires vigilant self-monitoring, which can lead to errors (Reason, 1990). Worse, VMS technology affords new abilities for automated surveillance and rigorous scrutiny, creating opportunities to make bad impressions with unforgiving, socially-inept computer algorithms that may report misinformation to peers and superiors.

Because there is little in the way of privacy-supporting technology, there is similarly little known about the failures of such technology. There are no rich taxonomies categorising such failures. Neustaedter et al discuss the actions that could be taken by a context-aware media

space as ‘privacy increasing’ and ‘privacy decreasing’ but do not deconstruct these terms in greater detail (Neustaedter et al, 2003). Moreover, it is unknown how accommodating users are of different kinds of failures in privacy supporting technology. Consequently, Neustaedter et al advocate that the system be designed to require users to give explicit consent before taking ‘privacy decreasing’ actions.

### 5.6.2 Decontextualisation prompts apprehension

When short segments of a conversation are examined independent of its totality, examiners are forced to invent contextual information needed to support its interpretation. The invented context can make the speaker look bad. This is yet another privacy-related implication of the decontextualisation of formerly clearly situated action (Grudin, 2001). Nardi et al mention that no one media space channel alone conveys the complete meaning of an event or utterance (Nardi et al, 1997). In the hospital media space in their ethnographic study, neurosurgical operating room staff used humour to relieve the stress had of a mentally demanding surgical task. The decontextualisation of such humour permitted skewed interpretation of it. Thus, the media space put the privacy of nurses and doctors—under the constant threat of malpractice litigation—at risk. Aware of the risk, staff felt compelled to eliminate such humour from their speech. Thus, the media space not only threatened privacy, it reduced *joie de vivre*. It could even be argued that the media space threatened patient safety because unrelieved tension disturbs mental focus and prompts errors in performance.

Nardi et al point out that these kinds of threats are not accounted for by merely providing appropriate feedback. In their experience, participatory design permits the identification and solving of these kinds of problems. It also repairs discrepancies between users’ perceptions of their own involvement in the design process (typically low) and designers’ perceptions of users’ involvement (typically high). This, in turn, reduces users’ resentment over loss of control over their privacy.

Technology affords new degrees of temporal and spatial freedom for information access (Palen & Dourish, 2003). It makes speech and actions that were once fleeting and available to only a few people present at the same place and time accessible to anyone, anywhere, and at any time (Grudin, 2001). For example, it is relatively easy to capture video for later replay and

review as part of a meeting capture and analysis tool (Tang et al, 2003). Recorded speech and video captured actions—even if not archived—can be edited convincingly to make it appear as though one did say or do things one did not, or omit words and actions so as to remove context and mislead or confuse downstream viewers.

## 5.7 Reflecting on the problems

The previous sections show that privacy issues arise out of human, social, environmental, and technical factors. Technical factors weigh heavily in problems related to deliberate privacy abuses, and not surprisingly there are many technical solutions proposed such as computer security and cryptographic methods and the filtration methods described earlier. On the other hand, human factors—especially the interplay between human and technical factors—weigh heavily in problems related to inadvertent privacy violations. There are fewer generalised technological countermeasures for dealing with inadvertent privacy threats than there are for deliberate threats, and there are more high level design problems without obvious solutions. In problems related to apprehension, we see that social factors dominate, concerning the placement of technology throughout society and the psychological aspects of technology use, disuse, and misuse. The discussion of these problems seems messier, vague, and completely removed from the practical matters of designing, building, and deploying a video media space that are immediately apparent when discussing the other problem themes.

More broadly, there is somewhat of a ‘chicken-and-egg’ problem here. Designing privacy supporting technology requires that they be implemented and then evaluated as privacy-supporting. Both the design and the evaluation however require that one be able to operationalise privacy, that is, reduce it to a model that relates observable and measurable inputs and outputs and transformations and decisions performed on them. This model of privacy is hard to uncover through introspection, but can be uncovered by experimenting with privacy supportive technologies. Thus, there is a kind of cyclical dependency produced by the co-evolution of CSCW understanding of privacy and technology and the design, development, and deployment of privacy-intersecting technology. As explained in section 5.2, this cyclical dependency has prompted many to draw theoretical frameworks for understanding privacy

from other disciplines. In the next chapter, we look at some of the diverse conceptions of privacy which can inform the design of video media spaces.

## Chapter 6—Perspectives on privacy

Many disciplines of study must deal with the notion of privacy: anthropology, architecture, behavioural psychology, law, sociology, as well as computer science. Technology designers can learn much from these other disciplines. Thus, the vocabulary we build for discussing the human factors relevant to the privacy-design link in media spaces draws from these varied areas. We begin with a broad overview of various themes in privacy research by drawing from Brierley-Newell's (1995) cross-disciplinary survey of privacy-related literature. She classified works discussing privacy as being 'person-centred,' 'place-centred,' or interested in person-environment interactions. By far, the major of works she examined are interested in the interactions between a person and his environment with balanced emphasis on the roles of each. In this section, we use her taxonomy as inspiration for our own survey of various conceptions of privacy that seem particularly related to the design of video media spaces.

### 6.1 Private/public dichotomy

"Private" is often defined as the opposite of "public:" **public** is to "being together" as **private** is to "being apart." Brierley-Newell found this to be the most fundamental and broadly cross-cultural conceptualisation of privacy (Brierley-Newell, 1998). Being apart is different from being alone. For example, one can be with one's lover and the two together are apart from a larger group. The part of one's life lived apart from society was not highly valued in some ancient societies (Hixon, 1987) and strong emphasis was placed on social involvement. Palen & Dourish (2003) call this the **disclosure boundary** tension (4.d.vii): a tension between one wanting/needing/choosing/being private versus public. This tension carries over to VMS

design. From an organisational perspective, the video media space is seen positively as it strives to increase the amount of ‘togetherness’ experienced by group members, even though the heightened collaboration and cooperative work may not be something desired by all individuals at all times. Because of this tension, there will be times—no matter how well the media space is designed—when it will be considered unwelcome by a user.

Private and public form a dichotomy because they are both inverse and complementary: each may be defined as not the other. This is a pervasive conception of privacy. The Greeks had their *idion* (private life) and *koinon* (public life) (Arendt, 1958). Goffman describes front and back stage performances (Goffman, 1959). Journalists have different ethical guidelines for disclosure of information pertaining to ‘public figures’ versus ‘private citizens.’ Media space literature trades awareness off for privacy e.g., Chapter 4. Schwartz’s (1968) macrosociological analysis of privacy characterises it as a “highly institutionalised counterpattern of withdrawal” complementing a pattern of social interaction.

Although conventional notions of private/public suggest that privacy is important for the satisfaction of personal goals, Schwartz’s analysis suggests privacy also subserves public (i.e., institutional) goals. More precisely, Schwartz suggests that privacy serves to stabilise institutions (societies) in which people are organised into status hierarchies. Horizontally, privacy stabilises relationships between people of the same status by providing them with opportunity to seek leave of and relief from others when too much social contact becomes irritating. Vertically, privacy stabilises the hierarchy by reinforcing status divisions. It also allows high status members of a hierarchy to conceal their flaws from low status members, preserving idealised impressions that reinforce social superiority, authority, and obligatory relationships. Overall, privacy conceals deviant behaviour which, if widely publicised, would destroy social order. Strict conformance is rarely possible and **deviance** (3.b.ii.5) provides the relief to “disobey in private to gain the strength to obey in public” (Brierley-Newell, 1995). A deluge of evidence that society’s rules were being disobeyed would weaken an individual’s resolve to conform.

Little is understood about the effects of video media spaces on these sorts of macrosociological functions of privacy. Consequently, little is known about how to design for these effects.

## 6.2 Privacy as an attribute of places and people

In architecture, privacy is often defined by features of the design and construction of architectural space: for example, the number of enclosing partitions, their height, the windows that make the space visually porous, and the intelligibility of human speech and the loudness of other noises passing through walls and openings. Schwartz notes that this kind of privacy can easily be changed by people (Schwartz, 1968). Treating privacy as an architectural attribute of space is useful for VMS design. It permits construction of architectural metaphors for privacy safeguards (Greenberg & Roseman, 2003), and it informs us that some aspects of privacy can be quantified as observable metrics.

There are other privacy metrics that are not so easily quantified. In particular, architecture not only defines a space, but it creates a social place full of social meaning (Harrison & Dourish, 1996). The social meanings given to a place determine its privacy. For example, public toilets are not very private in construction but can be very private in the sociological experience of their use. This fact has definite implications for video media space design. People perceive privacy in subtle, subjective, and social ways. Yet technology has historically had a hard time observing and quantifying phenomena that exhibit these properties. Furthermore, it is not known if these perceptions are attitudes that are learned (Altman & Chemers, 1985) or are culturally universal aspects of humanity (Brierley-Newell, 1998).

## 6.3 Privacy as an interpersonal process

As part of human experience, privacy is affected and controlled by human behaviours:

- verbal, e.g., telling someone across the VMS link to keep some information secret;
- para-verbal or non-verbal, e.g., pointing a VMS camera out the window; or,
- social, e.g., deciding, as a group, that it is taboo to turn on the VMS camera in the kitchen when someone is already in the kitchen .

One perspective of privacy identified in Brierley-Newell's survey is that these behaviours are part of a **privacy process** (4.b). Altman (1975) in particular sees it as a boundary-regulation

process which facilitates the negotiation of access to the self. The **self** (3.a) broadly refers to the totality of a person: her/his body, thoughts and personality, and information about her/him. The negotiation occurs between the self and the **environment** (4.e): the physical environment and also the social environment i.e., the people immediately nearby and society at large.

Altman's privacy process is a **dialectic** (4.b.i). The actual level of privacy attained is decided through a process of negotiation between the self and the environment. This dialectic is **normative** (3.b.ii). Altman draws a sharp distinction between desired privacy and attained privacy. People's desired privacy is constrained by the environment to socially accepted (normal) levels. What constitutes a **privacy violation** (4.c) is defined against the same set of norms, some of which may be codified as laws while others are part of the culture's tacit knowledge. Individual factors are also important. Each person possesses his/her own set of **privacy preferences** or 'personal norms' that determine his/her initial desired privacy level and subsequently influence the privacy dialect. Also, group norms change in response to changes in group membership and so are influenced by individual preferences. Making things even more complicated, there may be a number of norms that can apply in a given situation because one is typically involved in many groups simultaneously, or because of cross-cultural contact. The relationship between group norms and individual preferences seems complex and co-adaptive.

Altman's privacy process does not deny interactions between the self and the environment rather it **regulates** them. When one has too many interactions or, in other words, too little privacy, these interactions can be throttled. For example, a person turns off the media space to get away from others. When the connections with others have been cut so deeply that one has 'too much privacy' the privacy process can open access to the self so that a person gets the interactions he craves. For example, a person turns on the media space when he wants to chat with others. This process demands skill or, more likely, **power** (3.b.i.2) that not all persons share equally (Brierley-Newell, 1998) and power relationships become significant when addressing privacy problems in VMS design (Dourish, 1993).

Treating privacy as a process is important for VMS design because it permits consideration of observable metrics for evaluating the 'health' of the process. However, much of the process is cognitive, and it is difficult to design context-aware systems that can adapt to

<b>Too few interactions</b> <i>(Too much privacy)</i>	<b>Too many interactions</b> <i>(Too little privacy)</i>	<b>Just right</b>
Loneliness and boredom	Stress and anxiousness	Rest, release of stress
Desperation and hopelessness	Vulnerability to others, i.e., theft	Self identity and self-confidence
Productivity impairment and errors due to boredom	Productivity impairment due to distraction	Fulfilment of fundamental goals
Suicide	Underdeveloped ego	Self-evaluation (social
	Rage and misbehaviour	Accountability and
	“Looping”	Fantasy
	i.e., role separation failures	

Table 6.1—Negative aspects of insufficient control over privacy, and positive aspects of sufficient and necessary control over privacy. From Altman (1975), Brierley-Newell (1995).

changes in the environment affecting the internalised privacy process. It is possible to develop qualitative methods to permit observation of this process, which in turn might support evaluation of the effectiveness of particular media space designs. To this end, Altman’s theory holds potential heuristic value: because it has been specified so broadly, it can apply to many situations. Yet, Brierley-Newell speculates that this broadness also makes Altman’s theory the most criticised. For example, some critics argue that social interactionism may be better able to explain the privacy process. Within the CSCW community, Fitzpatrick’s Locales framework applies social interactionism principles to uncover and comprehend CSCW system design issues (Fitzpatrick, 1998). Although it has yet to be done, it is conceivable that methods for analysing privacy and design in video media spaces could be based on Locales.

## 6.4 Privacy as a need, right, and freedom

Researchers in behavioural psychology have studied individuals who routinely experience compromised privacy, such as the elderly and the mentally infirm living in institutions, and young children. They have characterised the outcomes of failures in the privacy process that yield harmful effects. A few of these effects are listed in Table 6.1. These extreme effects of course do not apply to the general population, most of whom are able to enjoy many benefits from a satisfactory amount of privacy. Some of these benefits are also given later in Table 6.3.

Perhaps because of these benefits people place great **value** upon privacy in our society. Privacy is often defined as a legal and moral **right** and as an inalienable **freedom** that no other

person or institution may lawfully or morally unduly curtail. Privacy is thus legally enshrined in various laws to: discourage “peeping toms,” prevent unjustified search, seizure, and confinement, punish slander and liable, and ensure contractual obligations to secrecy. This fact has relevance for science: Kelvin (1973) discusses barriers to the scientific study of privacy. When so much value is placed upon privacy the scientific manipulation of it for experimentation (needed to understand it) is seen as ‘morally suspect.’ Privacy can also be an equally difficult subject for the law to handle. For example, some advocates have suggested applying intellectual property law to protect the right to privacy but the fit is imperfect (Samuelson, 2000).

A privacy that is a right or freedom can be **violated** (4.c). Others’ actions may deny one this right or impair one’s exercise of it. Thus, it is a privacy violation when others’ actions prevent one from obtaining the privacy he needs, he normally enjoys, and society deems that he ought to enjoy. The normalised definition of violation is important. For example, Schwartz calls surveillance an institutionalised form of privacy violation (Schwartz, 1968). Others’ actions may prevent one from obtaining desired privacy, but this itself may not necessarily be considered a privacy violation. Privacy violations have outcomes such as the effects of too much or too little privacy given in Table 6.3. These outcomes vary in **severity** (4.c.iv), which is a subjective measure of how ‘bad’ the harm due to the outcome is.

Although the environment may permit others actions that will lead to a privacy violation, these people might not choose to invoke such actions. Hence, privacy can be threatened without necessarily being violated. Privacy **threat** (4.c.v) and privacy **risk** (4.c.i) are used almost synonymously and seem to include the **possibility** (4.c.ii) of a violation, the **probability** (4.c.iii) that it will occur, and the severity of the harm it causes. Risk is quite inescapable: abstractly, if there is insufficient control to outright deny the possibility that a violation can occur, then there is some risk. Practically, however, opportunities for violation are held in check by **policing** (4.d.iii): providing punishments, taboos, social consequences etc., to discourage others from doing things that violate one’s privacy. Schwartz cites Simmel’s claim that it is very tempting to intrude upon the privacy of another to deduce that institutions need privacy guarantees (norms) and relief or recourse to handle violations [Schwartz, 1968]. Some privacy violations are so severe that one is permitted actions to stop further harm and be

awarded damages to offset harm already done. Given that privacy violations arising from the deliberate and inadvertent misuse of video media space technology may be inevitable, one way that a design could support privacy is by supporting policing and recovery from violations in addition to providing safeguards to constrain misuse. We note, however, that it is highly likely that policing occurs even without explicit design-time support for it and it is an open question as to what policing behaviours groups employ to govern media space use.

## 6.5 Privacy as a balancing act

There is a tacitly held assumption in CSCW and social psychology that some degree of privacy risk is the inevitable cost of social life. As Palen & Dourish (2003) put it, some level of disclosure is needed to sustain social engagement. There is also a fundamental premise that, stress, tension, and irritation develop with time even among amicable social relations. People must balance the disclosure demands of social life against the privacy needs for maintenance of the self. Moreover, recalling Schwartz's (1968) discussion, a certain amount of privacy is needed to sustain social institutions and social interactions over time.

Aside from hermits and the like, people balance the benefits accrued from social interactions against the risks to privacy, engaging and withdrawing from others to satisfy both the need to be 'apart' and the need to be 'together.' Even though there is risk, there may also be **reward** (4.d.v): benefits to having less privacy than may be possible. There is a trade-off between risk and reward. Grudin (2001) mentions that this **risk/reward trade-off** (4.d.vi) is how privacy issues are resolved by both technology users and designers. He goes on to suggest economic-based decisions about which threads of the local context are captured, at what detail, and how they are presented cause disparities between threads that prompt privacy problems. He specifically mentions disparities in relevance and salience that confound risk/reward analysis. To make the situation even more cumbersome for designers, Bellotti (1998) reminds us that all of design involves making trade-offs and then dealing with the unforeseen consequences of the compromises made.

Another conception of privacy treats the risk/reward trade-off as an economic decision. This emphasises that privacy is not only valuable but also hard to obtain. Schwartz (1968) presents this view, calling privacy a "scarce social commodity" and an "object of exchange to

be bought and sold” indicative of civilisation and social status. He also claims that humans and their societies seem to require a definite ratio of secrecy relative to disclosure and that as a general rule people reveal confidential information in order to obtain some thing or receive some service, emphasising reciprocity and gratification through revelation. Other researchers have attempted to employ economic theory to observations of patterns of disclosure of confidential information e.g., Posner (1981) or incorporate economic factors such as incentive, supply, and demand into privacy-affective technology e.g., Acquisti (2002).

In most human activities reward exists commensurately with risk, yet many video media space designs ignore this relationship altogether. Nardi et al (1997) explain that benefit and threat are not constant; instead, these factors vary independently by person across channels over time. In their ethnographic study of a hospital media space, they found that the occupational roles played by a person determined how he/she used the media space and this in turn greatly influenced risk and reward for that person. Consider, for example, a video media space that connects home offices with corporate offices. Family members (e.g., spouses, children) routinely appear in the video media space but are likely strangers to most others in it and probably do not accrue much benefit from their own participation (Neustaedter & Greenberg, 2003). Thus, some designs bring people together in an indiscriminate way that disregards the need (or lack thereof) for social interaction (Fish et al, 1990; Fish et al, 1992; Greenberg & Rounding, 2001; Jancke et al, 2001).

People balance risk and reward in unmediated interactions but come up against problems when attempting to do so in mediated interactions. The technology itself, the ways it can be subverted, and the awkwardness of its interface may hinder their ability to port unmediated interaction skills to the virtual environment. For example, many video media space designs permit some form of surreptitious surveillance, i.e., the close monitoring the environment—usually the presence and activities of others—without revealing much about oneself. This kind of surveillance can come about from seemingly innocent actions. For example, in the CAVECAT media space a user could cover the camera lens to prevent others from seeing him and yet still see others (Mantei et al, 1991). Video media space designs themselves foster **disparity** (5.c.ii) between risk and reward such that reward does not accrue accordingly with risk or, conversely, risk does rise with reward. This disparity is analogous to

the work/benefit disparity noted by Grudin (1988) that is broadly applicable to all genres of CSCW systems.

**Reciprocity** (4.e.vi) is a simple rule that states that if  $A$  can access  $B$  via channel  $C$ , then  $B$  can also access  $A$  via channel  $C$ . Reciprocity is often enforced over video media space channels as a technological means for re-balancing this risk/reward disparity [Root, 1988]. Yet, reciprocity does not always hold for the physical environment, and sometimes breaking the reciprocity rule is beneficial. For example, it is possible to observe a person to deduce her/his **availability** (1.c.1) —willingness to engage in interaction— without disturbing her/him, such as by moving quietly and peeking around the corner of an open office doorway. Some VMS designs, such as the RAVE media, have explored privacy regulation in the absence of reciprocity but these design experiences underscore the need for multiple modalities of support for privacy in any one given system and across systems (Gaver et al, 1992).

Furthermore, Nardi et al (1997) found that reciprocity does little to address risk and resentment users and non-users develop towards the technology. They noted that these feelings often follow professional allegiances, much as Harper (1996) found regarding Active Badge use. Instead, Nardi et al (1997) recommend that designers perform careful analysis to determine the risk and reward for each person in the media space. Extra attention must be paid to conditions in which risk is high yet benefit is low. Unfortunately, the real-time diagnosis of such situations is difficult. For example, in interviews with operating room staff, the authors found that the most privacy-sensitive verbal exchanges were the ones least relevant to the surgery being performed. These exchanges were sensitive not because they revealed confidential information but because of the impression management problems that arise when they are presented out of their original context. Relevancy, however, can be just as difficult for computers to measure as sensitivity.

## 6.6 Summary: Focusing on an interpersonal process model

In this chapter we surveyed a number of phenomenological perspectives on privacy that have been cultivated in disciplines such as anthropology, psychology, and sociology. We premised our discussion on the belief that these varied perspectives each uniquely inform the design of privacy-supportive technology. What we have seen is that privacy can be:

- a basic human need,
- an institutionalised phenomenon,
- a state in which people find themselves,
- a quality of places, and
- a behavioural process governing interactions that seeks to balance risks and rewards associated with social interactions.

These perspectives on privacy can be integrated but this integration is not trivial. Privacy involves various aspects of the physical environment, human psychology, and social behaviour for in the maintenance of self and the regulation of social interactions. Bellotti (1998) contrasts normative definitions of privacy with operational ones. She points out that since operational definitions focus more on the capabilities people have for regulating privacy they are better suited for deconstructing the control and feedback problems in video media spaces. Like her, we will focus on operational aspects of privacy in the current chapter but our deconstruction will reflect a fundamental assertion that privacy behaviours follow normative, institutional and situational patterns.

Of the perspectives offered, the one pervasive in environmental psychology—that privacy is a process—holds great appeal because it accounts for the other perspectives as well. As already mentioned, Altman broadly characterises privacy as a boundary-control process regulating access to the self (Altman, 1975). This conception of privacy as a control process relates strongly to the overwhelming importance of feedback and control identified in CSCW literature on privacy in video media spaces e.g., Dourish & Bly (1993); Bellotti, (1998). Not surprisingly, it is also the foundation selected by Palen & Dourish (2003) in their deconstruction of privacy and technology design confluence.

Altman's is a theory of **interpersonal privacy**, and it makes a fitting selection for our uses because most of the privacy problems reported by media space users and researchers tend to be of an interpersonal nature. People—e.g., media space users and researchers, and law makers—seem to be highly skilled at rationalising about interpersonal privacy problems in highly situational local contexts. In such circumstances, people behave in ways that closely

match their preferences. People do not seem to be very skilled at rationalising about macrosociological privacy problems, or about problems that span wide temporal or spatial boundaries. While privacy is very important to people, their behaviours (for example, in e-commerce) often contradict spoken opinion (Spiekermann et al, 2001). While privacy is very important to institutions, it sometimes goes unprotected because it is also very hard to make good laws to protect it. Although there are macrosociological privacy problems inherent in video media spaces very little discussion of them exists upon which we may base our vocabulary. For now, the present discussion is constrained more or less to interpersonal privacy processes.

It is by no means an easy task to apply Altman's theory of privacy to the problem of designing a privacy-supporting video media space! Although it is a rather contemporary theory, it was developed long before the widespread deployment of the computationally powerful personal computer, sophisticated audio and video compression algorithms, high-bandwidth low-latency high-reliability multimedia internetnetwork, and massive rapid random-access storage facilities that are the technological infrastructure for video media spaces. Technology's threat to privacy is materially different today than from Altman's time and, unsurprisingly, Altman's theory largely ignores the privacy-technology relationship. Although there is an appealing sense of validity in seeking to inform design with conceptual frameworks that predate the problems faced today, the transition from this theory to design is not trivial. Altman's description of the process is extremely abstract: both the boundaries and the mechanisms by which they are controlled are purposefully left ambiguous. Yet it is exceedingly difficult to apply his theory to the problems faced by video media space designers without some manner of concrete link to tie it in. An important part of the value of the work done by Palen & Dourish and ourselves here is to provide these concrete links.

## Chapter 7—An integrated vocabulary for privacy and video media space design

In addition to the specialisation of Altman's theory, we must necessarily make some elaborations to it because as it stands Altman's theory does not account for all of the problems reported in VMS research. The over-arching elaboration to Altman's theory that we make incorporates Gavison's (1980) decomposition of privacy into three basic elements. *Solitude* relates to understanding how a person regulates social interactions. *Confidentiality* relates to understanding how a person manages others' access to information about her/himself. *Autonomy* relates to understanding how a person chooses to present her/himself when alone or in social situations.

Gavison emphasises the role of control in privacy management and that genuine control requires both an abundance of options to choose from and the power to ensure that one's choice is respected by others. Her discussion is rooted in law and the design of legislation to protect privacy. We call legislation a design problem so as to underscore parallels to the problems faced by technologists. Gavison's decomposition of privacy yields a powerful vocabulary which we use to disambiguate the many interrelated meanings of privacy discussed by Altman. Specifically, we transform Gavison's basic elements into modalities by which people control the self-environment boundaries described by Altman. We subsequently expand them to cover more of the problems encountered in video media space research.

The three control modalities we have found are:

- **Solitude**: control over one’s interpersonal interactions, specifically one’s attention for interaction (1).
- **Confidentiality**: control over other’s access to information about oneself, specifically the fidelity of such accesses (2).
- **Autonomy**: control over the observable manifestations of the self, such as action, appearance, impression and identity (3).

Casting these components of privacy as controls makes the discussion directly relevant and immediately applicable to understanding the problems researchers face in designing and building privacy-supporting video media spaces. In section 2 we cited many discussions that attribute privacy problems in video media spaces to inadequacies in control and its exercise. Moreover, by discussing privacy in terms of controls, we deconstruct the mechanical aspects of self-environment boundary regulation and ignore the much more difficult deconstruction of the boundary itself. This complementary approach has been taken up by Palen & Dourish (2003). In their framework they identify three boundaries which are congruent to but not direct parallels of the three modalities of privacy control we describe here. The **disclosure boundary** (4.a.i) is regulated mostly by confidentiality, but also by solitude. The **identity boundary** (4.a.iv) is regulated by autonomy. The **temporal boundary** (4.a.ii) spans both identity and disclosure and is regulated by the norms and preferences that are part of solitude, confidentiality and autonomy.

What is **control**, anyway? Dennet (1995) gives a technical description: “*A* controls *B* if... *A* can drive *B* into whichever of *B*’s normal range of states *A* wants *B* to be in”. Gavison points out two elements in control: the ability to make a choice (implying that a number of alternatives exist to select from) and the power to ensure the choice is respected. Control can be exercised through a normative dialectic as per Altman’s theory. Such control is founded upon individual and social human behaviours such as those discussed by Altman & Chemers (1985) and Langheinrich (2001). These behaviours are thus the low-level mechanical means by which control is exerted. An implication of a dialectic sort of control is that the processes are **satisficing**: there is no need for complete control in order to experience privacy.

All three modalities of control are negotiated concurrently. Behaviours used to exert one modality of control also have strengthening and weakening implications for the other two. Moreover, the privacy-related actions of one individual operate concurrently with those of all other individuals. Altman's notion of **attained privacy** is thus the net effect of all these mutually, complementary and competitively interacting privacy-affecting actions.

Privacy controls in our vocabulary are, as per Altman's theory, social. As soon as social interactions of casual or work topics are made possible—be it by spatial propinquity or by a media space—the role of privacy must be considered because privacy fundamentally concerns the regulation of these interactions. In addition to affording new opportunities for people to 'be together' when they want to feel connected with one another, a media space intended to support privacy must also afford opportunities for people to 'be apart' when necessary and to affect social relationships in intended ways. It is important that technology mirror intentionality because patterns of use and disuse of social technologies (e.g., video media spaces) convey social meanings that affect social relations (Harper, 1996). For example, in heterogeneous video media spaces where some users may not have cameras, such users are sometimes thought by others (with cameras) to be spying on the community (Coutaz et al, 1998).

Privacy is also a co-operative process. A person will sometimes do things to respect others' privacy and to help others respect his own privacy. Sometimes, though, it is difficult for a person to show respect for another's privacy yet also make that person feel included (Schwartz, 1968). While privacy violations occur regularly, gross privacy violations seem to occur less often than the environment permits. Sometimes group members take advantage of opportunities to violate the privacy of other group members, but often they do not. Given that group privacy is contingent upon the privacy of its individual members, some group members may even take steps to protect the privacy of other members and defend the group as a whole against outside intrusion. For example, even if Mike does not get a chance to close his office door before his lawyer calls him regarding a sensitive topic, his colleague Saul may sense Mike's privacy needs and close his door for him. This cooperative view of privacy differs markedly from the competitive view (common in computer science) which assumes that if an opportunity to violate privacy arises it will necessarily be capitalised upon. While this more

extreme competitive view may be useful for evaluating a system's fortifications against deliberate privacy violations, it can also lead to user interface designs which encourage inadvertent violations. For example, few VMS designs allow one user to protect another's privacy by changing her settings on her behalf, losing out on opportunities to defend against inadvertent privacy violations.

In the next three sections we will delve deeply into each of the three modalities of control—solitude, confidentiality and autonomy—to complete the construction of an integrated vocabulary for privacy. Each discussion starts broadly, with particular emphasis placed on human behaviours and the psychological and sociological processes related to the modality of control. As the human concepts become more fully expressed, we weave in factors related to VMS design, illustrating the relationship between environmental psychological theory of privacy and human life and CSCW theory of privacy and technology.

## 7.1 Solitude

Altman describes solitude when he discusses control over interactions between the self and the environment, particularly other people. Solitude controls help a person 'be apart' from others and is involved in many behaviours that are vital to human development, e.g., self-evaluation and ego development (Altman, 1975). As previously mentioned, we clarify that being apart is different from being alone: for example, two lovers can find solitude in each other's company, even in a crowded restaurant. 'Togetherness' is thus a continuum of states, and the extremes present failure conditions that yield negative behavioural, psychological, and physiological responses. For example, **crowding** (1.a.i.1) results when others are granted too much access to the self. **Isolation** (1.a.i.1) results when one cannot interact with others to the degree they wish. Both conditions indicate failures in solitude control.

### 7.1.1 Attention and distraction

To discuss other issues in video media spaces that closely relate to solitude, we generalise Altman's definition of solitude to include control over where one directs one's **attention** (1.a.ii) and how one controls **distraction** (1.c.ii). Most video media spaces require that users expend extra effort to attend to awareness information by presenting it in ways that potentially

distract or disrupt people. Thus, media spaces confound solitude. Thus presence and availability are regulated by solitude. Our broadened conception of solitude makes it strongly related to Rodden's (1996) model of focus for awareness. Although not originally conceived to tackle privacy problems, the Rodden model also relates to the other two modalities of privacy control and will be discussed later.

This extension also helps to explain 'camera shyness' problems in video media spaces (Lee et al, 1997). In co-located settings, people track the focus of others' attention as an informal awareness cue that helps determine availability. In particular, a person notices if another is looking at her, i.e., that she is becoming the object of others' attention. This prompts her to reflexively focus her own attention back upon herself, to monitor self-appropriation and track others' impressions. This state of heightened self-awareness can cause discomfort if maintained for prolonged durations (Duval & Wicklund, 1972). In our extended notion of solitude, technology can invade users' solitude (or permit users to invade others' solitude) by making it difficult for users to control how they direct their attention for self-work and interactions.

### 7.1.2 Verbal and para-verbal solitude controls

A variety of individual and social behaviours are used to regulate solitude. Verbal and para-verbal mechanisms for controlling solitude usually involve signalling availability, e.g., verbally telling another you wish to be left alone or hanging a 'do not disturb' sign outside a hotel door. Desires can be signalled in both the content (the meaning of the words spoken) and the structure (pitch, duration, volume etc. of voice) of speech (Altman & Chemers, 1980). Para-verbal means for signalling one's desired solitude include a posture or facial expressions and explicit gestures to beckon or dismiss others. While these mechanisms are very lightweight in face to face settings, they are easily impaired by limitations of VMS technology. For example, low-quality video (i.e., low resolution, low frame rate, many visible artefacts of compression) mask subtle para-verbal cues for communicating availability. Because such desires must instead be communicated with speech, video media spaces can make the process of signalling solitude desires more explicit and heavyweight. These changes alter social interpretation of the expressed desires.

### 7.1.3 Westin's four privacy states

Westin, another noted privacy theorist, decomposed privacy into four states (Westin, 1967).

— Solitude is a state of total isolation. (Note that Westin uses the word differently from what we have presented.)

— **Intimacy** (1.b.i.1) is the state in which a small group (e.g., lovers) isolate themselves from others.

— **Anonymity** (1.b.i.1) is the state in which one is physically co-present with others and yet not expect to be recognised by them and so free from interactions with them. It refers to a condition in which one can be “lost in a crowd.”

— **Reserve** (1.b.i.1) is the state in which we can ignore the presence of others who are nearby. It entails the use of psychological controls to shut out others. (Another meaning for reserve is personal restraint in dialogue and action to constrain interactions with others.)

We consider these four states to be four particular points along a spectrum of social interactions arising from typical exercise of solitude.

### 7.1.4 Affordances of space for solitude

To regulate solitude, one can go someplace to be alone. These places of **refuge** (1.b.ii.1) are where one can seek solitude and also safety from the stresses incurred through interactions with others. Refuge is needed for psychological repair (Altman, 1975). VMS design complicates refuge-seeking. Although places of refuge from the media space are typically nearby—it is prohibitively expensive to put cameras in every room and so the media space is usually present in only a few locations—the media space is usually present in a person's personal office. Awkwardly, the office is where most will retreat to find refuge. A place of refuge can be created by ‘pulling the plug’ on the video media space (Neustaedter et al, 2003). Unfortunately, this disconnected mode of operation is often misinterpreted in many media space implementations as an exceptional error case to which little developer attention is given. Consequently, most hardware and software infrastructures make reconnection so complicated that users are disinclined to ‘pull the plug.’

Conversely, when one craves social stimulation, one can go to places where others are. Place partially determines **accessibility** (1.c.i.2) i.e., the effort people must expend to engage others for interaction (Harrison & Dourish, 1996). Architectural spaces can often be reconfigured to raise or lower their permeability to light, matter, and sound. In changing these attributes, people control the affordance of space for interactivity. For example, an office door can be closed to reduce visual and auditory distractions from the corridor and serve as a physical barrier to others' entry. Doors permit fine-grained control because they can be fully closed, slightly ajar, or wide open. Indeed, this becomes a social cue indicating one's solitude desires. In contrast, video media spaces generally provide only one modality for interactivity (an audio/video channel) and offer few ways to configure this channel to signal the desired level of engagement.

People can also capitalise upon the ambiguity inherent in some architectural changes to regulate solitude. For example, a closed door ambiguously symbolises both absence as well as a wish to be left undisturbed (Root, 1988). People also capitalise on ambiguity when it is possible in computer-mediated environments. For example, Nardi et al (2000) reports that people use the inaccuracies of IM presence indicators as form of "plausible deniability," where they ignore requests for conversation from people because they know that the other person will be uncertain if they are really there.

### 7.1.5 Personal space

Space and social behaviour interoperate with respect to solitude. **Personal space** (1.a.i) refers to an invisible boundary in space around a person, separating him from others. The boundary's shape and size varies from moment to moment as part of the privacy dialectic. Although the boundary's characteristics are never made explicit, people show definite behavioural and physiological responses when others physically enter their personal space. **Territory** is similar, but usually implies a recognisably fixed spatial or psychological location, even if it is defined relative to its owner. Territories are important for the regulation of workspace artefacts and confidentiality and will be discussed later.

Distance	Modality	Interaction capabilities
Public distance (>5m)	Gross vision	Gross assessments of posture and large gestures; facial expressions and gaze not visible
Social distance (<4m)	Hearing	Speech content and structure
Personal distance (<2m)	Detailed vision	Posture; gestures; gaze; facial expressions involving eyes and mouth (e.g., wink, smile)
Interpersonal zone (<0.5m)	Touch and smell	Exchange, inspect, and manipulate artefacts; physical contact (e.g., handshake, hug); perfume

Table 7.1—Interpersonal distances and the interactions supported at each (Hall, 1966).

Personal space regulates solitude by reducing sensory stimulation due to the presence of or interactions with others. This, in turn, affects attention. At each distance, different sensory capabilities afford different modes for interaction. Hall describes four interpersonal zones, each with differing modalities for social interaction; these are given in Table 7.2 (Hall, 1966). Because of this relationship between distance and interaction, distance itself becomes imbued with social meaning (Altman, 1975). For example, consider when one person sits down at the same table as another. If the newcomer sits diagonally across the table and out of direct eye contact, he sends a solitude-related message that differs markedly from when he chooses to sit directly across the person and in easy eye contact.

Personal space, as a tool for solitude regulation, depends on having a range of **interpersonal distances** (1.a.i) at which people may space themselves. These distances define modalities for interaction that differ in both affordances for interaction and the attention or engagement needed to sustain such interactions. These distances are thus imbued with social meanings. Typically, in a video media space the camera position and display size dictates the visual distance between people; these are sometimes arbitrary and do not represent the desired social distance. For example, seeing a tightly cropped face shot on a large video monitor places someone visually close, but the mannerisms exhibited by that person may reflect actions of someone who is in fact quite far away. The concept of interpersonal distance in a VMS can be even further generalised to include engagement and connectivity. In a typical VMS, only two or three such distances are offered: full interconnectivity; connected to just one other person; and, disconnected from everyone. The limited choices for connectivity make the media space a crude tool for the selective expression of social interest for interactivity. Moreover, in

physically co-located settings, adjusting distances is very lightweight and can be continuously adapted by just moving around. In contrast, media spaces offer highly discrete choices selected using heavyweight GUIs and limit degrees of freedom, e.g., it is awkward to reposition the VMS camera because of limited cable lengths, lighting, shelf space, and similar factors.

## 7.2 Confidentiality

Confidentiality is the control of access to information about oneself, e.g., informal awareness cues, intentions, vital statistics, thoughts and feelings, medical history, criminal record. Controlling access is as much granting access as it is restricting it. **Secrecy** is similar to confidentiality but narrower because secrecy emphasises that the information is concealed from certain people. Secrecy modulates the communication of information to others, but this is only one aspect of confidentiality. Palen & Dourish (2003) use the term **disclosure** to describe deliberate control over what information is communicated, to whom, when, and how.

Confidentiality and solitude are of course related. Confidentiality directly regulates the outward flow of information and thereby indirectly others' attention, whereas solitude directly regulates one's own attention by indirectly regulating the inward flow of information from others. As noted earlier, there is a fundamental tension between confidentiality and the goal of the video media space to reveal informal awareness cues (the disclosure boundary tension described by Palen & Dourish). Hence, there is tension regarding confidentiality in the design of a video media space. Confidentiality and autonomy are related as information yields power to affect livelihood (e.g., coercion, competitive advantage), personal safety or autonomy (e.g., interference or intervention).

### 7.2.1 Sensitivity

**Sensitivity** (2.b.i.1) is a property of a piece of information that can be defined as a perception of how important it is to maintain control over access to it (Adams, 2000). Others' impressions of a person are predicated upon their knowledge of her, and so confidentiality is part of impression management (Goffman, 1959). The harms that could arise from breaches of confidentiality include embarrassment, damage to ego and identity, loss of others' esteem, and possibly impairment of livelihood. Video media spaces can, of course, easily reveal sensitive

information when they unintentionally capture and transmit a person's image that, for example, shows that person in a socially unacceptable act.

### 7.2.2 Fidelity

**Fidelity** (2.b.ii) is a perception of how faithfully a piece of information represents some truth. It includes both **precision** (2.b.ii.1) —how detailed the information is perceived— and **accuracy** (2.b.ii.2) —the confidence or certainty one places in the information, or the error in its perception. The same essential truth or description of circumstance may be perceived at a variety of fidelities. Also, people's perceptions of the fidelity of information about a person are situated in the context of the whole history of social interaction with that person (Palen & Dourish's temporal boundary). Information about oneself—the object of confidentiality—may be known by different individuals at different fidelities. Our conception of confidentiality is broadened to address VMS design issues by considering that confidentiality includes control over fidelity. Confidentiality is breached when a person is unable to control the fidelity at which others are able to access her/his information.

Video media spaces have several dimensions for video fidelity e.g., field of view, resolution, frame rate, codec quality, latency, jitter, etc. Technology places an upper bound on most of these parameters, and these bounds are usually much lower than in face to face situations. For example, although a person can move his head or body to very easily change his field of view to encompass virtually any area around that person, the field of view in a video media space is typically fixed because the cameras lack pan/tilt/zoom capabilities.

Despite these upper bounds, video is nonetheless a high-fidelity medium for informal awareness and casual interactions. This is both part of the appeal of video and a source of confidentiality problems. Undoubtedly, video offers more fidelity than is genuinely needed in many scenarios, even between intimate collaborators. Consequently, many video media space designs try to preserve confidentiality by discarding fidelity. The premise is that appropriate blurring can find a balance by providing just enough awareness information to be useful, while not too much to violate confidentiality (Figure 7.1). These techniques presume that sensitive information lays mostly in image details and so low fidelity overviews of the video pose less risk (Hudson & Smith, 1996). The manipulation of fidelity (especially timeliness) introduces

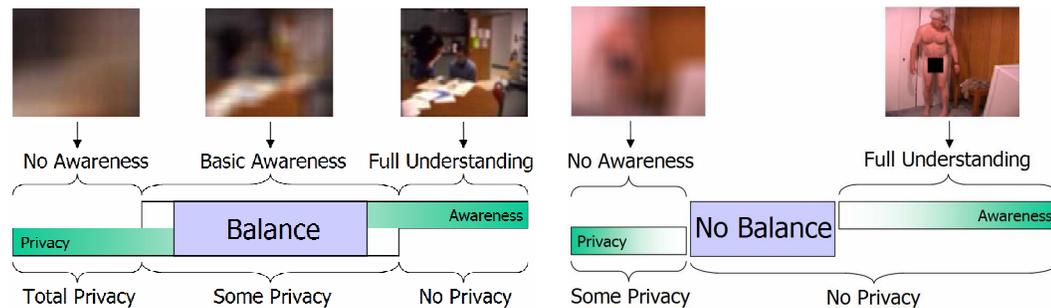


Figure 7.1—The blur distortion filter can operate at a variety of levels. Each level affects fidelity and risk, which in turn affect awareness and one’s ability to control confidentiality. The left part of the figure shows a mundane scene used in the Boyle et al filter study for which privacy could be balanced with awareness. The right part of the figure shows a risky scene used in the Neustaedter et al filter study for which privacy and awareness could not be balanced.

ambiguity that is incorporated into privacy control e.g., “plausible deniability” (Nardi et al, 2003). For example, distortion filters such as the blur filter shown in Figure 7.1 can operate at many levels, discarding a little or a lot of fidelity (Chapter 4). Of course, while fidelity is reduced, there is no guarantee that these techniques mask the sensitive information. For example, the second study in Chapter 4 questioned the effectiveness of a blurring video filter in extremely risky home telecommuting scenarios. There we found that the filter preserved privacy in only mundane scenes and the filter alone was ineffective at masking sensitive details from very risky scenes.

The perceived fidelity of information is not static. It is influenced by the **trust** (3.b.i.5) one places in the sender and the number of recipients. We also consider that information has properties such as **persistency** (2.b.i.2) and **transitivity** (2.b.i.3) that are relevant to confidentiality. It may change when it is transmitted between people, such as through oral or written statements or when it is permanently recorded. Hence, confidentiality also involves the regulation of the fidelity of information that third parties transmit about us. A significant factor responsible for the decontextualisation problem reported by Grudin and by Palen & Dourish is that the digital encoding of contextual information changes it in specific ways, some of which alter fidelity. Receiving a data transmission may increase the perceived fidelity of information, especially if it was previously not known and persisting data that is otherwise fleeting increases its perceived fidelity when it is reviewed. Imagine, for example, that Mike and Saul participate in a media space that archives the video streams and Mike thinks he saw Saul

passionately kiss someone who is definitely not his wife. If the video was not archived, Mike would be left with lingering doubts, but archival changes the persistency of the information and permits scrutiny which yields a more accurate (i.e., higher fidelity) view of the event.

There are larger theoretical issues here, as well. Grudin (2001) points out that persistence leads to the temporal separation of action and its effect or, as Palen & Dourish (2003) put it, a tension across temporal boundaries for performances. There becomes the logical separation of the text of a performance and its context, i.e., the audience. A common theme among theoretical frameworks for privacy is that typical privacy regulation behaviours assume people act in clearly situated local contexts and changes to this assumption can make the regulatory behaviours ineffectual.

### 7.2.3 Direct controls

Mechanisms for regulating confidentiality overlap greatly with those for solitude, emphasising their synergistic relationship. The principle means for confidentiality control involve keeping our bodies, possessions, and thoughts accessible to some but inaccessible to others. We consider possessions because things like diaries, driver's licenses and even automobiles reveal a great deal of sensitive information about a person and are used to mark status and individuality (Schwartz, 1968). Territoriality and personal space use distance to afford fine-grained control over others' access to our bodies and our things, e.g., the notepad example given earlier (Luff & Heath, 1998). Similar control is available over speech: a person directs his voice and modulates its volume so as to whisper into the ear of someone nearby without allowing others to hear what is said. This same technique is also used to preserve the solitude of others: for example, people whisper at the cinema because they do not want to disturb others. Private vocabularies can be used to talk openly among others yet obscure what is being said: e.g., pig latin among children and hand signals in baseball.

Architecture also plays a vital role in the preservation of confidentiality (minimising leaks out) as well as the preservation of solitude (minimising leaks in). Walls reduce access via visual and auditory channels. Walls may also be fortified with sound-proofing materials to preserve aural confidentiality as well as solitude. Window blinds may be raised or lowered and doors closed or open to modulate visual confidentiality. Video media spaces afford similar

opportunities for regulating confidentiality, e.g., turning down microphone volume so as not to be overheard, encoding information with cryptographic methods so others cannot eavesdrop, or using one of the filtration techniques in Chapter 5. These techniques, however, suffer from the feedback and control issues discussed in Chapter 5.

#### 7.2.4 Computers and confidentiality

Increasingly, computers are being used to store or transmit confidential information and **computer security** (5.a.i) holistically addresses many aspects of confidentiality. **Authorisation** (5.a.iv.1) is control not only over access, but also use i.e., a person's intention for using the system or the information it provides, or outcomes of access. **Data integrity** (5.a.vi.1) concerns ensuring that persisted information about oneself is not modified or transmitted information is not modified en-route. Both of these are obviously part of confidentiality. **Process integrity** (5.a.vi.2), availability, responsiveness, and reliability concern ensuring that computers perform their intended function when requested correctly and completely in an expected amount of time with no undesired side-effects. Process integrity is an important component of confidentiality because, as stated in the introduction to this section, confidentiality includes ensuring a person has all the access he/she has been granted. **Cryptographic methods** (5.a.ii) are used to provide access control and verify the identity of the receiver or sender of information and check the integrity of the message (e.g., with digital signatures).

Users generally have a hard time rationalising about computer security, and so they unwittingly fall prey to **malware** such as data-destroying viruses, service-denying worms, and **spyware** Trojan-horse software that offers some benefit of use but covertly gathers information on a computer user's habits, such as which web sites he visits and which music tracks he play. Computer systems may afford defences against confidentiality threats, but if the control is heavyweight or feedback inappropriate these confidentiality-preserving features may actually interfere with usual processes and behaviours. Users often deliberately circumvent computer-supported confidentiality when they become a nuisance. For example, security measures are often incomprehensible to set up and use, or they require great effort, or they do not supply sufficient feedback for people to know what is actually being transmitted (Balfanz

& Simon, 2000). Because this interferes with their work tasks, people often thwart computer security measures. For example, instead of carefully configuring access control lists for network shared files and folders—granting and revoking privileges on an as-needed basis—users often open files and folders up for full access by everyone, completely negating the value of the facility. Although VMS systems such as RAVE (Gaver et al, 1992) and CAVECAT (Mantei et al, 1991) included expressive languages for controlling access, little is known about the utility or usability of these kinds of user interfaces.

### 7.2.5 Indirect controls

People explicitly state (verbally or para-verbally) their confidentiality desires and perceptions on information sensitivity. For example, one person can tell another to “Keep this secret, okay?” Telling a person that it is important to keep a piece of information secret does not prevent that person from revealing it to others. Yet, people can choose to—and sometimes do—keep others’ secrets. People can intuit others’ sensitivity perceptions and from these infer self-imposed limits to behaviour. In contract law, stiff penalties dissuade breaches of confidentiality. The law also enshrines confidentiality in certain relationships, e.g., doctor/patient and lawyer/client, such that desirable limits are placed on the judiciary’s access to information obtained from questioning such confidants. Silence and ambiguous speech ensure confidentiality.

Information about others, including confidentiality preferences, are usually revealed over time. One gets to know another better with each subsequent interaction and access to information about another person accrues with the amount of social “work” invested to build and maintain the relationship with that person. Palen & Dourish (2003) introduce temporal boundary regulation by pointing out that future disclosures and interactions are patterned after past ones. A dialectic sort of privacy implies that the temporal context—including norms—is important to its regulation. Palen & Dourish introduce the notion of **genres of disclosure** (4.d.ii) to capture not only institutional (socially constructed) expectations regarding confidentiality but also situational ones that change with the temporal boundary. That is, genres of disclosure are loosely defined patterns of interactions that evolve over time. Indeed, the role of time in the evolving practice of privacy is significant and complex, motivating

further examination beyond the scope of this thesis. Anyway, because genres of disclosure are loosely defined it is possible to feel that one's privacy has been violated through others' misappropriation or misuse of confidential information and not just inappropriate disclosure.

The risk/reward trade-off mentioned before guides not only an individual's control over her/his own confidentiality but also how he/she treats the confidentiality of others. For example, breaching a close colleague's confidentiality could foster distrust that might break down the relationship. Institutionally, breaching a patient's confidentiality could cost a physician her/his license to practice medicine. Preserving privacy allows one to reap the rewards of social interaction and the denial of these rewards can act as a psychological mechanism for conforming to another's confidentiality desires. Obvious caveats to these claims—e.g., “blabbermouths”—exist but these do not detract from their generality. While people can keep secrets or assess sensitivity, a particular individual may not keep a secret well, or may ultimately choose not to respect the apparent sensitivity. Of course, the VMS may change the rules of engagement. For example, a VMS might permanently archive video/audio exchanges for later replay, rendering requests to keep information confidential meaningless. Verbally telling those people present to keep matters confidential does not preclude others from listening in later. By the same token, people willingly and unwittingly spread **misinformation** —unintentionally inaccurate information— and **disinformation** —intentionally inaccurate information designed to obscure the truth, i.e., lies. Technological safeguards against these kinds of confidentiality violations will probably never be perfectly effective. It might not even be desired to have perfect safeguards. For example, disinformation can be an important tool for protecting confidentiality when no significant harm results from its spread but significant harm can result if the truth is spread as in the telling of ‘little white lies.’ Confidentiality must regulate unintentional disclosures so that they do not weaken such disinformation. Nonetheless, it is important to incorporate into the VMS design various awareness and interaction channels that can be used to diagnose, police, and reprimand wilful and damaging violations.

## 7.3 Autonomy

Collectively, the freedom to choose how one acts and interacts in the world (freedom of will, also liberty) and the power to act in such a way are taken as the third modality of privacy control: autonomy. In law, **personal liberty** (4.e.vii) is often used synonymously with autonomy. Self-appropriation, described earlier, and autonomy point to the same basic control—control over one’s own behaviour—yet, autonomy incorporates behaviours that facilitate **self-definition** and identity. Accordingly, Altman places great emphasis on the importance of self-definition and the role privacy plays in it. As suggested by Table 6.3, autonomy and identity afford vital rewards for ego development. Many of the symptoms of privacy problems in video media spaces that were discussed previously can be blamed on systems’ poor support for managing behaviour, identity and impressions. Thus, an understanding of autonomy—which regulates these things—is needed to design a privacy-preserving VMS.

### 7.3.1 Preserving and constraining autonomy

Autonomy is like the ‘muscle’ of privacy in that it must be routinely exercised or it will atrophy. The simplest mechanism for preserving autonomy is to try to do as one wishes. One can communicate to others how important it is that he be allowed to do precisely as he wishes. Such signalling may be explicit in the content of speech or implicit in the structure of spoken language, facial expressions, and posture. Informal awareness cues for availability simultaneously reveal one’s autonomy desires.

Autonomy violations are often the most unbearable. Schwartz describes walls and doors as partitions that permit individuality. The violation of these barriers implies a loss of control over access to the self (Schwartz, 1968). He goes on to suggest that in order to be true to oneself one must deceive others, that is the public self must be sufficiently distinct from the private self to keep the two separate. Partitions permit this separation of front and back regions. In Schwartz’s analysis, doors are clearly more valuable than walls. Doors imply regulated separateness (freedom) analogous to Altman’s notion of porous boundaries between people. Walls imply forced separation (loss of freedom).

Autonomy can be impaired when technology robs media space users of the opportunity to choose when and how they participate in the media space community. While there are cases in which media space participation is effectively mandated by an organisation's culture in such cases the social fabric of the organisation has evolved through an extended period of use (Harper, 1996). Introducing video into home offices also engenders several different kinds of privacy fears, one of which is related to loss of autonomy. One of the advantages of working from home is the ability to set one's own schedule. Home workers often work at irregular times outside the typical "9 to 5" hours to better accommodate the demands of family life they hope to balance by working at home in the first place. A video media space that connects home and corporate offices blurs the clear separation between one's presence at home and one's presence at work. This could introduce social pressure to schedule one's activities at home to fit the work context, effectively robbing them of the opportunity to decide when they work.

Exercising autonomy does not imply that one "always gets one's way." Although the sanctity of autonomy is enshrined in law—people are granted the rights and freedoms needed to enjoy life, each according to her/his own will—both autonomy and our legal entitlement to it take part in a dialectic based on group norms. Each may do as he/she wishes, so long as her/his actions conform to group expectations. Indeed, as part of the normal regulation of autonomy, one routinely adjusts one's behaviour so that one may live cordially among others. Doing so ensures that long term plans come to fruition even if they are not done strictly as planned. This is essentially self-appropriation. Thus, autonomy is generally constrained rather than compromised by group norms. Yet, if group norms change faster than people can adapt, or insufficient feedback about the presence and activities of others is offered to support self-appropriation, autonomy can be compromised.

These constraints to autonomy illustrate how privacy controls are synergistic. Consider the following scenario in which Saul and Mike use a video media space to connect with one another. Saul's schedule today will alternate between working intensely on his own and discussing confidential matters on the telephone. Mike needs to chat for a half-hour with Saul about an upcoming deadline. Saul can trade his confidentiality off for his solitude if he uses the media space to provide Mike with sufficiently high-fidelity informal awareness cues so that

Mike can choose appropriate times to contact him. Similarly, Mike can put off engaging Saul for conversation—even though he really does not want to wait—to ensure that he does not disturb Saul and ultimately so that Mike can interact with Saul for the full length of time desired. Saul's availability becomes a constraint and a cue that helps Mike regulate his autonomy.

This example underscores that in video media spaces, privacy can be preserved by the judicious reveal of informal awareness cues, contributing to the disclosure boundary tension (Palen & Dourish, 2003). People mix deliberate disclosure of availability with deliberate concealment. Obviously, disclosure increases accessibility and so some unintended disclosures confound solitude, but the authors acknowledge that, as in our example, some deliberate disclosures actually limit accessibility. A tension arises because it is never immediately clear how little can be disclosed while sustaining interaction or how much can be disclosed without confounding solitude or confidentiality. This idea of appropriate disclosure increasing privacy is the foundation of work on distortion filtration (e.g., Chapter 4) yet often prior work contradictorily plays privacy and awareness off each other in direct opposition.

Beyond self-imposed limits to autonomy, others may directly constrain it. For example, institutionalised people often incur great losses in autonomy (Altman, 1975): drugs or physical restraints are used to prevent injury to themselves, staff, or other residents. Autonomy is constrained to enforce social protocol. Parents often restrict the autonomy of their young children to keep them safe and to socialise them (teach them how to behave properly in society). Barriers are erected to restrict access to dangerous places, or places where confidentiality is demanded, or prohibit certain behaviours in communal spaces: e.g., no smoking in restaurants. Constraints to autonomy are the primary means for punishing bad behaviour: adults who commit crimes are incarcerated and children who disobey their parents are grounded. These observations have implications for VMS design. Fundamentally, the single user interface to a social technology like video media spaces eliminates social governance of its use.

Media spaces allow people to transcend geographic constraints on observation and interaction, providing rewarding opportunities for remote collaboration but at the same time introducing problems as discussed earlier. Video media spaces do not erect barriers to

constrain users' autonomy so that they do not violate group norms. For example, a media space that connects home and corporate users is generally unable to switch its cameras off if the home worker appears in a bath robe. Disembodiment obscures feedback about the presentation of self, confusing decision-making regarding autonomy. Placing a mirror next to the camera intends to remedy this problem by showing a person how she actually appears to others. Yet, this is only a partial solution because the mirror shows nothing about the norms that drive self-appropriation.

### 7.3.2 Autonomy-confidentiality-solitude symbiosis

The second way in which autonomy is like the muscle of privacy regulation is that it provides people with the power to enact their privacy choices, i.e., to control information access and direct attention for interactions. Solitude and confidentiality intrinsically depend on autonomy in a readily understood way. Yet, the converse is also true: one cannot have autonomy without solitude and confidentiality. Solitude is needed for self-reflection and the formulation of future plans (Altman, 1975). Solitude also affords a person with confidentiality needed to perform socially unacceptable acts. Confidentiality is also needed to preserve autonomy when others can use privileged information to thwart one's short- and long-term plans. Because of the symbiotic relationship between solitude, confidentiality and autonomy, when a VMS design impairs the regulation of one kind of control, the other two may also be negatively affected. For example, when cameras are ubiquitously embedded into every corner of our physical space, their pervasiveness makes it difficult for people to find opportunities to be apart from others (i.e., regulate solitude) and thus limits choices for autonomy where they cannot do some desired behaviours because they are being watched.

Some important autonomy-related terms can be borrowed from Goffman's (1965) framework for self-presentation. People are **actors** who have **fronts** (3.a.i) which serve as conduits for the social expression of self and team identities. A front is manifested in actions, utterances and interactions as well as various verbal and non-verbal **signifiers**(3.a.iii): **social setting** such as location, scenery, props; **appearance** such as costume and props, posture, expressions, gestures; and, **manners**. These signifiers have social meanings which contribute to the front. As such, fronts can become institutionalised and the audiences' expectations of a

front become part of the front itself. Fronts are carefully constructed and maintained (for example, by confidentiality) to ensure homogeneity between **performances**. The **back** (3.a.ii) is a secondary presentation of the self to only the team (for team fronts) or the individual her/himself. Here, deviance occurs and the self is maintained.

Bellotti's framework discussed previously focuses on the **usability** of video media space control and feedback affordances to support the kind of self-appropriation process developed by Goffman, in particular about what contextual information is captured by the system. Much is known about the expected **utility** of awareness cues (i.e., feedback) needed to support group interactions. Comparatively little is known about the expected utility of privacy control mechanisms. In this regard, the terms in our vocabulary borrowed from Goffman help. Many of the signifiers he discusses—both subtle and obvious—are visual in nature. His framework establishes the theoretical footing for linking visual information and impaired control over visual confidentiality to problems in autonomy, confidentiality, and solitude. These links inform the design of techniques to modulate the fidelity of specific visual signifiers (e.g., scenery, props) with an understanding of their utility, that is, the kinds of privacy problems the techniques can be specifically expected to address.

### 7.3.3 Identity

We broaden our conception of autonomy to include control over **identity** (3.a.i.1) and its expression, e.g., a person's likeness (visual physical appearance and mannerisms, and the sound of one's voice) and names (e.g., signature or seal). National identity cards, passports, driver's license, credit cards, and so forth are tangible artefacts revealing identity. These exist separately from a person's body and may be held in possession or reproduced by others. Electronic equivalents include email addresses, personal web pages, and network IDs. These make up part of one's **digital persona** (3.a.i.2) (Clarke, 1994). While there are legal safeguards to discourage others from mishandling one's conventional identity, such as civil penalties for libel or unauthorised use of one's identity to promote a product or service, these are still sadly lacking in the electronic medium. With no recourse to reprimand violators, computer system users must turn to **privacy-enhancing technologies** to protect their online identities, usually by preserving the confidentiality of one's digital persona (Burkert, 1998).

Identity is highly relevant to VMS design. Dissociation relates to identity because the virtual embodiments of people—which signal presence and afford means to interact with others and access information about them—do not, unlike our corporeal bodies, reveal identity. Computer security also relates. **Impersonation** (5.b.vi.4) is the act of assuming the identity of another, usually without authority. **Identity theft** (5.b.vi.3) is a form of impersonation that usually involves theft of documents used to **authenticate** (confirm the identity of) an individual. Confidentiality guards against this type of crime, but vigilance is required to keep identifying information and authenticating documents out of the hands of malicious individuals. Just as reserve promotes confidentiality, minimising the amount of identifying material that exists physically separate from an individual preserves her/his control over her/his own identity. Detractors of national identity cards often use a similar claim: reducing one's identity to a single, physically separable and easily reproducible form facilitates identity theft. Oddly enough, certain privacy-preserving techniques used in video media spaces can create situations that confuse identity. For example, distortion filters that greatly blur an image, or substitute actors in the video with stock images can make one person unintentionally appear as another (Crowley et al, 2000).

#### 7.3.4 Pseudonymity

A person is typically involved in a number of intersecting and disjoint social worlds. One maintains an identity for each such world. Although we can recycle much of one identity for another, keeping distinct identities separate is a core privacy task. **Pseudonyms** (5.a.ii) are alternate identities which one creates and uses for interactions with an environment. Pseudonymity is one mechanism for keeping identities separate. Often, each identity is used in a distinct social world and little is revealed that relates one identity to the others. Transportation and telecommunication technologies facilitate pseudonymity by allowing social circles to extend across large geographic ranges and population bases, decreasing the likelihood that a person who is part of one social world is also part of or communicates with members of another. Also, some telecommunication technologies permit anonymity by allowing one's interactions with the environment to proceed in a way that limits the reveal of identifying information. Video media spaces are at odds with pseudonymity because much identifying information is communicated in the video image of one's face and body. While video

manipulation techniques could conceivably replace a person's real visage with an artificial one, such algorithms are tricky to implement in practice, require considerable setup for creating replacement images for multiple identities, and likely reduce the value of the video channel for expressive communication.

### 7.3.5 Role Conflict

People often assume different **roles** (3.b.i.1) as they move between social worlds. A single person may have the role of a stern leader when working with underlings, a supplicant when working with her boss, a parent when with her children, a lover when with her mate, and a slob when alone at home. **Role conflict** (5.b.v) (Adler & Adler, 1991) can result when previously non-overlapping social worlds collide and one is forced to assume two previously distinct roles simultaneously, exposing each to people whom one would rather not. The classical example of role conflict in the non-mediated environment is when parents go to visit their children at their college dormitory: the children must simultaneously play the role of 'children' in the eyes of their parents and 'adults' in the eyes of their peers.

Role conflict can be a major problem in video media spaces. The purpose of the media space is to connect physically distributed people, but its users will likely inhabit quite different physical contexts. By virtue of connecting two physically disjoint spaces—each embodying their own, possibly different sets of privacy norms—the media space creates opportunities for role conflict akin to problems with self-appropriation. Moreover, there is an analogue of role conflict for privacy norms: decontextualisation confuses which norms apply in a given circumstance (Palen & Dourish, 2003). These problems are particularly evident when the VMS connects both home and corporate offices. The home worker must simultaneously play the role of an office worker (because he is connected to the remote office site), a disciplinarian parent and intimate partner (when children or mates enter the home office) and a relaxed home inhabitant (when he is alone at home and forgets he is connected). Role conflict fosters opportunities for inadvertent privacy violations and contributes to the apprehension participants feel towards the media space.

### 7.3.6 Focus and nimbus

The tripartite conception of privacy as presented can be reinterpreted using Rodden's (1996) focus/nimbus model for awareness. While not developed for privacy, the symbiotic link between awareness and privacy suggests that it could serve as a model for privacy regulation and negotiation. **Foci** correspond roughly to attention and so solitude can be thought of as foci regulation. **Nimbi** correspond to embodiments and socially constructed personas, and to one's relationships with information and artefacts in the environment. Regulation of nimbi therefore roughly corresponds to confidentiality and autonomy. **Awareness** (1.c.1), which is a functional composition of focus and nimbus, is analogous to the dialectic negotiation of privacy boundaries.

Rodden uses set notation to describe focus, nimbus, and awareness and the operations that can be performed on them. This abstract representation decouples awareness, focus, and nimbus from conventional spatial metaphors ascribed to them. It also makes it conceivable that his model might someday be incorporated into quantitative methods for analysing privacy. Other quantitative methods drawn from economics limit analysis to confidentiality, while holistic methods are often highly qualitative. Even though privacy is composed of qualitative phenomena it is still very appealing to have some reliable quantitative methods for analysing it. Yet, the development of these models is entirely non-trivial in part because the Rodden model does not account for some important topics:

- normalised, institutionalised character of privacy expectations;
- history of interactions as predictor for future interactions;
- technological factors such as disembodiment, dissociation, and decontextualisation;
- information properties like fidelity and sensitivity;
- apprehension and self-appropriation;
- role conflict; and,
- policing and reprimand

Most importantly, the Rodden model itself does not provide guidance concerning how user interfaces for controlling foci and nimbi should be designed.

## 7.4 Conclusion

Chapter 5 motivated the need for a comprehensive descriptive theory of privacy embodied as a vocabulary of terms to support unambiguous description of how privacy is affected by video media space design and use. Chapter 6 presented various perspectives on privacy that grounded the development of this descriptive theory. Finally, this chapter described privacy in VMS design as a process that intends to regulate the interactions between a person and her physical and social environment. The process consists of three synergistically interrelated modalities of control.

- **Solitude:** control over social interactions, specifically control over the allocation of attention for interaction and engagement.
- **Confidentiality:** control over information access, specifically control over the fidelity at which others access information about oneself.
- **Autonomy:** control over one's own behaviour and the expression of identity.

The controls are exercised as part of a normative dialectic that utilises well-understood environmental constraints for interactivity as affordances for privacy regulation. The dialectic is highly situated action and incorporates contextual cues that may be communicated explicitly or consequentially as people work alone and interact.

Through the three chapters of this act, we have discussed some of the myriad of ways video media space technology disrupts privacy regulation. Principally, technology lifts or changes environmental affordances and constraints for interactivity such that privacy regulating behaviours fail or are compromised. The changes affect the signalling and perception of situational privacy cues, causing interactions to become decontextualised in time, space, and privacy norms. Technology also alters social perception of an individual's action. As a result, technology permits both deliberate and inadvertent privacy violations and prompts apprehension about the presentation of one's social self.

At the conclusion of this act, let us revisit the thesis problem and goal developed in Chapter 1:

**Thesis Problem #3:** There is no comprehensive vocabulary of privacy terms—one that integrates conceptions and theories of privacy from many disciplines—to support unambiguous description of how privacy is affected by video media space design and use.

**Thesis Goal #3:** Integrate privacy theories and observations from many disciplines of scientific inquiry to produce a vocabulary for describing privacy and a video media space's effect on it in an unambiguous and comprehensive manner, accounting for at least the privacy issues reported in previous literature.

**Status of Goal #3:** Completed.

This chapter builds a vocabulary for talking about privacy in a holistic yet unambiguous way. The vocabulary is informed by theoretical frameworks for understanding privacy drawn from CSCW, environmental psychology, sociology, and behavioural psychology. Although it has been given a broad theoretical footing, the vocabulary is incomplete because scientific understanding of privacy in individual and social human life is still incomplete. Nonetheless, this holistic perspective **fundamentally changes our understanding of the nature of privacy** and its role in people's daily lives. This will, perhaps, be the most significant contribution to come out of this thesis.

Rather than deal with specific issues, the focus here has been on providing the vocabulary to communicate the totality of privacy. I chose to attack this problem because I found I lacked sufficient theoretical understanding of the nature of privacy to be able to design a prototype video media space that I felt confident might preserve privacy enough to merit evaluation. The perspective of assembling a descriptive theory of privacy came only after I immersed myself in sociological psychological literature on privacy and tried to apply what I learned to the problems described in CSCW literature on video media spaces.

Although there has been a considerable corpus of work relating privacy problems to the design of social technologies, there is tremendous work yet to be done to advance the state of our understanding from individual words that describe privacy, to axioms that explain what 'privacy-preserving' means, to models that will drive the design and verification of privacy

supporting social technologies. Nonetheless, if design can be thought of as a discussion about the way things are and the way they ought to be, then this act and the next will show that there is value in assembling a vocabulary to facilitate such discussion.

## PREFACE TO ACT III

### USING THE DESCRIPTIVE THEORY OF PRIVACY TO ANALYSE VIDEO MEDIA SPACE DESIGN

The chapters of the previous act assembled a descriptive theory of privacy embodied as a comprehensive vocabulary that permits unambiguous description of privacy-related phenomena and issues connected with the design of video media spaces. This chapter concerns the value of that theory for design.

The central idea posited here is that a comprehensive descriptive theory of privacy supports design by permitting unambiguous explication of tacitly held assumptions hidden in a design or its implementation. In making these assumptions visible, the theory helps designers understand the merits and demerits of the design, implementation, and evaluation of privacy safeguards in their systems, and circumscribe constraints on the circumstances of use in which their safeguards will be successful. This understanding in turn indicates directions for further iterative or exploratory design.

There are limits to this theory. It does not lead directly to design ideas, their implementation, or their evaluation. Rather, there are many kinds of theories, each with a distinct capacity to inform design. As a descriptive theory, the privacy theory I have assembled informs the analysis method of the design, implementation, and evaluation of systems. In particular, systematic analysis may reveal assumptions hidden in the design, or the implementation, or the evaluation. A method for such systematic analysis is a valuable, unique, and important contribution.

This act presents this method. Chapter 8 develops the method while Chapters 9 and 10 present case studies that illustrate the method in use. My primary goal in doing these case studies is to show the application of the method as a means of illustrating the value of the descriptive theory upon which the method is predicated. The particular observations and insights gained by performing the analysis are considered a secondary goal.

## Chapter 8—Analysing video media space design and privacy

In Chapter 1, I identified a problem to address in this thesis that concerns how the descriptive theory can be applied to understand privacy in the context of a video media space:

**Thesis Problem #4:** There is no systematic method for applying the concepts in the privacy vocabulary to understand and inform the design of privacy-preserving video media spaces.

**Thesis Goal #4:** Develop a systematic method of applying the terms in the privacy vocabulary to describe and analyse the effect of a video media space's design and use on privacy.

In this chapter I seek to develop a method for the systematic analysis of video media space design and privacy. I want the method I develop here to satisfy the following criteria:

—applicable to each phase of the design process, i.e., during task analysis, idea generation, prototype design, implementation, or evaluation.

—directly incorporates a vocabulary that embodies the descriptive theory;

—conforms to the ways descriptive theories can and cannot inform design; and,

The kinds of objects of analysis to which the method will be applied are varied:

- techniques, schemes, and algorithms for privacy safeguards;
- prototype video media space and privacy safeguard implementations;
- laboratory evaluations of these systems and their safeguards; and,
- frameworks that guide the design of privacy-supporting systems.

The procedure I develop is repeatable in the sense that any analyst can use the method to uncover hidden assumptions. However, each analyst following the procedure may produce different results. That is, the procedure is used to guide and structure the analysts inquiry, but still relies on the analysts using their insight to interpret how best to apply the theory.

This begs the question: is the power in the theory or in the theorist? Undoubtedly, there are many deep and meaningful concepts given in the theory and the vocabulary that embodies these empowers the analyst to describe the privacy-design relationship in a video media space. The initial Partition and Describe steps of the analysis procedure I develop here are rote, and similar results are expected for any analyst applying the procedure, according to each analyst's familiarity with the vocabulary terms. The Reveal and Summarise steps are more open ended and the results are expected to vary according to the analyst's insightfulness. In the case studies that follow, the analyst is me: the same person who compiled the theory. My goal here is to illustrate the method being applied by *someone*, and not expressly resolve the question of whether the method empowers *everyone* to produce similar results. I will take this matter up in Chapter 11 when I discuss directions for future work that come out of this thesis.

## 8.1 Theories that inform design

A typology is way to analyse things by organising them into categories. For example, McGrath's circumplex model of group tasks is an example of a typology (McGrath, 1984). Taxonomies are like typologies in that they are both categorisations, but usually taxonomies refer to hierarchically organised categories that afford increasing speciation with descending levels in the hierarchy. I am introducing a typology of theories to characterise the power of my descriptive theory and how that power can and cannot be soundly used. The typology is based on Knudtzon, Thomas & Shneiderman (2000). It lists five operational categories of theories, each differentiated by the kinds of analytical operations it permits. The typology is depicted in

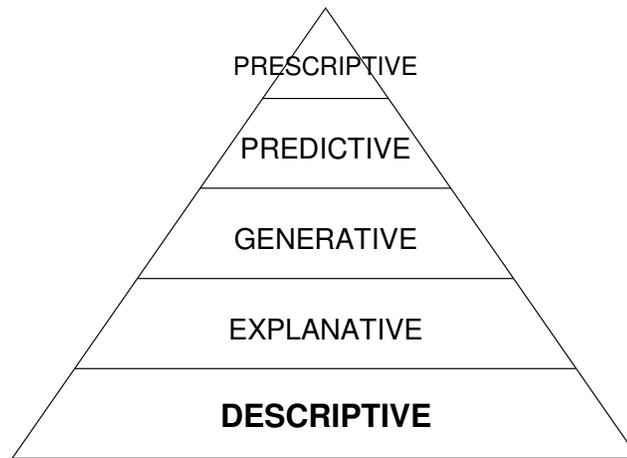


Figure 8.1. Pyramid of theories.

Figure 8.1 as a five-layered pyramid of theory types: descriptive, explanative, generative, predictive, and prescriptive.

This layered arrangement of the theories is premised on the idea that no single theory is genuinely sufficient to inform all of design. Instead, design involves many different kinds of analytical thinking. Theories in the bottom-most layers of the pyramid are the most comprehensive in terms of the diversity of phenomena discussed, but are also the least focused on solving design problems. Layers above build upon the knowledge provided in the layers below. While each level narrows the range of phenomena or factors included it provides more and more guidance about how things ought to be designed.

**Descriptive theory:** permits unambiguous description of phenomena. An example of a descriptive theory from another discipline is Bohr’s model of the atom which can be used to describe the atomic structure of elements using terms such as “electron,” “nucleus,” “charge,” and “orbit.” (As it turns out, Bohr’s theory of atomic structure includes other kinds of theories, too, and I am merely accentuating here its descriptive nature.) A descriptive theory of privacy provides a vocabulary of terms to deconstruct and describe the larger phenomenon into its constituent, interrelated parts. The label “descriptive theory” is sometimes used elsewhere in a much broader sense. For example, scientific theories are often categorized by a dichotomy of descriptive versus normative (i.e., prescriptive) theories.

**Explanative theory:** permits logical explanation of why phenomena occur. An example of an explanative theory from another discipline is the kinetic theory of gasses which explains why gases have pressure as a result of being composed of tiny particles repeatedly colliding with other particles and with the walls of the gas container. An explanative theory of privacy would, for example, explain why people find cameras in the home to be privacy invasive, such as by providing a set of cause-effect relationships among phenomena. An explanative theory is generally narrower than the descriptive theory that informs it because in many cases it is possible to disentangle related phenomena to give each constituent part its own unique name, but not yet possible to piece them back together with orderly causal explanations. Explanative theories are more focused on design than descriptive theories because they permit the causal linking of choices made during design to observations gathered about privacy and the design.

**Generative theory:** permits generation of prototype designs or hypotheses for experimentation. It provides an organisation of measurable phenomena yielding the dimensions of the “design space” and their scales. Experimentation and iteration is possible with the systematic modulation of these design or experimental factors so as to cover the interesting permutations. An example of a generative theory from a related discipline is the classical same/different place/time typology for groupware (Dix et al, 1998). The label “generative theory” is often used by others in a much broader sense. For example, linguists speak of descriptive versus generative theories of language. Generative theories are typically narrower than the explanative theories that inform them because some of the causal factors discussed in an explanative theory cannot be directly observed and the tools and methods needed to measure have not yet been developed or they are immeasurable. Generative theories are more focused on design than explanative theories because they give discernable shape to the design space and facilitate systematic exploration of it.

**Predictive theory:** permits prediction of the response of some phenomenon under different conditions. It provides a set of formalised functional relations among measurable phenomena. The classic equation  $E=mc^2$  is an example of predictive theory: it can be used to predict the energy potential of a nuclear warhead. A predictive theory of privacy would predict the effects of a particular design decision on human behaviours such as preference or opinion. A predictive theory is narrower than a generative theory which can enumerate it because a

predictive theory typically accounts for only a small subset of the dimensions categorised in the generative theory. Predictive theories are more focused on design because they help designers set expectations for their designs and track their progress meeting them.

**Prescriptive theory:** permits prescription of a formula of recommendations for improving a design. It provides rules for design decisions, typically guidelines or heuristics, or methods for identifying, prioritising and resolving design problems. A prescriptive theory from another discipline is a set of guidelines on environmentally-sensitive design of human communities. A prescriptive theory of privacy could provide resources such as a list of minimum requirements for control or feedback, as in the case of Fair Information Practices (OECD, 1980). A prescriptive theory is narrower than the predictive theory which informs it because it concerns only how things ought to be in the typical case, rather than trying to account for how things might be in arbitrary cases. Prescriptive theories are more focused on design because they state explicitly the criteria that designs must satisfy and often give a procedure for getting the design to satisfy them.

All five forms of analysis permitted by the theories in this typology can be incorporated throughout an iterative design process. Descriptive and explanative theories help designers understand the problem domain. Generative theories help designers recognise opportunities for innovation. Predictive theories help designers make sound selections among competing design ideas. Prescriptive theories help designers identify areas for improvement between iterations of the design. These systems of knowledge represent only one small part of the process of design. Beyond theoretical knowledge, design capitalises on such things as practical experience, creativity, inspiration, aesthetic sensibilities, instruments, components, methods and heuristics. It is also important to recognise that the distinctions between practical and theoretical knowledge become blurred as we focus more on the design task. For example, prescriptive theories are most often used in frameworks that include not only other kinds of theories but also guidelines, measurement instruments and methods.

Although completeness and correctness are caveats to consider when evaluating this typology, these concerns are somewhat moot insofar as the purpose for which I am using it. Even if there are systems of knowledge about privacy and design not categorised by this typology, the typology nonetheless shows that there are kinds of analysis that are hard to do

with a descriptive theory alone. Even an incomplete or incorrect typology makes clear some of the things a descriptive theory cannot be used for. In the next section, I build a method for utilising the descriptive theory to deconstruct and describe notions of privacy “embedded” in a system or safeguard’s design.

## 8.2 Revealing assumptions hidden in design

In this thesis, I take **assumptions** to be statements of knowledge<sup>2</sup> about the world and a design problem in it which are used to logically inform design decisions. As such, assumptions play a vital role in design. The world is very complex; probably too complex to design effectively for with our present tools and methods. Sometimes, the complexity is found in subtleties that are peripheral to the design problem (e.g., users, their contexts, and their goals). Other times, the complexity fully permeates every aspect of the design problem. In either case, designers strip out the unwieldy complexities while striving to keep many of the core properties of the design problem intact. In doing so, they create for themselves simplifying abstractions or approximations of the design problem and then work to solve these. They hope that the elided complexity will not meaningfully affect the workability of the final design in actual use but, in the case of privacy, this is not often possible.

Thus, designers ought to take great interest in understanding the assumptions they use to inform their designs. These assumptions are not intrinsically “bad things” indicative of sloppy design. They are necessary predicates that facilitate design and successful design requires the careful construction of good assumptions *as well as* the careful construction of artefacts that fit with those assumptions. There is a problem, however: assumptions are usually hidden. Their existence may be tacitly known and left unacknowledged or they may be altogether unknown. Their revelation brings the designer both joy and burden: there is never a clear path to a new assumption if an existing one is found wanting, and if an assumption must be amended or rejected then the decisions it informed must be also reconsidered and the implementation of those decisions must be reworked.

---

<sup>2</sup> “Knowledge” bears little relation to “truth:” assumptions may be always correct, always incorrect, or any conditional variation between these polar extremes.

The descriptive theory of privacy and design produced in Act II gives designers a rich vocabulary with which they may describe how their systems support or deny privacy regulation: the mechanisms by which the support is realised, and the circumstances in which it is and is not realised. Through the process of assembling this description, designers may start to piece together an understanding of what must be true about the world—the system, its users, their goals and their contexts—in order for privacy to be preserved. This knowledge is the logical ground for the design. These statements about the designers’ expectations and understanding of the design problem *are* assumptions embedded in the design.

The challenge I undertake in this chapter is to present a method by which designers can systematically describe the interrelationship between privacy and their systems so that they may themselves begin to understand what they have assumed about privacy—constraints as to what is and is not part of privacy, and biases about the importance of various parts of privacy—when making their design.

### 8.3 Analysis vocabulary

The procedure by which the case studies presented in this chapter will be analysed is rooted very deeply in the descriptive privacy theory developed in Act II. As explained earlier in this chapter, the descriptive theory provides a vocabulary—compact handles for complex concepts—for describing systems and their effects on privacy. The method I develop in this chapter directly consumes the theory by way of its vocabulary.

Table 8.1 lists the vocabulary terms I will use in the subsequent case studies. This table was developed by enumerating terms introduced in Act II roughly in the order in which they are presented. A user of this table will therefore have to know and understand what these words mean. Those with a deeper understanding will likely be more perceptive when applying these terms, compared to those with a shallow understanding. The table itself adds value, for it brings together strongly related terms that are presented far apart in Act II for matters of comprehension and style

Table 8.1 Analysis vocabulary used in the case studies.

- 1) **SOLITUDE**
  - a) Physical Dimensions
    - i) Interpersonal Distance
      - 1) isolation to crowding
    - ii) Attention
      - 1) focus to periphery
  - b) Psychological Dimensions
    - i) Interaction to Withdrawal
      - 1) anonymity and reserve to intimacy
    - ii) Escape
      - 1) refuge
      - 2) fantasy
  - c) Presentation Dimensions
    - i) High-level Awareness
      - 1) availability
      - 2) accessibility
    - ii) Distraction
      - 1) relevance
      - 2) salience
- 2) **CONFIDENTIALITY**
  - a) Information Channels
    - i) Medium
      - 1) aural
      - 2) visual
      - 3) numeric
      - 4) textual
    - ii) Processing
      - 1) sampling
      - 2) interpolation
      - 3) aggregation
      - 4) inference
    - iii) Topic
      - 1) information about the self
      - 2) personally identifying information
      - 3) activities
      - 4) whereabouts
      - 5) encounters
      - 6) utterances
      - 7) actions
      - 8) relationships
  - b) Information Characteristics
    - i) Basic Characteristics
      - 1) sensitivity
      - 2) persistence
      - 3) transitivity
    - ii) Fidelity
      - 1) precision
      - 2) accuracy
      - 3) misinformation
      - 4) disinformation
- iii) Certainty
  - 1) plausible deniability
  - 2) ambiguity
- c) Information Operations
  - i) Basic Operations
    - 1) capture
    - 2) archival
    - 3) edit
  - ii) Use
    - 1) accountability
    - 2) misappropriation
    - 3) misuse
  - iii) Scrutiny
    - 1) surreptitious surveillance
    - 2) analysis

### 3) **AUTONOMY**

- a) Social Constructions of the Self
  - i) Front
    - 1) identity
    - 2) digital persona
    - 3) appearance
    - 4) impression
    - 5) personal space
  - ii) Back
    - 1) flaws
    - 2) deviance\*
    - 3) idealisations
  - iii) Signifiers\*
    - 1) territory
    - 2) props
    - 3) costumes
  - iv) Harms
    - 1) aesthetic
    - 2) strategic
- b) Social Environment
  - i) Social relationships
    - 1) roles
    - 2) power
    - 3) obligations
    - 4) status divisions
    - 5) trust
  - ii) Norms
    - 1) expectations
    - 2) preferences
    - 3) social acceptability
    - 4) conformance
    - 5) deviance
    - 6) place

### 4) **MECHANICS OF PRIVACY**

- a) Boundaries

- i) disclosure
- ii) temporal
- iii) spatial
- iv) identity
- b) Process Characteristics
  - i) dialectic
  - ii) dynamic
  - iii) regulation
  - iv) cooperation
- c) Violations
  - i) risk
  - ii) possibility
  - iii) probability
  - iv) severity
  - v) threat
- d) Behavioural and Cognitive Phenomena
  - i) self-appropriation
  - ii) genres of disclosure
  - iii) policing
  - iv) reprimand
  - v) reward
  - vi) risk/reward trade-off
  - vii) disclosure boundary tension
  - viii) disinformation\*
  - ix) reserve\*
  - x) Signifiers\*
    - 1) implicit
    - 2) explicit
- e) Environmental Support
  - i) situated action
  - ii) reflexive interpretability of action
  - iii) constraints
  - iv) transitions
  - v) choice
  - vi) reciprocity
  - vii) liberty
  - viii) refuge\*
  - ix) Embodiments
    - 1) rich to impoverished
  - x) Cues
    - 1) feedback
    - 2) feed-through

## 5) COMPUTERS AND PRIVACY

- a) Support Methods
  - i) computer security
  - ii) cryptography
  - iii) pseudonymity
  - iv) Access Control
    - 1) authentication
    - 2) authorisation
  - v) Content Control
    - 1) distortion filtration
    - 2) publication filtration
  - vi) Reliability
    - 1) data integrity
    - 2) process integrity
    - 3) stability
- b) Problems
  - i) inadvertent privacy infractions
  - ii) apprehension
  - iii) resentment
  - iv) the four 'D's : decontextualisation, disembodiment, dissociation, desituated action
  - v) role conflict
  - vi) Deliberate abuse
    - 1) misappropriation
    - 2) misuse
    - 3) identity theft
    - 4) impersonation
- c) User Interface Issues
  - i) degrees of temporal/spatial freedom for information access
  - ii) risk/reward disparity
  - iii) Feedback and Control
    - 1) believability
    - 2) socially natural qualities
    - 3) utility of privacy countermeasures
  - iv) Effort
    - 1) cognitive
    - 2) physical
  - v) Control Granularity
    - 1) fine- to coarse-grained

## 8.4 Analysis procedure

The four-step analysis procedure laid out below begins by assembling a description of the relationship between the various parts of privacy and the object of analysis. The object of analysis might be, for example:

— the inspection method used to identify and prioritise privacy risks in a system;

- the “on paper” design of a countermeasure to one of these risks;
- the prototype implementation of the countermeasure; or,
- the evaluation of the countermeasure in a semi-controlled laboratory setting, or in a field observational study.

The description can look at the “fit” between the object of analysis and all the various parts of privacy discussed in Act II. Such a description reveals what is and is not part of privacy insofar as the object of analysis is concerned. This kind of description is relevant to any kind of object of analysis. Concurrently, the description can look at the “effect” of the object of analysis on privacy regulation. Such a description reveals how the object of analysis supports or confounds privacy regulation. This kind of description is mostly relevant to the “on paper” design of countermeasures as well as their implementation in prototype systems.

Both styles of description draw heavily from the vocabulary given in Table 8.1. If an omission is noted in the description it illustrates an area for further iterative or exploratory design. Also, the act of preparing this description prepares the analyst for the deeper reflection of revealing hidden assumptions. What might be considered an assumption varies with the kind of object of analysis, for example:

- constraints on the kinds of risks found with an inspection method;
- conditions on the occasions in which a countermeasure will succeed; or,
- factors which confound the ecological validity of an evaluation of the countermeasure.

The analysis procedure concludes with an unstructured exploration of these hidden assumptions and a concise summary that circumscribes limits on the meaning of privacy as embodied in the analysis object.

#### 8.4.1 Step 1: Partition

In this step, the analyst roughly determines which of the three modalities of privacy regulation (solitude, confidentiality, autonomy) are the primary foci of the object of analysis. It is important to draw a distinction between main and secondary effects. This step will highlight areas for deeper analysis. By way of contrast, this step also highlights gross omissions. On the

one hand, these omissions represent a stopping condition on our analysis procedure: they are areas where deeper analysis might not be warranted. On the other hand, these omissions represent new areas for innovation and exploration. This step is not intended to require much thought or time.

#### 8.4.2 Step 2. Describe

For the areas highlighted in the previous step, proceed systematically through the descriptive theory vocabulary terms and their underlying concepts given in Table 8.1. Use those terms to describe the object of analysis subject in an unambiguous manner, focusing on the description of main effects. In assembling this description, a broad and deep picture of the stated connections between privacy and the object of analysis emerges. Subtle omissions or discrepancies between privacy as conceived in the descriptive theory versus privacy as embodied in the object of analysis highlight areas for future iteration on the design. This step requires time to complete, but because it unfolds in a systematic way with the vocabulary serving as a guide, there is no requirement that the analyst have particularly strong analytical skills to successfully complete this step.

#### 8.4.3 Step 3. Reveal

The goal of this step is to reveal *hidden* assumptions, not just highlight omissions. There is no easy recipe—no systematic procedure—for completing this step. The analyst must reflect upon the way the object of analysis accounts for different vocabulary terms. Depending on what is being analysed, the things the analyst must consider can be quite different.

— **If it is a prototype privacy safeguard design or implementation**, the analyst's goal is to answer the question: when will the safeguard succeed and when will it fail? To answer this question, the analyst must circumscribe conditions on the operation of the privacy safeguard by examining expectations regarding system use: who will or will not use it, in what contexts, and for what purposes.

- **If it is study to gather data or test hypotheses about privacy in the world**, the analyst's goal is to answer the question: what precisely is learned about privacy through the study and when is this knowledge valid? To answer this question, the analyst must pay attention to how variables are been defined, controlled, manipulated or observed, and how observations are or are not related together in the analysis procedure.
- **If it is a framework for identifying and analysing privacy risks**, the analyst's goal is to answer the question: which risks will be found and which countermeasures will be emphasised? To answer this question, the analyst must find the limits on what is considered part of privacy in the framework, what areas of the technology are up for consideration, and what kinds of use or misuse are evaluated.

Even though the description produced as the output of step 2 will richly inform this step, it nonetheless requires some insight, intuition, and skill on the part of the analyst.

#### 8.4.4 Step 4. Summarise

Synthesise the assumptions revealed in the previous step to generalise about when they hold and do not hold true. This step summarises the description and the assumptions to produce a sketch of the merits and demerits of the object of analysis and the conditions these merits and demerits exist. Although this step, like the previous, has little in the way of systematic guidance towards performing this kind of analysis, this step is considered easier to complete because it is simply bringing together and restarting the most salient points produced in the previous two steps.

### 8.5 Conclusion

In this chapter, I have developed a method by which I may systematically analyse and describe how privacy is affected by video media space design and use. This method will make apparent tacitly held assumptions regarding the nature of privacy and the kinds of privacy problems that will arise. The method is rooted strongly in the descriptive theory of privacy I presented in Act

II and directly utilises the vocabulary for unambiguous and holistic description of privacy that embodies the theory.

In the next two chapters, I present case studies which illustrate the application of this method. The first case study deals with the distortion filtration technique itself. It is done in greatest depth to best illustrate the technique. The next two case studies deal with the evaluations of the technique performed for Chapter 4 and the COLLABRARY toolkit presented in Chapter 3. I have selected these case studies because they allow me to revisit work presented in Act I with the understanding of privacy that is gained in Act II. This tells a compelling story about the potential of the descriptive theory of privacy to fundamentally change how one understands privacy.

In Chapter 10 two more case studies are presented. The first is on Neustaedter & Greenberg's (2003) HOME MEDIA SPACE prototype privacy-preserving video media space. The second is Hong's (2004) *privacy risk models* framework for analysing privacy problems. I have selected these case studies to show how the analysis method I have developed here can inform, augment, and be augmented by others' work in this field.

## Chapter 9—Case studies (1)

In this chapter, I will present case studies that illustrate the application of the analysis method presented in the previous chapter. Three different but related objects of analyses are considered:

- the distortion filtration technique (Section 9.1);
- the evaluation the technique given in Chapter 4 (Section 9.2); and,
- the COLLABRARY toolkit presented in Chapter 3 (Section 9.3).

The first case study—on the technique itself—is done in great detail so that it can serve as a good initial example of the technique in practice. As vocabulary terms are used in the analysis description, they are set in boldface type and the terms' number as given in Table 8.1 appears in parentheses. The focus of this chapter is to fully demonstrate that the descriptive theory of privacy and design developed in Act II informs the design of privacy-preserving video media spaces by revealing hidden assumptions that limit the efficacy of privacy safeguards, like distortion filtration. The actual assumptions revealed in the analyses themselves are considered to be of secondary importance.

These first three case studies deal with objects of analysis that can be applied to many scenarios of use. The analysis performed here is at a generalised level that is independent of the social context of use. Yet there is another level of analysis which does consider the social context in which the object of analysis is used. For example, in the first case study on the distortion filtration technique I do not describe the privacy needs or customs of the specific (if hypothetical) individuals using a media space that incorporates the technique. This type of

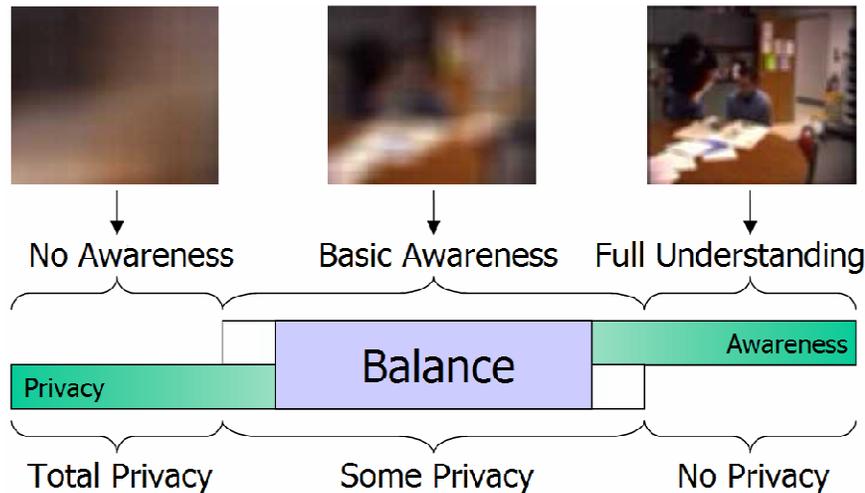


Figure 9.1—The first case study examines distortion filtration (like the blur filter) being used for this idea that privacy and awareness are opposing design goals that might be balanced. (Figure repeated from Chapter 4.)

analysis ought to yield far richer observations and should perhaps be recommended, but the assumptions and omissions revealed may depend on the peculiarities of the social context assumed in analysis. Since my goal here is to illustrate the analysis framework I have developed in a way that shows its broad applicability as well as its power to reveal hidden assumptions and omissions, I will abridge description of the social context

## 9.1 Case study #1: The distortion filtration technique

In this case study, I analyse the distortion filtration technique, as illustrated in Figure 9.1. Filters, like the blur filter pictured, remove some of the visual information, but not all of it. Usually, only high-precision visual information removed especially in the blur technique.

### 9.1.1 Step 1. Partition

The distortion filtration technique directly affects **confidentiality (2)**. Although the highly synergistic relationship among the different modalities of privacy control implies that the distortion filtration technique will

- 1) SOLITUDE
- 2) CONFIDENTIALITY
- 3) AUTONOMY
- 4) MECHANICS OF PRIVACY
- 5) COMPUTERS AND PRIVACY

Figure 9.2 Vocabulary for the Partition analysis step.

also affect **solitude (1)** and **autonomy (3)**, these effects are incidental consequences of the regulation of confidentiality.

### 9.1.2 Step 2. Describe

#### 9.1.2.1 Main effect: Confidentiality

Because the distortion filtration technique is largely focused on supporting confidentiality, I will first examine that modality of control in greater depth. Distortion filtration operates solely on the **visual information channel**

**(2.a.i.2)**. The information is obtained directly from **sampling (2.a.ii.1)** the visual field, rather than being interpolated, aggregated, or inferred from multiple other context sources. The distorted video image contains some personally identifying information, namely people's faces, but mostly contains information that is the kind that in the vocabulary is termed **information about the self (2.a.iii.1)**: the actions and activities, whereabouts, and encounters of a person that may or may not be known to or identified by an observer.

The distortion filtration technique seeks to modulate the **fidelity (2.b.ii)** specifically **precision (2.b.ii.1)** of the visual information to be communicated. The hope is that by decreasing the precision the **sensitivity (2.b.i.1)** of the information can also be decreased. At the same time, the technique also decreases the **certainty (2.b.iii)** that observers feel in their understanding of what they see, and that this creates opportunities for phenomena like **plausible deniability (2.b.iii.1)** and **misinformation (2.b.ii.3)** to occur. The technique itself is a kind of **edit (2.c.i.3)** operation that occurs after capture. While the technique is futile against misappropriation and misuse, it can make **scrutiny (2.c.iii)** and analysis a lot harder and surreptitious surveillance a lot less informative.

#### 9.1.2.2 Secondary effects: Solitude and autonomy

As mentioned before, the distortion filtration technique has secondary effects on solitude and autonomy. In terms of solitude, the technique aims to disclose some solitude-related

- 1) **SOLITUDE**
- 2) **CONFIDENTIALITY**
  - a) Information Channels
    - i) Medium
    - ii) Processing
    - iii) Topic
  - b) Information Characteristics
    - i) Basic Characteristics
    - ii) Fidelity
    - iii) Certainty
  - c) Information Operations
    - i) Basic Operations
    - ii) Use
    - iii) Scrutiny

Figure 9.3 Vocabulary to describe confidentiality.

information notably **high-level awareness (1.c.i)** cues such as **availability (1.c.i.1)** while concealing some autonomy-related information. It is conceived to be used in situations when **back-stage performances (3.a.ii)** mistakenly occur on the **front-stage (3.a.i)**, and so it seeks to regulate **appearances (3.a.i.3)** and **impressions (3.a.i.4)** while concealing **flaws (3.a.ii.1)** and **deviance (3.a.ii.2)**.

The technique works on every pixel in the video image, and so it also modulates the fidelity at which visual **signifiers (3.a.ii)** like **props (3.a.ii.2)** and **costumes (3.a.ii.3)** are perceived. The technique is mostly a countermeasure against **aesthetic harms (3.a.iv.1)**, although one could imagine scenarios in which **strategic harms (3.a.iv.2)** are also prevented. These countermeasures work by making the video image more **socially acceptable (3.b.ii.3)**.

### 9.1.2.3 Mechanics of privacy

As is apparent from the above discussion of the effects of the technique on confidentiality, plus also autonomy and solitude, the technique can be employed in the regulation of the **disclosure boundary (4.a.i)**. By itself, the technique lacks all of the characteristics of the **privacy process (4.b)**, however a context-aware user interface which incorporates the technique could conceivably support the **dynamic (4.b.ii)**, **dialectic (4.b.i)** and **cooperative (4.b.iv)** characteristics of the privacy process. The technique is to be employed to offer support for **self-**

- 1) **SOLITUDE**
  - a) Physical Dimensions
    - i) Interpersonal Distance
    - ii) Attention
  - b) Psychological Dimensions
    - i) Interaction to Withdrawal
    - ii) Escape
  - c) Presentation Dimensions
    - i) High-level Awareness
    - ii) Distraction
- 2) **CONFIDENTIALITY**
- 3) **AUTONOMY**
  - a) Social Constructions of the Self
    - i) Front
    - ii) Back
    - iii) Signifiers\*
    - iv) Harms
  - b) Social Environment
    - i) Social relationships
    - ii) Norms

Figure 9.4 Vocabulary to describe solitude and autonomy.

- 1) **SOLITUDE**
- 2) **CONFIDENTIALITY**
- 3) **AUTONOMY**
- 4) **MECHANICS OF PRIVACY**
  - a) Boundaries
  - b) Process Characteristics
  - c) Violations
  - d) Behavioural and Cognitive Phenomena
  - e) Environmental Support

Figure 9.5 Vocabulary to describe mechanics of privacy.

**appropriation (4.d.i)**. Long-term use of the technique will institutionalise its application.

We can already see the technique becoming part of a **genre of disclosure (4.d.ii)**: for example, the faces of youths charged with crimes are normally distorted when aired on television. The technique reduces the **risk (4.c.i)** of **violations (4.c)** by decreasing **probability (4.c.iii)** and **severity (4.c.iv)** of violations: also, when applied to an extreme level of distortion, the technique can eliminate the **possibility (4.c.ii)** of violations, too. Very little of the distortion filtration technique is accords with the kinds of support for privacy that the environment outside the computer offers. The technique supports **choice (4.e.v)** because it can be applied at a variety of different levels. Moreover, the technique supports **liberty (4.e.vii)** by providing some **refuge (4.e.viii)** when applied to an extreme level.

- 1) **SOLITUDE**
- 2) **CONFIDENTIALITY**
- 3) **AUTONOMY**
- 4) **MECHANICS OF PRIVACY**
- 5) **COMPUTERS AND PRIVACY**
  - a) Support Methods
  - b) Problems
  - c) User Interface Issues

Figure 9.6 Vocabulary to describe computers and privacy.

#### 9.1.2.4 Computer and privacy

The distortion filtration technique is a **content control method (5.a.v)** for **computer support (5.a)** for confidentiality. It is to be used as a countermeasure mostly against **inadvertent privacy infractions (5.b.i)** and **apprehension (5.b.ii)**. Some distortion filtration algorithms are one-way transforms, meaning the original video image can never be recovered. When these techniques are applied at an extreme level of distortion, the technique may also thwart **deliberate abuses (5.b.vi)** by rendering the visual information impotent in **misuse (5.b.vi.2)** or **misappropriation (5.b.vi.1)**. These effects, however, are contingent upon how well the user interface which incorporates this technique addresses key **user interface issues (5.c)** regarding **control (5.c.ii)** over the level of distortion used, such as **utility (5.c.ii.3)**, **effort (5.c.iv)**, and **granularity (5.c.v)** of such control.

#### 9.1.2.5 Omissions: future work

A number of obvious omissions have become visible in assembling this description. These omissions indicate opportunities for future work. Within confidentiality, the distortion filtration technique could be broadened to cover additional information media and topics. For

example, future iterations on the technique could muffle human speech, or show social networks inferred from email correspondence as fuzzy vertices and edges in a directed graph, or apply a blur-like filter to printed text. Also, there is now an apparent need for methods for manipulating accuracy instead of precision, such as context-aware algorithms for generating disinformation. Then, these methods need to be used within an architecture that makes fidelity gradually decrease with time for persistent data or with each subsequent re-transmission for transitive data, or be a function of the age and magnitude of social interconnections between actor and audience.

Beyond confidentiality, there are several critical omissions regarding solitude and autonomy. Little is known about the distracting nature of distorted images or the effect of fidelity on perceived interpersonal distance and its consequences on intimacy in social relationships. There is now an apparent need for model-based distortion filtration methods: those that model a scene in terms of actors, props, costumes and so forth, and selectively blur specific elements rather than uniformly. Since the distortion filtration technique is intended to conceal deviance, we can consider the application of machine learning algorithms whereby the computer can infer from example what constitutes normal behaviour and automatically apply the distortion filtration effect when its artificial intelligence detects deviance.

### 9.1.3 Step 3. Reveal

In Chapter 8 I claimed that there is no systematic method for revealing hidden assumptions in this step. I gave a concrete suggestion on how to proceed, however: reflect on how different vocabulary terms are accounted for in the preceding description. Since the object of analysis is a privacy safeguard, the goal of this step is to circumscribe limits on when distortion filtration will succeed and when it will fail. To begin, I will list below the major vocabulary terms that come out of the description for each of the five different sections of the vocabulary. I have grouped them to show strong relationships among the words.

— Group 1: high level awareness (solitude)

— Group 2: back-stage performances, signifiers, aesthetic harms (autonomy); inadvertent privacy infractions (computers and privacy); disclosure boundary (mechanics of privacy)

— Group 3: precision, sensitivity (confidentiality); severity (mechanics of privacy); social acceptability (autonomy)

The terms in Group 1 describe an important reward of video in a media space. Video media spaces provide high level awareness cues during periods of loosely coupled work so as to facilitate the efficient microcoordination of periods of tightly coupled work. The terms in Group 2 describe an important reward of distortion filtration in a video media space. The distortion filtration technique conceals irrelevant back-stage performances during loosely coupled work. These two rewards are set up as opposing forces in a disclosure boundary tension. Ideally, the distortion filtration technique should be applied just enough so that inadvertent aesthetic privacy violations are prevented, yet the rewards of video may still be enjoyed. The terms in Group 3 describe the aspects of the back-stage performances that are manipulated by the technique so as to resolve this disclosure boundary tension. In this step of the analysis, I will try to find logical conditional relationships on these aspects that must be satisfied in order for the distortion filtration technique to work.

#### 9.1.3.1 Precision and sensitivity

Distortion filtration filters out high-precision information specifically because it is assumed that the high-precision details of an image contain the sensitive aspects of the scene. This is a direct relationship: as precision increases, sensitivity is also assumed to increase. Correspondingly, it is assumed that the severity of the aesthetic harms that would result from unintended disclosure of these details also increases with the precision.

Based on this initial assumption—that the details are what is risky in an image—low-precision overviews, like the kind generated by the distortion filtration technique, are subsequently assumed to reduce the both the probability of a privacy violation and the severity of harms arising from unintended disclosure. Furthermore, it is assumed that the low-precision overview is useful for balancing disclosure boundary tensions by concealing sensitive information while disclosing low-sensitivity information.

This assumption breaks down when there is still sensitive information in even the lower-precision distorted view. For example, Figure 9.6 (a) and (b) show two different visual scenes filtered with a blur filter at the same level. While (a) looks harmless, (b) looks risky. (Try

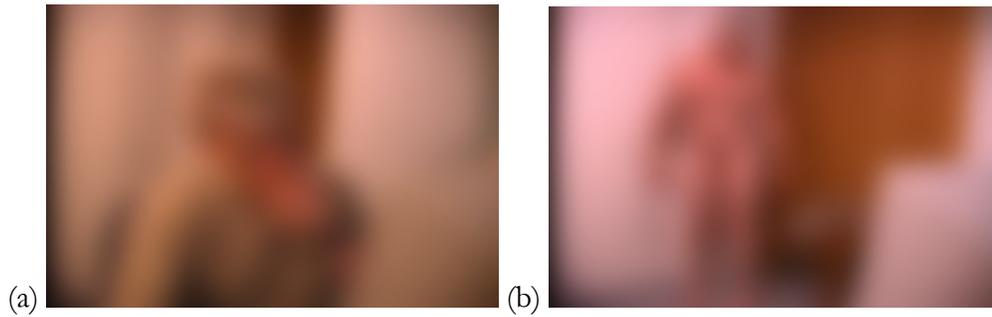


Figure 9.7 The distortion filtration technique does not preserve privacy when the sensitive aspects of a scene are of low precision.

viewing these images in colour or from a distance.) There is some low precision visual information in (b) such as the shading around the groin area and the body hue (if shown in colour) that does not conform to the assumptions made about precision and sensitivity, and hence the distortion filtration technique does not provide very good support for confidentiality in this case.

#### 9.1.3.2 Precision, sensitivity, and informal awareness

To enjoy the reward of increased informal awareness with a video media space that incorporates distortion filtration, it is necessary for the video to reveal enough informal awareness cues to support microcoordination of tightly coupled interactions *both before and after filtration*. Video captured under poor conditions may not provide enough awareness itself, even if it was unfiltered, or what little awareness information there is might be removed by the filter.

Moreover, in order for distortion filtration to work, there has to be visual information in a video field that both useful for maintaining informal awareness and is not sensitive. Sometimes, such as when a person wishes to be alone, even minimal or momentary informal awareness cues like presence and location can become very sensitive. The point of the distortion filtration technique is to conceal all but the most essential informal awareness cues, but if these cues are sensitive, then the technique does little to actually preserve privacy. For example, imagine one colleague has disinformed another over the phone that she has already left to meet him at his office with the intention of not leaving for yet a while. Even if she turns the camera towards the wall and applies distortion to an extreme level, incidental visual traces of her presence, such as office lights being turned off or on might contradict the

disinformation she gave earlier. Such an incidental disclosure resulting in an inadvertent privacy violation with aesthetic privacy harms could not be prevented by the distortion filtration technique because the sensitive information (e.g., global illumination changes resulting from lights turning off or on) is already very low-precision.

#### 9.1.3.3 Misinformation and social acceptability

It is understood that sometimes high-precision contextual details need to be invented—imagined—by viewers so that they can make sense of a low-precision visual scene. In the vocabulary given in Act II, these imagined details are called **misinformation (2.b.ii.3)**. In order for distortion filtration to preserve confidentiality, the sensitivity of the misinformation and the severity of the harms that arise because it is generated must necessarily be less than that of the real information that is concealed by the filter.

Thus, the distortion filtration technique makes an implicit assumption that this imagined misinformation favours social acceptability. However it is easy to consider situations in which the misinformation is riskier than the truth. Furthermore, the distorted image—the fact that something unknown is being concealed—can titillate and activate voyeurism and curiosity. This may draw greater attention and scrutiny to the actor and his performance, affecting solitude and autonomy, as well as confidentiality.

#### 9.1.3.4 Signifiers, precision, and conformance

A media space may join an actor in one social setting with an audience in another social setting. Although these social settings are usually very similar (e.g., both may be personal offices) they are nonetheless separate. Expectations regarding the social acceptability of performances may vary between settings. Furthermore, changes in one setting will not be reflected automatically in the other. The distortion filtration technique attempts to relieve the problem that result from such decontextualisation by reducing a performance to a low-precision form that is likely to be considered acceptable in the audience's setting.

In order for this to be true, it is assumed that socially acceptable performances in the various social settings in a media space differ only in high-precision visual information, since these are the visual differences that are filtered by the technique. When this assumption does

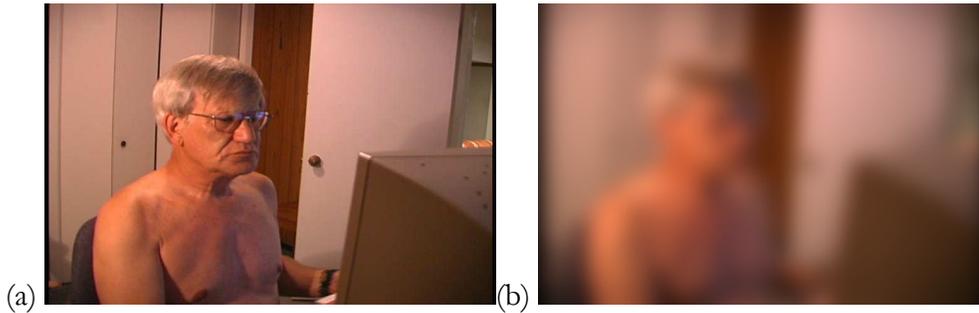


Figure 9.8 Some performances are social acceptable in the actor's setting, but not the audience's setting and even heavy filtration does not change this.

not hold, though, distortion filtration cannot adequately safeguard against inadvertent autonomy problems. Figure 9.7 (a) and (b) shows an actor engaged in a routine activity in a home-office social setting. Even though the image in (b) is heavily filtered it is still largely inappropriate for viewing by colleagues in a corporate office.

#### 9.1.4 Step 4. Summarise

Reflecting on the description of the distortion filtration technique developed in Step 2 and the assumptions discussed in Step 3, in this step I seek to summarise the merits and demerits of the technique and the circumstances in which they arise.

The essential merit of the technique is that it conceals only high-precision information. This ought to make the distortion filter technique useful for regulating the disclosure boundary by supporting the efficient microcoordination of tightly-coupled interactions while reducing the probability and the severity of the aesthetic privacy harms that arise from inappropriate disclosure. This merit applies when four conditions are met:

- ample informal awareness cues are encoded in unfiltered low precision information;
- the sensitive information in a visual scene is encoded in only high precision information;
- the differences between a socially acceptable performance in the actor's setting versus the audience's setting(s) lay only in high precision information; and
- any misinformation (e.g., details imagined to help make sense of the low-precision visuals) casts the actor favourably.

Perhaps one of the reasons why the distortion filtration technique is so popular is that these conditions are typically—but not always—met in common VMS use cases.

There is a second merit attainable with the distortion filtration technique. It can establish a place of refuge in the media space but when it is applied to an extreme level of distortion. Ideally, no meaningful awareness information should be discernable from such “zero fidelity” views. It is this total deprivation of outside access to the self that affords the refuge.

Two critical omissions have been noted in this case study. First, consider the converse of the second condition given above. Sometimes innocuous yet useful information is encoded in the low-precision visual information that gets filtered out by the technique. Thus, when the technique is applied, it conceals not only deviance and flaws, but also conformance, informal awareness cues, and the context for spontaneous and serendipitous casual interactions. Because of this, even when all four conditions listed above are met, the success of the technique still hinges largely on the user interface by which its application is controlled to ensure that useful low-precision information is not inappropriately filtered. Second, there is little understood about the broader, long-term effects of distortion filters on solitude and autonomy. For example, we do not yet understand the harm to relationships that might arise from a lack of intimacy or an erosion of trust brought on by distortion filtration. Consequently, we do not have a very well-formed picture of the demerits of the technique.

### 9.1.5 Conclusion to case study #1

The analysis performed in this first case study exemplifies the method developed in Chapter 8. This method is strongly grounded in the descriptive theory of privacy developed in Act II. In this case study, the vocabulary obtained out of the descriptive theory has been used to reveal and articulate a considerable number of assumptions hidden in the very concept behind the distortion filtration technique.

- The vocabulary and method permitted description of the expected benefits and risks of the distortion filtration technique for regulating privacy and informal interactions.
- The vocabulary pointed the way towards explanation of the mechanism by which this benefit is accrued.

- The vocabulary-based method used in this case study permitted the explication of significant caveats (in the form of assumptions) to these mechanisms, risks, and benefits.
- Reflection over the assumptions revealed permitted prediction of the success or failure of the technique in a given use case
- The assumptions and summarised reflections permitted the generation of use cases to test for the success or failure of the technique (failure boundary conditions).
- The method permitted the identification of critical privacy-related omissions, in the form unaffected aspects of solitude, confidentiality, autonomy, or environmental mechanics that support these modalities of control.
- Ideas for iterating over a design that incorporates the distortion filtration technique were generated as a by-product of the analysis procedure.
- New directions in supporting privacy regulation in video media spaces were revealed through the analysis procedure.

## 9.2 Case study #2: Evaluations of the distortion filtration technique

In Chapter 4, I presented the results of two studies evaluated the distortion filtration technique. I revisit that work in this case study.

In the first study, mundane low-risk office scenes were distorted using the blur and pixelize filters at ten different levels. A scene was shown to a study participant at the lowest level of fidelity first and gradually increased as the participant answered questions to determine the awareness cues that were visible. When all ten levels had been viewed, the participant rated the privacy-preserving potential of each level. A threshold detection analysis procedure was used to identify filter levels at which basic and detailed awareness cues were perceptible and levels at which privacy was reasonably preserved.

In the second study, high-risk home office scenes were distorted using the blur filter at similar levels as in the first study. Scenes that included partial or full nudity were included to address a methodological weakness in the first study (that the filters were not tested in privacy-sensitive situations). A similar experimental design was used, except: (1) at each level

participants in the second study were also asked to rate the privacy threat to the actor in the scene and the actor's family members; (2) when all levels for a scene had been viewed, participants were asked to choose a level that made the scene acceptable for viewing by office colleagues; and, (3) when all scenes has been viewed, participants were asked to perform a forced sort of the scenes according to privacy risk. The second study analysis also identified ranges in which basic and detailed awareness cues were visible, and where privacy was preserved.

The first study found that a balance between privacy and awareness was possible with the blur filter around levels 3~5 for mundane, low-risk scenarios. The second study found that such a balance is not possible with the blur filter when used on high-risk scenarios. Taken together, the studies suggest that there is no general purpose range of filtration levels that adequately balance privacy and awareness design goals in all situations.

### 9.2.1 Step 1. Partition

Much like the distortion technique itself, the first and second evaluation studies of the distortion filtration technique focus on **confidentiality** more so than **autonomy** or **solitude**. Consequently, much of the analysis of the technique naturally also applies to its evaluations. However, the evaluations expand upon the treatment of privacy given in the technique in an important regard. Because both studies make use of semi-controlled experimental designs (instead of field observations) the study protocols necessarily make explicit assumptions about contextual factors surrounding the anticipated uses of the technique; in the jargon of scientific experiments, these kinds of assumptions are called "control variables." As a result, we can be a bit more specific than in the previous case study and "close" some previously "open-ended" points of discussion.

### 9.2.2 Step 2. Describe

**Solitude** is given the same limited treatment in the evaluations as it is given by the technique itself. The only aspects of solitude tracked are an observer's ability to track some visual **informal awareness cues**, but these measures relate more to confidentiality than they do solitude. In addition, the use of non-interactive pre-recorded video scenes in the study

protocols precludes the tracking of effects on psychological dimensions of solitude (e.g., **intimacy**) or its presentation (e.g., **distraction**).

Since the evaluations focus primarily on the **confidentiality** of **visual information about the self**, much of the confidentiality analysis done for the technique itself also (unsurprisingly) applies to the evaluations of it. How the technique affects confidentiality in terms of the **channels**, **properties**, and **operations** permitted remains the same as per the previous case study. An important distinction that arises in this case study is that since both evaluations utilised (different) stock palettes of video scenes some characteristics of the visual information tested like the **topics** of information disclosed and their **sensitivity** become independent variables. A stated caveat of the first study's results relates to the fact that it used mundane video scenes; the second study used more variation in the topics and sensitivity of the scenes presented. Other characteristics, such as **persistence**, and **transitivity** become control variables through careful wording of the experimental protocol given to participants.

The level of filtration is treated as an independent variable in both studies and correspondingly the **precision** and **accuracy** (as judged by the experimenters, not the participants) at which visual information is perceived and participants' self-reported **certainty** in the accuracy of their observations were treated as dependent variables. In this way, the studies sought to verify a hypothesis that the filters do in fact reduce the fidelity of visual information that can be accessed from video. While the first study also varies the distortion algorithm used as an independent variable, the analysis procedure was the same for both and it is as though the analysis procedure had merely been repeated for the second algorithm. The first study aggregated precision, accuracy and certainty into a single estimate of "understanding" with three levels (none, basic, and detailed) in its analysis while the second study keeps certainty separate and also utilises binary values for the accuracy metric in its analysis.

For the most part, the two evaluations handle **autonomy** differently. They are similar to each other in that in both studies, the filters' effects on **identity** (i.e., **personally identifying visual information**) could not be tracked because unfamiliar actors were featured in the video scenes used. Although the first study asks participants to determine the gender of the actors

depicted, this measure is lumped in with other awareness related measures in the analysis. In fact, the first study is largely unconcerned with any aspect of autonomy.

The second study remedies this by considering the distortion filtration technique's ability to make some **costumes** and **actions** appear more **socially acceptable** (when filtered versus unfiltered) as its main reward. The second study's third hypothesis seeks to take the loss of precision incurred through distortion and relate it to an increase in the **conformance** to (or decrease in **deviance** from) socially held **norms** and **expectations** for **self-appropriation**. The study protocol specifically asks participants to choose the least distorted level of filtration that still makes the video scene viewed "socially acceptable." In this way, the study is highly concerned with autonomy insofar as **social constructions of the actors' selves** and the risk of **aesthetic harms** to the actors are concerned.

As mentioned earlier, however, the use of semi-controlled experimental methods gives fixed values to the **social environment** that predicates participants' understanding of the prevailing norms for socially acceptable performances. In both studies a carefully worded verbal protocol given to participants fixes specifically the imagined **place** in which they as the audience find themselves viewing the actor's performance, and it also fixes the **roles** and **status relationships** between themselves and the actor viewed. The place is also reinforced by physically seating the participant in a room and environment that are architecturally and socially similar to that to be imagined. The choice of study materials seeks to create a conflict between the norms inferred from these fixed contextual factors and the actors' visible costumes, actions and manners.

Since both evaluations try to obtain measures of privacy, it is useful to describe what is measured, and by omission, what is not measured because these inclusions and omissions relate to the mechanics of privacy. The first study does not deconstruct privacy in any way: it treats it as a scalar value that is both subjective (i.e., differs by the person experiencing it) and relative (i.e., filter level  $A$  affords  $X$  units more/less privacy protection than filter level  $B$ ). It simply asks study participants to "rate the privacy afforded by each level of the filter." There is simply no way to tell if "privacy" as held in the minds of participants as they answered the study questions might have dwelled largely on **visual confidentiality**, even though immediately prior to giving these privacy assessments the first study's protocol asked targeted

questions to gauge the visibility of elements in a scene's visual composition. The second study on the other hand very clearly deconstructs privacy as the risk of **aesthetic harms** arising from an untimely **loss of visual confidentiality** in which a normal **back-stage performance** occurs on the front. Although risk is not further deconstructed into **probability** and **severity**, it is likely that the "line of privacy" methodology used in the study measures the relative severity between different video scenes.

As with the technique itself, the studies are concerned with the regulation of the **disclosure boundary**. In the analysis performed on the first study's result, ideal ranges of filtration—those that permit basic understanding while being perceived as preserving privacy—are found. These ideal ranges could then be incorporated into a **genre of disclosure** for loosely coupled work that seeks to resolve the **risk/reward trade-off** by concealing mistakes in self-appropriation. The use of a single filtration level as a default when not engaged in tightly-coupled interpersonal interactions, however, is not a very **dynamic** means of regulating confidentiality. Indeed, only narrow ranges were identified the first study, suggesting that if the technique is to be used it must be incorporated into a complete system that also supports other characteristics of the privacy process, such as a **cooperative dialectic**, or other forms of environmental support, such as **rich embodiments** that permit understanding needed to support **situated action**, or introduce latencies in dramatic changes to the context in order to smooth out sharp privacy transitions.

While the distortion filtration technique attempts to safeguard against **inadvertent privacy infractions** that arise from the **decontextualisation**, **disembodiment**, and **dissociation** of action inherent to video media spaces and offset users' and non-users' **apprehension** towards the technology neither study tracks apprehension. Moreover, since a semi-control lab experiment design was used, neither study is able to collect observations on inadvertent privacy infractions that were or were not ameliorated that could be obtained in, say, a field observation study.

### 9.2.3 Reveal

Recall from the previous case study that the distortion filtration technique safeguards against aesthetic harms that arise when back stage performances unintentionally occur on the front

stage. In the fourth step of the previous case study, a list was generated giving the conditions in which the distortion filtration technique will be useful for balancing the disclosure boundary tension. We can use this list to analyse the stock video footage used in the initial and follow-up studies, to see if the footage used covers these conditions.

The first two conditions derived in previous case study concern the source of the risks that media space participants are exposed to. The first condition is that if there are differences between a **socially acceptable performance** in the audience's **setting** and the actor's setting then these differences are encoded in **high-precision** information. In this way, the audience is given an inappropriate glimpse into a back-stage performance that could result in aesthetic harms. At the very least, in order to adequately evaluate the distortion filtration technique with stock footage, some of this footage must be of actors in social settings similar to the audience and some of the footage must be of actors in a difference social setting from the audience.

The first study featured mundane video scenes of typical looking office workers in typical looking office environments engaging in typical looking office work. All of the videos used show actors in social settings that are very similar to the one constructed for study participants as audience. Thus, even though the actors are featured by themselves, the audience is not given a glimpse into a back stage performance. No aesthetic harms could arise from any observation of these performances.

The second study tried to remedy the shortcoming of the first study by using scenes of people in a typical home office environment doing things that are typically done in the home although not always in a home office environment and are certainly atypical for a corporate office environment, such as putting on clothes after bathing. The scenes were selected to create an obvious disjunction between the actor's social setting and the audience's social setting, or rather between what is considered normal in the actors context and what is considered normal in the audience's context. Undoubtedly, participants as audience were given a glimpse into a back stage performance. Moreover, the study looks at differences in both high-precision details (such as the working-with-no-shirt scene) and in overview (such as the changing-clothes scene).

Despite this seemingly adequate coverage, important omissions have been revealed in our description, and these omissions merit further elaboration. Although the second study does present participants with back stage performances, these performances do not reveal **flaws** or **deviance** that, if known, could be imagined to have serious repercussions for the actor portrayed. Even the “riskiest” scene—changing clothes—shows the actor in a room in his or her own home, costumed in a manner and engaged in activities that are socially acceptable in such a setting when one is by oneself. The few **props** visible are mundane and carry little confidential information. **Strategic harms** are also not tested by any of the video scenes used. Lastly, since there is no a-priori relationship between actor and audience in the semi-controlled laboratory study, it is not possible to measure the effects of lapses in careful regulation of visual confidentiality on the maintenance of **idealizations** that stabilise social relationships.

These omissions point to a hidden assumption that concerns the expected genre of disclosure to be used with the video media space, and the kinds of risks to be minimised. Since the first study does not really utilise videos that contain sensitive visual information, this discussion focuses on the second study.

First, there is an assumption that the **privacy risks** are greater for a person at home than for a person at work. Second, there is an assumption that working at home in the presence of a video media space connecting to office colleagues is risky because a person at home routinely appears with costumes or actions that while appropriate for a home setting when the actor is alone are nonetheless inappropriate for viewing by a social peer in an office setting. Thus, while the authors of the second study report state that “participants usually associated threat with the person’s particular **activity** or **appearance**” it is important to note that the scenes presented were engineered thusly. Along similar lines, when the authors report that “...the Kissing scene posed the highest risk to family members” we must note that it was the only scene tested to actually feature a visible family member. Finally, there is an assumption that making video appropriate for anyone to view suffices to eliminate the **privacy threat**. This assumption holds if and only if the privacy threat stems exclusively from the fact that the video is inappropriate to view. Moreover, this assumption excludes consideration of modalities of **repair** and **recourse** that may be available.

Finally, both studies use a repeatable experimental method in which a participant is shown a video scene and then asked to think about the privacy of the actor shown. In order for such a method to be sound, it is necessary to assume that people can cogitate about privacy in desituated contexts. For example, the evaluations assume that a person can reliably imagine themselves as the actor seen and from this can reliably determine the privacy risks faced by the actor and their relative severities. If unfamiliar actors and settings have an effect on such privacy assessments, then it is assumed that the effect will be to decrease the upper bound on the “safe” range of fidelities.

#### 9.2.4 Summarise

Both studies claim to evaluate the distortion filtration technique’s effects on awareness and privacy, and in the analysis for this case study we have narrowed in on exactly what effects are tracked.

Neither study tracks “awareness” per se: instead, they track the visibility of visual information that is expected to be useful for tabulating informal awareness. Because no interaction occurs between a study participant and the actors featured in the studies’ stock video footage, these measures relate more to **confidentiality** than they do solitude. To more precisely paraphrase one consistent result of the studies, the distortion filtration technique affects the precision at which visual information can be accurately and confidently discerned from a video scene. This effect could be used as part of a multifaceted approach to regulating confidentiality, but its usefulness in this capacity is not really evaluated.

Both studies try to track a filter’s effects on privacy, but there are notable omissions. In the first study, there is a lack of detailed deconstruction of the dimensions of privacy to track and there is limited diversity in the **sensitivity** of the performances featured in the video footage shown to participants. These methodological problems mean that results of this study are not very useful for evaluating the fitness of the distortion filtration technique as a privacy safeguard. In the second study, these methodological problems are partially addressed, but there remain omissions that constrain our ability to really generalise from the results of the study. The important conclusion from this case study’s analysis is that the distortion filtration technique should be viewed by designers not as a panacea but rather as but one tool for

regulating confidentiality and autonomy that could be widely applicable but only narrowly beneficial.

### 9.3 Case study #3: The COLLABRARY toolkit

In Chapter 3, I presented the COLLABRARY: a toolkit intended to support the rapid prototyping of video media spaces that incorporate novel approaches for preserving privacy. In this case study, I analyse the privacy-technology link as realised through the features in the COLLABRARY and their implementation. Doing so helps cast some light on what precisely which aspects of privacy can be preserved by media spaces built atop the COLLABRARY. As it turns out, the overwhelming majority of the concepts and characteristics of privacy as captured in the vocabulary are entirely absent from the COLLABRARY, calling in to serious question the privacy-specific value of the toolkit.

#### 9.3.1 Step 1: Partition

Computers deal with information, which is principally the domain of the **confidentiality** modality of privacy control. Accordingly, most of the features of the COLLABRARY that are intended to facilitate the implementation of novel privacy-preserving safeguards are focused largely on narrow aspects of confidentiality. Some of the information to be regulated, however, is either part of a person's **digital persona** or is **information about the self** that is regulated by the **autonomy** modality of control. Some of the information is conveyed to facilitate human **interpersonal interactions**, which are the domain of the **solitude** modality of control.

#### 9.3.2 Step 2: Describe

**Solitude:** The COLLABRARY supports regulation of solitude mainly by providing facilities which can be used to construct a video media space for **informal awareness** and **casual interaction**. Recall from the vocabulary that solitude regulation requires that interaction be supported as well as withdrawal. In general, interaction is clearly favoured over withdrawal in the COLLABRARY's feature set. Uncoupled work (withdrawal) has typically been the only thing supported well in software; the "hard work" performed by the COLLABRARY is making tightly coupled, highly interactive work possible.

To support **withdrawal**, the COLLABRARY handles unexpected device or network disconnection as a normal mode of operation and it is easy for end-users to construct a place of refuge in a video media space constructed with the COLLABRARY by simply “pulling the plug.” Also, some of the **psychological** and **physical dimensions** of solitude could be approximated using the COLLABRARY. For example, distortion filtration can be used to signal **intimacy** or **reserve**, and the size of video frames can be easily changed to correspond to changes in the **degree of coupling** between people collaborating over a distance and the **attention** available for interpersonal interactions.

Although the COLLABRARY provides the rudimentary channels needed to support informal awareness and casual interaction and the rudimentary facilities needed to support withdrawal from interactions, there is little else in the way of support for the regulation of solitude within the COLLABRARY itself. For example, **availability** is a critical informal awareness cue for regulating interactions, yet the COLLABRARY provides no “availability tracker” component to alleviate the end-programmer of the hard task of aggregating various signifiers of availability. Lacking also are Bayesian statistical tools to help sort out relevance or salience (Horvitz et al, 2003). A person’s **gaze**—a high-quality signifier for attention—cannot be conveniently tracked with the Collabrarity.Camera component. The Collabrarity.Speaker component does not make available any indication of the potential **distraction** caused by the sound it is being asked to generate. **Objective signifiers** for solitude such as the number of people are present in a room and which of them are looking at the video media space are also unimplemented.

**Confidentiality:** The COLLABRARY affects confidentiality because the components it provides support information **capture**, **aggregation** and **interchange** (but not **inference**) along **aural** and **visual**, **numeric**, and **textual channels**. Although the information that can be captured by the COLLABRARY components includes both **personally identifying information** and **information about the self** including **activities**, **whereabouts**, and **encounters**, these topics are not exposed to end-programmers in these terms. Instead, the COLLABRARY affords end-programmers the direct access to captured multimedia in the form of audio samples or video pixels.

Because of this mismatch between the “data” of privacy and the “data” of the COLLABRARY, some basic information characteristics such as **sensitivity** and **fidelity** of capture are not provided to the end-programmer. Along the audio channel, for example, the **intelligibility** of human speech present in recorded sound samples is not measured by any of the COLLABRARY components. The **topic** of conversation is not made available because the COLLABRARY offers no speech recognition components and, consequently, neither is the **sensitivity** of the topic of conversation made available. In the video channel, the COLLABRARY does not make known to end-programmers, for example, a map of which pixels in an image correspond to which elements of a **performance**, e.g., **actor**, **costume**, **prop**, or **scenery**.

Most of the COLLABRARY’s built-in support for visual confidentiality comes as a palette of **distortion filters**. As per the previous case study, these can give the end-programmer a convenient way to alter the **precision** of a video frame or block of audio samples. All of the algorithms implemented are model-free (i.e., assemble no understanding of the true topic of information conveyed) and so cannot be used to change the **medium of representation**. For example there is no way in the COLLABRARY to “render” a video frame as a verbal description with different sensitivity or fidelity characteristics, nor is there any way to alter the **accuracy** of the representation of a person in a video frame while keeping **precision** constant (which would be useful for supporting careful and deliberate **disinformation** as a confidentiality mechanism).

The COLLABRARY provides the `SharedDictionary` component so that end-programmers may, with little effort, prototype distributed systems that transmit or persist captured information. Yet, the shared dictionary is not very **accountable** (Dourish, 2001) because it provides no cues to end-users that would help them accrue and maintain awareness of what information is being transmitted to whom, when, and for what purpose. Moreover, when information is stored in a shared dictionary it remains at the same **fidelity** no matter how long it is kept there or no matter to how many people it is transmitted. There is little support for tracking the **transmission**, **use**, **misuse**, or **misappropriation** of information captured or transmitted by the COLLABRARY components.

The shared dictionary does not use any schemes for **authentication** or **authorisation** (i.e., **access control**). Thus, any person on the network may connect to a shared dictionary

and the COLLABRARY itself provides no end-user notification nor requires no **consent** for any operation it performs. Also, public-key cryptosystems are not (by default) used as part of the shared dictionary's wire protocol. These design traits combined make the COLLABRARY open to **surreptitious surveillance**. **Scrutinising** analysis of the information captured or carried by COLLABRARY components is made possible by the convenient direct access to multimedia data offered, but only extremely simplistic analyses are supported directly: for example, the `Photo.Subtract` method and `Photo.PSNR` properties can be used to quickly estimate the amount of motion between successive frames of a video sequence.

**Autonomy:** The COLLABRARY affects autonomy because it accommodates media that is part of one's **digital persona**; the selfsame media signals one's **identity** and provides the **embodiment** by which a person appears to others, interacts with them, and makes his or her **impressions** upon them. The COLLABRARY shared dictionary supports **pseudonymity** in that end-users' network addresses are known only to the centralised server the only uniquely identifying information used for client computers is a pseudo-random ID.

The COLLABRARY is designed to be used in distributed systems which connect people in markedly different **social settings**. As described in the vocabulary, such circumstances of use can lead to discontinuities in the social **norms**, **expectations**, and **protocols** that govern interpersonal interactions. These concepts, however, are omitted in the COLLABRARY. The COLLABRARY provides no algorithms to determine if a person looks **acceptable** to another, or to determine what kind of **impression** he or she is making, or to pick out **flaws** in one's appearance. While these rather subjective qualities are still challenging for digital computer algorithms to assess, more rudimentary and objective measures are not implemented in the COLLABRARY, either. For example, consider that a basic machine-learning system might be able to classify a performance (represented in video frames) as being **conformant** or **deviant**. Furthermore, the COLLABRARY does not provide routines to deconstruct a video frame into **signifiers** of the **social setting** such as **territory**, **props**, or **costumes**; hence, the COLLABRARY does not readily support the detection of meaningful changes to the social setting. For example, a simplistic approach that could have been implemented in the COLLABRARY might have been to ask the user to select the closest match from a list of canonical example social settings. At the very least, the end-programmer might have some

approximation (no matter how inaccurate) of the social setting to be incorporated into the user interface.

**Mechanics of Privacy:** The COLLABRARY directly affects every boundary across which privacy is regulated. As mentioned previously, regulation of the **disclosure** and **spatial boundaries** is complicated by the easy ability afforded by the COLLABRARY to capture and distribute multimedia recordings of performances in ways that are wholly unaccountable by end-users. Regulation of the **identity boundary** is made more complicated because of the fact that the COLLABRARY it is geared to support the development of distributed **presence** and interaction systems making problems like **role conflict** more likely.

Regulation of the **temporal boundary** is made more complicated because the COLLABRARY can be used to **record** and **archive** parts of performances for later viewing but much of the social setting that serves as the context for a performance is not captured at all, making problems like **decontextualisation** and **desituation** more likely. The `.transient` and `.lifetime` metadata functionality in the shared dictionary, described in Chapter 3, might be useful in regulating the temporal boundary by putting constraints on the persistence of data, but once the data is made available to program components outside the COLLABRARY these **persistence** rules are easily violated.

The design of the COLLABRARY does not take into account any of the characteristics of the **privacy process** and it is up to the end-programmer who designs the user interface that utilises the COLLABRARY to, on his or her own, support such things as **dialectism** and **cooperation**. In a similar manner, concepts of privacy related to **violations** such **risks** and **severity** are wholly ignored in the design of the COLLABRARY. All information is given equal treatment and it takes the same amount of effort to mediate a performance regardless of the **threat** involved.

The implementation of the COLLABRARY is not predicated upon any of the behavioural or cognitive phenomena used by people to regulate privacy. For example, the COLLABRARY offers no components that would help programmers offer useful **feedback** for **self-appropriation** such as, say, knowing when a particular person's computer has subscribed to a piece of information in the shared dictionary, indicating that this person could now observe a

remote performance. The COLLABRARY is not adept at handling multiple concurrent **genres of disclosure** and it requires significant end-programmer discipline to support behaviours like **reserve**. Since the components in the COLLABRARY are entirely unaware of violations (when they will happen or if they have happened and how severe the harms resulting from them will be or were) end-programmers are left to their own devices to implement schemes that utilise **feed-through**, **feedback**, and **inspection** to police the media spaces they devise. The authentication-less architecture used in the shared dictionary further complicates **policing**, and offers few **reprimands**: only the server can disconnect a client computer already connected to the shared dictionary and, still, there is no way to prevent them from immediately reconnecting. It is clear that the COLLABRARY is designed with the same assumptions about risk and reward as were revealed in the analysis of the distortion filtration technique; as a result, we find that beyond the distortion filtration algorithms there is no particular support for managing the **risk/reward trade-off** or the **disclosure boundary tension**.

The COLLABRARY ideally supports 2D video and audio **embodiments** for interactivity. While on the one hand these seem particularly rich, on the other hand the COLLABRARY is bandwidth constrained and so only small video frames may be used at full-motion frame rates. The COLLABRARY lacks much specific support for **environmental affordances** for privacy and it is up to end-programmers to implement much of this in the user interfaces they design. The COLLABRARY makes programming easier, which should encourage and facilitate end-programmers' development of many options for regulating privacy each with different merits, demerits, and circumstances under which they apply. In doing so, the COLLABRARY supports the basic environmental affordance of **choice** for privacy regulation, but it is still up to end-programmers to think up and implement these choices. Some features, like the animation of transitions to **smoothen** important changes in the social setting for a performance, or providing mirrored feedback to oneself to support **reflexive interpretability of situated action**, are left for end-programmers to implement.

As mentioned earlier, for a toolkit that is intended to support the prototyping of privacy-preserving video media spaces, the COLLABRARY has remarkably poor **security** and does nothing to make conventional security **safeguards** easier to rapidly integrate into a prototype. For example, although there are **cryptographic** components, they are awkward to use and

aren't used by default within, say, the shared dictionary. There is little in the way of **access control** in the shared dictionary because it does not utilise any **authentication** scheme. There is little in the way of **content control** beyond the provided **distortion filtration** techniques and aside from these there's nothing to safeguard against **inadvertent privacy infractions**. Moreover, some forms of unchecked **deliberate abuses** (e.g., **misappropriation**, **misuse**, and **impersonation**) are made possible by certain COLLABRARY design choices that de-emphasise security and **reliability**. End-users' apprehension about **inadvertent** and **deliberate privacy violations** is likely heightened by the extensive abstraction and information hiding done within the COLLABRARY shared dictionary making the system very unaccountable (Dourish, 2001) and unbelievable. While rapid-prototyping could ameliorate end-users' feelings of **resentment** by increasing their involvement at earlier stages in the development process, support for this hinges on a software engineering methodology that provides for it.

### 9.3.3 Step 3: Reveal

The COLLABRARY was developed without the benefit of the vocabulary and in the description above several omissions from its feature set are accordingly noted. Yet, the unstated assumptions about what privacy is and how it is to be preserved that guided the design of the COLLABRARY and its feature set stem not solely from a lack of a holistic perspective of privacy fostered by the vocabulary. Some of these assumptions stem from a necessary narrowing of perspective due to the constraints of hardware and software technology. Some of these assumptions arise naturally from the selection of audio/video channels as the primary embodiment and medium for interaction over distances. In seeking to explicate these hidden assumptions, I build a picture of what the COLLABRARY can do well and what it cannot do well.

To begin, I will take a general framework for studying awareness and computer mediation of awareness put forth by Greenberg & Johnson (1997) and reinterpret their framework as it might apply to privacy. This reinterpretation teases apart four distinct ideas:

- there is some critical **information** about the **performance** needed to regulate privacy;
- there is a need for a way for digital computers to **capture** or **infer** this critical information;

- there is a need for a way for digital computers to transcode and **communicate** (or **store**) this critical information; and,
- there is a need for fitting ways of **presenting** or acting upon the critical information to alter the **performance** or its observation.

The vocabulary is valuable for addressing the first point in that it enumerates the high-level critical information people use to regulate privacy. The critical information related to **solitude** includes **availability**, **relevance**, **salience**, **distraction**, and **attention**. To borrow from Goffman's dramaturgical model of social interactions, these comprise a model of the **actor** and the **audience** for a performance. The critical information related to **confidentiality** include topics like **whereabouts**, **actions**, **utterances**, and **interactions**, as well as qualities like **sensitivity**, **intelligibility**, **fidelity**, **accuracy** and **precision**. These comprise a model of the content of a **performance**. The critical information for **autonomy** includes **identity**, **impression**, **roles** and **signifiers** of these things (**props**, **costumes**, **manners**, and **territories**) as well as **expectations**, **acceptability**, **conformance** and **deviance**. These comprise a model of the **social setting** for a performance. The critical information for the **mechanics of privacy** includes **genre of disclosure**, **risk**, **reward**, **violation**, **severity**, **policing**, and **reprimand**; these comprise a model of the **intention** for **interaction** and **withdrawal**.

The critical information enumerated by the vocabulary consists mostly of highly subjective or inter-subjective phenomena. It has been repeated time and again that computers are not adept at capturing, inferring, or transcoding subjective information and because of this the COLLABRARY is quite limited in the kinds of critical information it can capture. The choice made in the COLLABRARY is to focus on providing audio- and video-mediated embodiments for awareness and interactivity. In doing so, the COLLABRARY is able to *incidentally* capture the subjective critical information enumerated by the privacy vocabulary. People are very skilled at perceiving the subjective critical information in a performance given the objective audio/video representation of it, but there are few algorithms that enable computers to do the same, to go beyond the raw sensation to construct the kinds of models of the performance described earlier. Hence, the COLLABRARY does not provide the end-programmer with components that tell her how **available** a local end-user is, but it does provide the end-programmer with

components that provide a video image which would allow a remote end-user to determine the availability for himself.

Consequently, any privacy regulation operations performed with the COLLABRARY must act indirectly, altering not the information or the performance but rather its audio/video representation. The COLLABRARY is thus ideally suited for the prototyping of techniques for regulating privacy that make inferences from captured audio/video or manipulate the raw audio samples or video pixels. The way this support is realised in the COLLABRARY is by utilising architectures and APIs that offer end-programmers effortless read/write access to the individual pixels in video frames and samples in sound buffers.

The distinction between the **information** (perception) and its **representation** (sensation) has over-arching implications on the kinds of privacy regulation that can be achieved with the Collabrury. For example, there is no way to with the COLLABRARY components and methods alone **edit** out the socially unacceptable qualities of a video frame from a performance while leaving the rest intact. Instead, the kinds of manipulations automated by the COLLABRARY are less targeted. For example, the distortion filters provided by the COLLABRARY affect every pixel in a video frame, be it a part of the socially acceptable aspects of the scene or otherwise. Similarly, the COLLABRARY automates making a video frame more socially acceptable by given the end-programmer methods useful for comparing the video frame on a pixel-by-pixel basis against a database of stock video frames determined beforehand to be socially acceptable (perhaps by the end-user) and use as a substitute the frame from the database that is the closest match as given by the peak signal-to-noise ratio of the pixel-wise difference image. A further elaboration of this technique that is also possible with the COLLABRARY (although not automated to the same degree) is the use of eigenfilter values in the comparison. Although the COLLABRARY is itself not suited to prototyping techniques for preserving privacy that utilise the kinds of subjective information enumerated by the vocabulary, it still can be claimed that the COLLABRARY supports these techniques in that it provides a foundation for the prototyping of algorithms that infer this subjective information from captured audio and video or manipulate the audio/video presentation of an model of the performance.

Continuing the examination of the COLLABRARY using the framework derived from Greenberg & Johnson, it is clear that the COLLABRARY shared dictionary is ideally suited to the **communication** of digital encodings of the objective representations of the critical information needed for privacy regulation. Although the shared dictionary's ability to automatically marshal complex data types means the end-programmer can use nearly any encoding that makes the information available for digital computers to process, the COLLABRARY is particularly suited to encodings like compressed audio and video. The hierarchical organisation of data in the shared dictionary gives the end-programmer a way to **aggregate** many kinds of critical information into higher-level abstractions of the critical information that models the performance.

With respect to the **transmission** of critical privacy-regulating information, some omissions described previously are worth repeating. The omission of automated encryption of data transmitted over the network reflects an emphasis on developing safeguards against aesthetic harms instead of strategic harms. Also, the shared dictionary in particular does not easily handle multiple concurrent representations of the same piece of information. Although end-programmers are free to devise schemes whereby some people access information in the shared dictionary at a different fidelity than others, the fundamentally broadcast nature of the shared dictionary means that—in typical scenarios—all fidelity levels are accessible to every computer even if all but one of the fidelity levels are disregarded by the application. Finally, although the shared dictionary's use of pseudorandom identifiers for client computers provides some pseudonymity, end-programmers can easily obviate it by merely storing person- or computer-specific information in the dictionary, such as a video frame from a connected webcam.

In regards to the last item raised in the framework derived from Greenberg & Johnson, the COLLABRARY itself provides few mechanisms for presenting or acting upon the critical information needed to regulate privacy. Although the COLLABRARY provides a **Speaker** component for the rendering of audio, end-programmers are assumed to use some external GUI toolkit for presenting video and consequently the COLLABRARY is designed to interoperate well with the image display features of conventional GUI toolkits. The assumption that belies these design decisions is that critical information for privacy regulation

captured as video will be presented as video and information captured as audio will be rendered as audio. The presentation is tightly coupled to the representation.

Finally, expanding upon the four items derived from Greenberg & Johnson's framework, the COLLABRARY does not offer any guidance for the design of privacy safeguards. This guidance could come in the form of a prescription or in the careful arrangement of affordances and constraints (make "good" privacy practises easy and "bad" privacy practises difficult). Yet, as discussed in the previous chapter there is no prescriptive theory upon which the COLLABRARY could be based. It is not fully understood which practises are "good" although it is generally simpler to tell if a given practise is bad.

#### 9.3.4 Step 4: Summarise

The COLLABRARY supports the rapid prototyping of video media spaces that incorporate techniques for preserving privacy that analyse and manipulate audio- and video-based representations of performances by (a) greatly decreasing the effort end-programmers must expend to gain direct read/write access to pixels in video frames and samples in audio buffers and, (b) providing a flexible mechanism for organising and sharing captured audio/video and other digitally encoded information inferred from these sources. Although end-programmers can leverage these features to implement novel approaches for regulating privacy, this constitutes all but the most rudimentary level of support that a toolkit could offer. These restrictions on the ways the COLLABRARY supports the development of privacy safeguards for video media spaces arise out of ignorance of a holistic perspective on privacy, out of choices that fall naturally from an emphasis on audio/video embodiments for awareness and interaction, out of limitations of hardware or software technology, and out of a lack of prescriptive theory on how to design computer support for regulating privacy.

## 9.4 Conclusions

This chapter presented three case studies that illustrate the application of the analysis method developed in Chapter 8.

**Case study #1** looked at the distortion filtration technique itself. This case study showed that in the technique itself there are embedded certain assumptions about what will

cause a privacy violation that affect the circumstances in which the technique will actually preserve privacy. The important conditions on the success of this technique are that the sensitive information in a scene must be encoded in high-precision visual information, that differences between the actual performance and a socially acceptable performance must be encoded in high-precision visual information, and that there is minimal risk for misinformation.

**Case study #2** looked at the evaluations of the distortion filtration technique that were performed for Chapter 4. This analysis revealed that the methodology employed in these evaluations constrained the scope of privacy to only the most rudimentary aspects of confidentiality.

**Case study #3** revisited the COLLABRARY as a toolkit for rapidly prototyping video media spaces. This analysis revealed that the COLLABRARY toolkit does not expose privacy-related information at the same conceptual level as the vocabulary e.g., sensitivity, attention, social conformance. Consequently, end-programmers can only manipulate privacy indirectly, manipulating representations of these things in terms of pixels. As was the case with the distortion filtration technique and the evaluations of it, this constitutes only the most rudimentary level of support for privacy—specifically confidentiality with secondary effects on autonomy.

In these three case studies, I have re-examined the work I presented in Act I of this thesis—dealing with low-level technological factors for privacy in video media spaces—using the concepts and vocabulary developed in Act II. The next two case studies will examine work by others outside this thesis.

## Chapter 10—Case studies (2)

This chapter continues the case studies that illustrate the application of the analysis method developed in Chapter 8. These case studies have been chosen because they position my work along side that of others and illustrate an important contribution I have made with this method. The method supports our efforts to comprehend our systems' effects on privacy and to **share our understanding with others**. In turn, these case studies identify opportunities to improve the method and expand the descriptive theory of privacy in video media spaces.

### 10.1 Neustaedter & Greenberg's HOME MEDIA SPACE

Neustaedter & Greenberg (2003) describe a prototype reactive video media space for use in a home office so that a home-based telecommuter may stay connected with a colleague elsewhere: the HOME MEDIA SPACE (HMS). Their system was *reactive* in that the media space operation changed in response to input from various environmental sensors that were part of the installation (Figure 10.1). These reactions were designed to safeguard against inadvertent privacy violations that can easily arise when persons in two very distinct social settings such as a bedroom in the home which doubles as an office, and an office in a corporate headquarters.

The camera state (capturing or disabled) and orientation (facing into the room or against the wall) were key media space parameters controlled in HMS. Sensors determined the presence of the telecommuter and the presence of family members in the room into which the media space was installed. Heuristics used these sensors to decide when the camera should be turned off and pointed it to face the wall, to preserve privacy. Physical sliders (see the figure)



Figure 10.1—The Home Media Space prototype included environmental sensors that were used to autonomously enable/disable capture. The camera was mounted on a servo motor so that it could be rotated to face the room or face the wall. (Figure reproduced from Neustaedter & Greenberg, 2003.)

added further fine-grained control over camera angle, the level at which a blur filter was applied, and the video frame rate.

Particular attention was paid to the design of believable feedback cues regarding system operation and lightweight explicit and implicit mechanisms for controlling media space operation. While their paper discusses the design and its rationale, it does not discuss any evaluation of the system with users. The goal of this case study is to demonstrate how the descriptive theory of privacy presented in Act II can permit detailed understanding of the kinds of privacy problems that various proposed safeguards might best address (in addition to those left unaddressed) and circumscribe limits or conditions on these safeguards' efficacy.

### 10.1.1 Partition

The HMS uses a variety of sensors and heuristics to control the capture and broadcast of video and audio in a media space. By *heuristics*, I mean a set of rules that describe a desired state for media space operation given combinations of inputs from the environmental sensors. In their paper, Neustaedter & Greenberg make use of various scenarios for HMS usage as a vehicle for explaining the heuristics themselves and the rationale behind them. In this first step of partitioning the privacy space, I will examine the prototype functionality and not the case

studies; the case studies will be examined in the second step of describing in detail the work's relationship to privacy.

This prototype certainly deals with **solitude** in that it permits mediated interactions over distances, allowing people who would otherwise be kept apart the opportunity to come together. Like all media spaces, emphasis is placed on **togetherness** and **interaction** instead of **withdrawal** and **reserve**. Also, the effects of the prototype on attention are not discussed in the authors' paper.

The HMS prototype certainly deals with **confidentiality** in that it regulates the **capture** and **transmission** of **sensitive** audio and video **information**. (The authors themselves use the term "sensitive" to describe the contents of performances mediated by the prototype.) The prototype permits **modulation** of **disclosure** by way of starting or stopping capture or by applying a blur distortion filter, but it does not permit **disinformation**. The prototype does not **store** video, and so there are few issues with **temporal boundary** and **desituation** of action over time.

The prototype also certainly deals with **autonomy** in that some of the control over confidentiality afforded in autonomous reactions or through explicit gestures is exerted so as to ensure that audiences never see **inappropriate performances**. Such inappropriate performances—specifically inappropriate **costumes** or **behaviours** or the **presence of non-users**—are considered more probable in HMS usage because of the stark contrast between the actor's **setting** (home) and the audience's setting (office). Other autonomy-related factors beyond **impressions**, such as **role conflict** and **liberty**, also relate to the prototype, yet because the prototype is intended for use between people who already know each other very well, a topic like **pseudonymity** has no place in it.

The HMS also certainly deals with some of the **mechanics of privacy** in that it seeks to support various **behaviours** for **signalling privacy needs** and for regulating confidentiality and solitude. Also, the scenarios used to explain the prototype's heuristics examine many aspects related to risk. Not all mechanics are covered, of course: for example, there are no facilities to **cooperatively** regulate privacy in the event the contextual sensors or reactive heuristics fail. Finally, as a software system, the prototype attempts deals with issues peculiar to

**computers and privacy**, such as **risk/reward disparities**, but not others such as **deliberate misuse** and **misappropriation**.

## 10.1.2 Describe

### 10.1.2.1 Solitude

As a tool to mediate **awareness** and **interaction** via audio and video channels, the HMS prototype concerns itself with many aspects of solitude. The awareness itself becomes an important resource for **cooperatively regulating solitude**. Perhaps the most important reward that a video media space connecting a home-based worker with office colleagues offers is that it can alleviate isolation for the tele-commuter, allowing them to maintain dual **presences** at both home and work. This reward may be enjoyed even though the system as described connects only dyads and the scalability of the solution is not examined. This reward comes with risk, however, to the user. Seeking escape from their media space partner (that is, seeking refuge) becomes more error-prone as the user must remember to check the state of the camera and turn it off when appropriate. If left unchecked, there is the possibility that unconscious **back-stage performances** such as **fantasy** can be made into lapses in desired **self-appropriation** in the **front**.

Along the **psychological dimensions** of solitude, it is clear that like any media space there is a definite bias towards facilitating interaction in the HMS. Indeed, another reward of system usage is that it can help sustain **intimate working relationships** while partners who are normally together become distance-separated, e.g., during maternity or paternity leave or while school-aged children are at home from school during summer or winter breaks in the school year. The HMS is different in that it also provides sensors, heuristics, and control gestures for quickly disabling capture and transmission in the media space. While on the one hand this “ejection seat” style withdrawal lacks social grace, it is **immediate** and **believable** (e.g., the camera turns around to face the wall) and these qualities should fit well with its use as an urgent response to safeguard against **inadvertent privacy problems** as they arise and are diagnosed.

The assumed primary reward of video for distance-separated colleagues described in the previous chapter—the microcoordination reward conjecture—applies since the HMS is

designed to provide partners with **informal awareness** cues during loosely coupled work that leads into and out of casual interactions during tightly coupled work. Much like any other system that intends to present useful information in the **periphery** of one's **attention** (e.g., ambient displays), for this reward to be enjoyed the HMS must provide a reasonable balance between **salience** and **distraction**. The only heuristic in the HMS regarding this is that when the user is in the office it is assumed that awareness and interactions shift towards the focus of the user's attention. None of the sensors in the HMS measure or approximate the attention-related demands and needs of the users, however. Consequently, the support within the HMS for **anonymity** (in the sense of "going unnoticed") or **reserve** depends largely on how peripheral its display is and such support is further complicated by the fact that it is difficult for one user to gauge the extent to which his nimbus enters the focus of the other user.

#### 10.1.2.2 Confidentiality

Much of the analysis related to confidentiality in the distortion filtration case studies also applies to the HMS, as does the analysis of confidentiality in the COLLABRARY, since the HMS was prototyped using the COLLABRARY toolkit. The confidentiality features of the HMS aim to safeguard against the unintended delivery of back-stage performances on the front stage to an audience intended to see the front-stage performance. These features are important because it is expected that the large differences between front and back stage performances between the home settings and office settings in which the HMS is to be used increase the **risk** associated with such lapses (certainly the **probability** and possibly also the **severity of harms**). The HMS adds to this body of analysis the idea of heuristics for autonomously regulating capture. This is an important addition because it helps facilitate **refuge**, eases **apprehension**, and makes the system more **accountable**.

The HMS provides audio and video channels for mediating **embodiments**. Much of the control over confidentiality provided by the HMS is control over **sampling** in the audio and video channels. By far, video is given the deepest treatment as it is the channel kept on most of the time whereas audio uses a push-to-talk model. The information sampled by the context sensors (or inferred from them) in the HMS is used to regulate audio/video sampling; this information is neither transmitted nor presented to others and thus cannot be used as

interactive media in their own right. Sampled audio is presented as-is; sampled video may be passed through a blur distortion filter but it is also presented as video.

Neustaedter & Greenberg motivate their work on a privacy-preserving reactive video media space by pointing out that some of the **topics** in audio and video (as discussed in the previous chapter) may contain **sensitive** information. They specifically focus on the sensitive information in video which is characterised as **costumes** and **actions** that while entirely appropriate for the home might be considered **inappropriate** by office colleagues. These kinds of sensitive information lead to **aesthetic harms**; sensitive information such as a user's daily routine or personal habits and the security features in his or her home—information which may lead to **strategic harms**—is also transmitted in the video channel but is not considered in the authors' discussions.

In the HMS, all information is ephemeral (not archived) and is broadcasted synchronously and in real-time. There is no need to consider **temporal boundaries** or the role of **temporal precision** or **accuracy** in establishing **ambiguity** that permits **plausible deniability**. The system incorporates manual control over the application of a blur **distortion filter** that modulates video **precision**. While the content ambiguity fostered by the filter creates some chance for **misinformation** and plausible deniability, precision modulation offers little chance for controlled **disinformation** (as discussed in the previous chapter). It is not clear how the system would broadcast information when scaled to three or more users. Although it is audibly and visibly apparent to the user when the system is capturing audio or video, there is no way to diagnose or safeguard against the **misuse** of captured information. The HMS itself provides no means for autonomous **scrutiny** of the captured video and since it is a closed system (meaning that there is no way to extend it with third-party add-on extensions) there is no way that analysis of the video or audio could be automated. Scrutiny by humans—particularly superiors at the corporate office—is a concern, however. The authors assume that the system would be used by work peers so that issues related to **status** might be ignored. The system autonomously disables capture in a number of circumstances as part of its normal operation and these periods of **refuge** might safeguard against scrutiny.

### 10.1.2.3 Autonomy

Much of the analysis of the preceding chapter which deals with video as an embodiment for **social constructions of the self** applies to the HMS since video (and, to a lesser extent, audio) is the only embodiment it provides. For example, the HMS uses the COLLABRARY and consequently lacks a model of what constitutes a **socially-appropriate** performance or detect when a performance **deviates** from such **norms**. Further elaboration is necessary, however, because the home setting for HMS usage fosters autonomy-related problems.

First, the fact that it is a media space connecting home to office obfuscates which **norms**—ones for the home or ones for the office—apply to performances. The HMS tries to address this issue by allowing the user to apply a blur distortion filter and by autonomously disabling capture in certain situations. Second, it introduces **obligatory relationships** from work that were previously kept outside the home and vice versa challenges these work relationships with the demands of home life. These two factors combined create opportunity for **role conflict**. This includes not only conflict in norms and **expectations** for **behaviours**, particularly **actions**, **reactions**, and **interactions**, but also **costumes** and **manners**.

The authors particularly emphasise the **aesthetic harms** to one's socially constructed self that might arise from such role conflict, providing the distortion filter, for example, as a safeguard. Furthermore, the authors make explicit an assumption that the system is to be used by intimate collaborators who are peers, but even close colleagues at work may not have access to the details of each other's home life. The authors consider only the actor's **costume** as a potential **risk**. The **props** in the room however ought to be considered as needing to be regulated by the actor's self-appropriation. The video camera captures the actor and the **props** in the setting with roughly equivalent resolution. In doing so it ascribes each the same importance, yet the **props** in the room in which the actor works are not nearly as important for maintaining an intimate working relationship over distances. Indeed, home offices are often in small, cramped rooms which might be too untidy to ever let a visitor in the home see. During periods of tightly coupled interactivity, a distortion filter applied uniformly across the image might be **distracting**; selectively distorting the entire scene except the actor might be preferred.

Finally, the HMS is designed to disable capture and transmission when a **non-user** such as a family member enters the media space. Doing so helps support the autonomy of the non-user by maintaining the refuge-like qualities of the home. One could speculate that this also helps alleviate problems with non-users' **apprehension** of the privacy-insensitive nature of video media spaces.

#### 10.1.2.4 Mechanics of privacy

The HMS permits users the ability to transcend **spatial boundaries** for interaction. This is, after all, the way the rewards for using the HMS are gained. Yet, these boundaries are also part of privacy regulation: spatial boundaries permit people to be apart (solitude) and reduces the **fidelity** of information **access** (confidentiality) which in turn supports **liberty** (autonomy). Spatial boundaries become eliminated from the HMS user's repertoire of privacy regulating mechanics. As stated in the previous sub-section, the HMS may provide an office colleague access to the details of a person's home life that he might not otherwise have (e.g., if their working relationship never transcended the workplace) and consequently, the **disclosure boundary**, which here would concern the revelation of information in order to sustain an intimate collaborative relationship, is made more difficult to regulate. Role conflict makes **identity boundaries** more difficult to regulate. **Temporal boundaries** for privacy regulation are not affected though because all captured video and audio is **ephemeral**.

When they designed HMS, Neustaedter & Greenberg paid careful attention to the behaviours which people employ as part of regulating privacy. The features they engineered reflect the **dynamic** and **regulatory** nature of privacy. The inclusion of an audio channel supports explicit **dialectism**. Although the HMS may autonomously disable capture when its heuristics decide capture is inappropriate, similar control is not offered to the audience; thus, the HMS denies some aspects of the cooperative nature of privacy regulation.

Neustaedter & Greenberg also use a selection of scenarios to ground the HMS design in reality. These scenarios portray the very real risks that the authors imagine HMS users might face on daily basis. The scenarios range from the user picking his nose on camera, to a non-user entering the home office (which doubles as a bedroom) to get dressed after bathing. These scenarios depict risks in which the harms are **aesthetic**, of low to moderate **severity**,

and can be easily imagined as probable to occur. In fact, the ease at which these violations could occur likely adds to users' and non-users' apprehension of the system. The inclusion of non-users in some of the scenarios makes them useful in examining the efficacy of the safeguards given a **risk** and **reward** that ought to be traded off much differently than for users.

When dissected with the analysis vocabulary, though, the limits of the scenarios become visible. All of the actions and costumes in the scenarios are **normal** and **appropriate** for the actor when alone in the home setting. While these actions and costumes might be abnormal in the presence of the office colleague as audience (even if the colleague were with the actor at home) or if the actor was at the office (even if alone), they are not outrageously or offensively **inappropriate**. They are perhaps so mundane that it is likely easy for audiences to sympathise with the actor, reducing the cost of the aesthetic harm. In the end, the harms faced by the user are mostly momentary embarrassment, and if a non-user is involved in the violation there may also be some arguments or resentment or other strain on the domestic relationship.

The system's features prescribe a **genre of disclosure** for different situations to compensate for lapses in self-appropriation made possible by the separation of actor and audience. In the system, the genre of disclosure is defined by sampling parameters (capturing or not capturing, looking at wall or looking into room) and affects both the level of interaction possible and the risk of inadvertent privacy harms. Three genres of disclosure are prescribed, as given in Table 10.1. Additionally, within the Started genre, parameters like frame rate and field of view and precision as modulated by the blur distortion filter are under manual control.

A video media space necessarily has in its design support for a set of idealised genres of disclosures. Usually, there are at least two: on (full disclosure) and off (non-disclosure). Other genres are possible by, for example, incorporating the content control techniques discussed in the privacy vocabulary. The genre of disclosure establishes a pattern of interactions. They are predicated upon **institutional signifiers**—normalised expectations for performances—based upon such things as actor and audience roles and the setting of the performance. These institutional factors give a rough approximation of the interaction needs between actors and their audiences in a particular setting. A video media space has idealised notions of who the actors may be, who the audience may be, and where the setting is. Therefore, the media space design is essentially a function  $F(\text{actor}, \text{audience}, \text{setting})$  genre . These baseline expectations are

Genre	Capturing	Image shown	Interactivity	Risk
Stopped	No (camera pointed at wall)	Wall	None	None
Paused	No (camera pointed at wall)	Room	Minimal	Low
Started	Yes (camera pointed into room)	Room	Rich	High

Table 10.1—Genres of disclosure in Neustaedter & Greenberg (2003) HOME MEDIA SPACE.

subsequently modified by situational factors, such as interaction trajectory, appearance and preference. If we abstract away situational factors we can establish for particular  $\langle \text{role}, \text{role}, \text{setting} \rangle$  tuple a genre of disclosure. Thus, if we idealise the genre of disclosure we idealise the pattern of interactivity supported by it. Similarly, people may have their confidentiality, solitude, and autonomy put at risk by these interactions and so by idealising the genre of disclosure we also idealise estimates of the risk inherent in it.

The system provides audio as an explicit signifier of privacy **desires** and capitalises on some objective facta to determine a few implicit signifiers that regulate automated transitions in genre of disclosure.

Beyond genres of disclosure, which really only provide institutionalised control over disclosure, there are few other **behavioural** or **cognitive mechanics** for regulating privacy directly supported in the HMS prototype. We see similar omissions regarding **policing**, **reprimand**, and **reward** as with the distortion filtration technique and its evaluation (Chapter 9). The HMS confounds **disinformation** as a mechanism for regulating confidentiality. It is difficult to tell how well **reserve** is supported because the effects of the HMS on **attention** are so unknown.

Neustaedter & Greenberg have engineered into the HMS prototype great consideration for environmental mechanisms for regulating privacy. The mirror image supports **reflexive interpretability of action** when a person clearly sees himself in the image. Changes in the environment that affect privacy and its regulation are signalled with transitions marked by well-thought out feedback cues. Naturally, the HMS is constrained to video only for embodiments and this in turn affects how people can negotiate disclosure and identity boundary tensions.

In sum, the greatest motivator for users' and non-users' apprehension about the HMS could be that for at least part of the day front-stage performances must happen in a place where almost exclusively back-stage performances happen, yet the environment—the spare bedroom or whatever it is that is serving as a home office—has not been re-architected in the physical environment or in people's conceptions and customs to introduce new constraints to liberty.

#### 10.1.2.5 Computers and privacy

The HMS is built using the Collabratory and it does not add any of its own **computer security** features. Hence there are no safeguards against eavesdroppers, hackers or other people with malicious intent. Other security measures, such as **pseudonymity** and **access control** within the set of intended users might be inappropriate because the set of intended users are already intimate collaborators. The **content control** measures—a distortion filter and an on/off publication filter—however fit well with the kinds of inadvertent **aesthetic harms** that are expected to arise in HMS usage. **Reliability** is not singled out in the authors' descriptions of the system as an area which has been specifically dealt with.

The authors seek to address **inadvertent privacy infractions**, thinking these are the ones people are most readily apprehensive about. As a result, many of the safeguards they have engineered intend to address the “four ‘D’ problems” of **decontextualisation**, **disembodiment**, **dissociation**, and **desituated** action. They also seek to ameliorate problems that arise from the unique placement of the media space into the home such as resentment from non-users and role conflict. In contrast to this emphasis on inadvertent violations, there is little done in the HMS prototype to safeguard or address **deliberate privacy violations** such as **identity theft** or **misuse** of information.

User interface issues come into place and some are addressed head-on by the prototype, others omitted, and still more are genuinely irrelevant. The increase in **spatial degrees of freedom** for information access is part of the **reward** of using the system. The authors try to balance risk and reward by detecting situations in which non-users are at **risk** and then selecting a genre of disclosure in which they are not at risk. They have marked **transitions** with **salient** and **believable feedback** (for example, the motor which rotates the camera into

or out of view makes a distinct and audible sound) and have sought to decrease **cognitive** and **physical effort** in regulating control by devising **lightweight explicit control** user interfaces as well as providing some automated control. The **granularity** of this control is fairly coarse and although the blur filter control is rather fine grained, it permits only one degree of freedom. It is unclear how well their user interfaces will maintain their lightweightness if the system is scaled to three or more users.

### 10.1.3 Reveal

The designers of any media space must necessarily make assumptions about the circumstances in which their systems will be used. Here there is an assumption that no single hardware/software system will do everything and so in fixing these circumstances designers establish a domain for their system and bring about the specification of software. The descriptive theory of privacy reveals four categories of such “circumstances of use” and provides the vocabulary needed to articulate each. Much as with the previous case studies, the vocabulary-based analysis can reveal opportunities to design new systems by permitting systematic alteration of each of these predicative assumptions.

**Actor and audience:** designers need to know users and non-users. This is true of any system to be designed, although the characteristics of the user population that are of greatest interest to the designer may vary with the system domain. In the case of the HMS, the users are distance-separated information workers who share an intimate a-priori collaborative relationship. The non-users may be other individuals who occupy the same physical space as the users (specifically family members, but we could also imagine office colleagues, visitors to the home or office, or cleaning staff). A group of non-users omitted from the HMS design includes others on the Internet who can intercept HMS network communications.

**Reward:** designers need to know what purpose the system will serve. This is true of any system to be designed. In the preceding sections, the rewards discussed include more satisfying collaboration via casual interactions, improved microcoordination of casual interactions via enriched informal awareness, maintenance of a sense of intimacy in the relationship, and reduced feelings of isolation and increased presence at the workplace for a tele-commuter.

**Social setting:** designers need to know the context in which the system will be used. This is true of any system to be designed, although the contextual details that are most interesting to the design may vary with the system domain. In the HMS one person will be in a home office social setting while the other will be in a corporate office social setting.

**Performances:** designers need to know what people will do with the system. This is true of any system to be designed, although the interesting details of use may vary with the system domain. In the HMS and the case studies used to ground it, the scenery and props are assumed to be inoffensive and normal for the actor's setting although they might be inappropriate for the audience's setting; the costumes in the corporate office are assumed to be fully clothed at all times and those in a home office setting may vary from fully clothed to fully naked, but are assumed to be inoffensive and normal for the actor in his or her setting; and, the actions and utterances are assumed to be inoffensive and normal for the actor in his or her setting and are assumed to not disclose information peripheral to the collaborative activity in which the users are engaged. (Other signifiers of course disclose information peripheral to the collaboration.)

All of these assumptions determine the kinds, severity and likelihood of various privacy risks and the kinds, costs, and efficacy of countermeasures for these risks. In other words, the problem domain determines much about the problems to solve and the solutions for them. The next task for the designer is to make assumptions—or, perhaps better put, decisions—about which risks to counter and how to counter them. Here there is an assumption that we cannot satisfactorily counter every risk. We might not know about some risks. We might not have effective countermeasures for some risks. We might not have sufficient resources to implement an effective countermeasure for every risk, and so we must trade off the cost and benefit of each countermeasure. Also, we might misunderstand the severity or likelihood of some risks and consequently make bad decisions in this trade-off. The vocabulary permits articulation of the risks so that ideally no risk goes unknown. Also, the vocabulary permits description of the means by which countermeasures manipulate risk and this might help facilitate generation of new countermeasures. Unfortunately, as stated at the beginning of this act, the vocabulary provided by a descriptive theory of privacy does not guide the designer through the process of deciding which risks to counter or evaluating the countermeasures.

<b>Feature</b>	<b>Inattention slip conditions</b>
autonomous activation and deactivation of capture	no one is in the room -or- a non-user is in the room
lightweight manual capture override	is diagnosed very quickly by the user
blur distortion filter	see previous case studies

Table 10.2—Conditions in which HMS privacy preserving features will counter inattention slips in self-appropriation.

In the case of the HMS, the risks to be countered are inattention slips (Reason, 1980) in self-appropriation and the countermeasures involve the semi-autonomous regulation of capture and transmission to ensure conformance to a designer-specified genre of disclosure. We can describe this genre of disclosure simply. The audience is permitted to see the user and the room when the user is present in the room and is alone. Capturing is autonomously started or stopped by the system to conform to this rule when the user enters or leaves the room or when a non-user enters or leaves the room (regardless of if the user is present). Furthermore, the audience may see things from only a user-chosen angle and at a user-chosen precision. The actor (user or non-user as the case may be) can intentionally permit or deny the audience from viewing things at any time, overriding autonomously activated or deactivated capture. Anytime the camera starts or stops capturing, an audible cue is heard. Given this, the inattention slips in self-appropriation that are successfully countered by the features in the HMS are stated in Table 10.2.

#### 10.1.4 Summarise

Neustaedter & Greenberg's HOME MEDIA SPACE is a prototype media space to be used between dyads where one person is in a home office and the other is in a corporate office. These circumstances of use create opportunities for inattention slips in self-appropriation and foster risk-reward disparities for non-users who may co-occupy the physical space in which the HMS operates. To counter these problems, the HMS includes three privacy-preserving features: autonomously regulated capture; manual capture override; and, blur distortion filter. These features will counter inattention slips in self-appropriation when they are diagnosed quickly by the user, or involve the presence of a non-user, or when the sensitive aspects of a scene or differences in ideal and actual self-appropriation are in high-precision visual

information and any misinformation arising from the loss of this high-precision visual information casts the actor favourably. There are still many other privacy problems not countered by the HMS prototype. Furthermore, it is easy to imagine situations in which the conditions on the HMS countermeasures' success may not be met. Despite these caveats, however, if the HMS countermeasures are successful when these conditions are met, then the countermeasures will go a long way in resolving users' and non-users' apprehension towards inadvertent privacy violations that may arise when using the system.

## 10.2 Hong's (2004) privacy risk model framework

Hong et al<sup>3</sup> (2004) develop a method which can be applied by designers to identify and prioritise concrete interpersonal privacy issues in a ubiquitous computing system. By concrete, they mean privacy issues that can be describe with sufficient specificity so as to permit formulation of a plan to alter the system design or its implementation and resolve the issue. They call their method *privacy risk models* and it consists of two modes of analysis to be performed sequentially.

The first, their *privacy risk analysis* method, has the analyst answer a compact set of questions about the system: users, their relationships, their motives for using the system, the user interface, and its implementation. These questions are listed in Table 10.3.

Although ubicomp systems are mentioned specifically as the primary domain for their own investigations, the questions are sufficiently general that they may be applied to nearly any privacy-sensitive computing system. The answers to these questions comprise a description of the privacy risks inherent in the system or its use. Not every imaginable risk will be found: the authors have engineered into their set of questions a pronounced bias towards serious or likely risks. If these are fixed, then a reasonable level of privacy protection that is commensurate with other system factors might be reached.

---

<sup>3</sup> From time to time, for brevity and clarity, I may label the framework by the primary author's name only, e.g., Hong's framework.

The second part of their method is called *privacy risk management*. This method also includes a set of questions that guide the analyst (Table 10.4). It permits cost-benefit type analysis to prioritise the risks identified in the first mode of analysis. The analyst roughly assesses for each privacy risk identified: the likelihood that the violation of the risk will occur; the severity of the harms that would arise if the violation actually happened; and, the cost to implement a countermeasure.

Unlike my top-down theory-driven approach, Hong uses a bottom-up practice-driven approach to arrive at their method. They base their method previous work in task analysis and in computer security threat models, both of which are methodologically similar to their approach, as well as in fair information practices, which provide general guidelines that are a good but still imperfect fit with the unique problems in interpersonal privacy. They also examined a variety of different systems and compiled a list of common patterns of privacy issues shared across them. From these starting points they developed an analysis method and the materials used therein.

In this case study, I will apply my framework for analysing privacy to another framework for analysing privacy. As with the previous case studies, from this analysis emerges a description of the kinds of privacy issues that will be addressed by the object of analysis—previously prototyped countermeasures and now a privacy analysis framework—and the

#### *Social and Organizational Context*

- Who are the users of the system? Who are the data sharers? Who are the data observers?
- What kinds of personal information are shared? Under what circumstances?
- What is the value proposition for sharing personal information?
- What are the relationships between data sharers and data observers? What is the relevant level, nature, and symmetry of trust? What incentives do data observers have to protect data sharers' personal information?
- Is there the potential for malicious data observers? What kinds of personal information are they interested in?
- Are there others stakeholders or third parties that might be directly or indirectly impacted by the system?

#### *Technology*

- How is personal information collected? Who has control over the computers and sensors used to collect information?
- How is personal information shared? Is it opt-in or opt-out? Do data sharers push personal information or do data observers pull information?
- How much information is shared? Is it discrete and one-time or continuous?
- What is the quality of the information shared? With respect to space, is the data at the room, building, street, or neighborhood level? With respect to time, is it real-time, or is it several hours or even days old? With respect to identity, is it a specific person, a pseudonym, or anonymous?
- How long is personal data retained? Where is it stored? Who has access to it?

Table 10.3—Questions used in Hong's Privacy Risk Analysis method.

conditions under which they will and will not be addressed. The converse analysis—using Hong’s method to analyse mine—cannot be done due to the fact that all of the questions used in Hong’s analysis framework assume the object of analysis is an implementation of a computing system design. In the final steps of the analysis in this case study, rather than picking out the assumptions that belay Hong’s model, I compare their analysis method to my own to show how they are similar and different with the explicit intent of showing ways either may be improved.

### 10.2.1 Partition

Looking at the questions in the analysis framework (Tables 10.3 and 10.4), it is immediately evident that “privacy” as conceived in the framework is exclusively **confidentiality**. Although **solitude** and **autonomy** may be affected, they are not analysed as part of the framework. Correspondingly, the synergistic relationships between confidentiality and the other modalities of privacy control are not analysed, either. Hence, in the next step of the analysis, I will not use the vocabulary terms from those two sections. There are additional concepts connected to confidentiality that are explicitly considered in Hong’s query-based analysis framework, most notably **risk/reward trade-offs** and **access control**. Correspondingly, in the vocabulary-based analysis method I use here, these concepts fall under the **mechanics of privacy** and **computers and privacy** categories. In the next step of the analysis, I will use vocabulary terms from those two sections in order to give complete coverage of what Hong’s analysis framework reveals.

#### *Managing Privacy Risks*

- How does the unwanted disclosure take place? Is it an accident? A misunderstanding? A malicious disclosure?
- How much choice, control, and awareness do data sharers have over their personal information? What kinds of control and feedback mechanisms do data sharers have to give them choice, control, and awareness? Are these mechanisms simple and understandable? What is the privacy policy, and how is it communicated to data sharers?
- What are the default settings? Are these defaults useful in preserving one’s privacy?
- In what cases is it easier, more important, or more cost-effective to prevent unwanted disclosures and abuses? Detect disclosures and abuses?
- Are there ways for data sharers to maintain plausible deniability?
- What mechanisms for recourse or recovery are there if there is an unwanted disclosure or an abuse of personal information?

Table 10.4—Questions used in Hong’s Privacy Risk Management method.

### 10.2.2 Describe

**Confidentiality:** Hong’s framework will ostensibly work with any **medium**. Unlike the vocabulary based method I use—which was developed for privacy-sensitive video media spaces—Hong’s method targets a much broader class of ubiquitous and context-aware applications. While my vocabulary based method could be applied to ostensibly any medium for information, it is clear that a substantial portion of the vocabulary applies solely to video and audio channels. The questions in Hong’s method lead the analyst to especially consider the sources of information. The vocabulary further deconstructs the mode of processing that was used to obtain the information: **sampling, interpolation, aggregation, or inference**. In the case studies, Hong et al concern themselves mostly with direct sampling of the environment. In a similar way, while Hong’s framework asks “What kinds of personal information are shared?” the vocabulary further deconstructs **topics** of information for confidentiality into categories such as **information about the self, personally identifying information, activities, whereabouts** and so forth.

Hong’s analysis framework does not specifically target **sensitivity** as a characteristic of the information to be regulated by confidentiality; instead, this vocabulary concept is handled in their risk management framework as “damage.” Similarly, certainty is not specifically questioned in Hong’s analysis framework, but it is given first rate treatment in their risk management model: they ask if there are ways for data sharers to maintain **plausible deniability**, underscoring the value of **ambiguity** to confidentiality.

Other information characteristics are analysed specifically in Hong’s framework, even to a greater degree of detail than in my vocabulary-framework. On the one hand, this greater detail is a natural consequence of the fact that Hong’s framework focuses specifically on confidentiality, but at the same time also indicates places where the vocabulary used in my method could be enriched slightly. For example, my vocabulary-based method lists **persistence** as an information characteristic. Hong’s method looks at freshness, lifetime, and update frequency: persistence deconstructed, if you will. With respect to **transitivity** as an information property, Hong’s questions tease apart the notion of intentional transmission and incidental transmission that is a consequence of the sampling or the use of third-party networks for intentional transmission. For **fidelity**, the questions in the framework examine

data “quality” which is further broken apart into what would be termed **content precision**, and **temporal precision** and **temporal accuracy**; the accuracy of the content, however, is never explicitly questioned. Consequently, the roles of misinformation and disinformation in confidentiality regulation are not revealed with Hong’s analysis method. Hong also considers the quality of personally identifying information, which we can consider as a special case of the more generalised content precision dimension. Hong’s teasing apart of fidelity into temporal and content dimensions for precision and accuracy is a valuable addition to my analysis method’s vocabulary.

The questions in Hong’s analysis framework do not directly address issues related to operations that can be performed on information, including use. While they look at **capture** and **archival**, they do not ask the question “Can the data be **edited**?” They ask about the value proposition for sharing information but not for observing the information; the subtle difference is important if we consider **misuse** and **misappropriation** of information, and if we broaden the scope of analysis to include solitude-related phenomena like **attention** and **distraction**.

Perhaps the most useful distinction between Hong’s framework and the one I provide deals with the people in the system. Hong asks about who is sharing and who is observing the data but not what these people can do with the data. This teases apart a distinction lost in my vocabulary based method, which categorises individuals as “**users**” or “**non-users**.” This user/non-user distinction reflects participants’ intentions and designers’ anticipations about participation and underscores a bias that some degree of reciprocity is a fundamental aspect of data interchange (an appropriate bias for video media spaces but not necessarily all ubicomp systems). Hong’s model offers a different, but comparable perspective: data observers include not only intended system users, but also malicious eavesdroppers or third parties who might incidentally be given the power to observe the data, such as equipment owner/administrators or other stakeholders in the system. It is worthwhile to enrich my vocabulary for analysis with these concepts borrowed from Hong’s framework. There is also an opportunity for Hong’s framework to be enriched with the distinction from my vocabulary: it is not clear where unintended data sharers fit in Hong’s analysis and including them would make Hong’s framework even more ideally suited to surveillance applications, for example.

Another useful distinction between Hong's framework and my vocabulary-based approach is that Hong's framework offers a very good deconstruction of **accountability**. While on the one hand we can consider accountability in the sense of auditing of data use or misuse (as in the case of the vocabulary) Hong's framework also addresses accountability in information capture, looking at distinctions such as **consent** (opt in or opt out) and **notice** (push or pull, one-time-only or continuous). These distinctions are absent from the vocabulary because in video media spaces the only opt-out option is complete **withdrawal** and because of the network architecture used most data is pushed continuously. Finally, while the case studies used to demonstrate the analysis framework mention over-monitoring as a problem in the systems revealed through the analysis, the questions in the framework do not specifically target this kind of problem. In my vocabulary-based them, over-monitoring would be termed scrutiny.

**Mechanics of privacy:** Naturally, in considering confidentiality, Hong's analysis framework concerns itself primarily with the regulation of the **disclosure boundary** and the **tension** inherent in social intercourse therein. However, they also concern themselves with the use of **temporal boundaries** for confidentiality, and how **identity boundaries** can be affected in the regulation of confidentiality. While **spatial boundaries** and their role in privacy regulation are almost invariably affected by the kinds of context-aware ubiquitous systems they propose (e.g., location-aware services) this boundary is not the subject of any specific questions in their framework.

All of the central characteristics of the privacy process presented in the descriptive vocabulary are at least envisioned by Hong et al as they phrase the questions in their framework. They envision confidentiality as a **cooperative** process that is predicated on both **control** and **choice**, and that it is exercised in a dialectic. They particularly note that **normalisation** of privacy **expectations** and habits. While they recognise that privacy regulation is highly **dynamic**, they balance their approach with the pragmatic view that default application behaviour must be singled out for special analysis. Where Hong's analysis method and my analysis methods depart, however, is that in my vocabulary-based approach the analyst is directed to specifically question if the object of analysis supports dynamic, cooperative, and dialectic regulation of privacy whereas Hong's method does not.

**Violations** are given extensive treatment in Hong's analysis framework: in fact, a goal for using Hong's framework is to identify and enumerate the opportunities for violations to occur with the system under analysis. As with my descriptive theory of privacy, **risk** in Hong's model includes not just the **possibility** of a violation, but also the **probability** that it will arise and the **severity** of harms that would come about if it the violation actually occurred. In fact, the likelihood and damage dimensions of risk are given top-level handling in their risk management model. Hong's risk management model also includes a concept absent from my vocabulary: the cost of implementing a particular safeguard, be it detective (e.g., **policing**), preventative (e.g., a **constraint** on use) or restorative (e.g., **reprimand** or **repair**). Another immediately related concept absent from both my vocabulary and Hong's risk management model is the efficacy of a safeguard.

Most of the **behavioural** and **cognitive** phenomena discussed in my descriptive theory of privacy and listed in my analysis vocabulary are also assumed in Hong's analysis framework and, to varying degrees, revealed by answering the questions in their framework. They examine: **genres of disclosure** as normalised **expectations** regarding data sharing; policing; reprimand and repair (called "recourse" and "recovery"); the risk/reward trade-off, which becomes a **cost/risk/reward trade-off** in their risk management model; the **disclosure boundary tension**; and, the use of **explicit signifiers** for privacy **needs** and **desire**. Their framework considers only the role of intentional control to operate safeguards and hence it would be difficult to understand the operation of, say, the autonomous control features in Neustaedter & Greenberg's HOME MEDIA SPACE (previous section) with Hong's method alone. The phenomena that are omitted in Hong's framework are **self-appropriation**, **disinformation**, **reserve**, and **implicit signifiers**. These phenomena play an important role in the privacy of video media spaces, hence their inclusion in my analysis framework. Their omission in Hong's framework may not represent a deficiency in Hong's framework as much as indicate that these phenomena may deal more closely with **solitude** and **autonomy** or might be rather specific to video media spaces and their ilk.

Hong's framework directly questions **trust** as a component of relationships; oddly, this concept has been included in the descriptive theory but was omitted in the table of vocabulary terms used for analysis. The reason for its de-emphasis in the analysis vocabulary is precisely

because video media spaces are assumed to be used among people who trust each other, and so there's little need to question it; this assumption of course does not hold for the analysis of the general class of ubicomp systems to which Hong's framework is geared. The descriptive theory is ultimately a lot richer than any table of vocabulary terms and it is hard to come of with a single all-satisfying list of terms. For example, along with trust, the descriptive theory of privacy indicates that reciprocal obligations and mutually shared risks and rewards are additional cognitive phenomena that support privacy regulation.

Compared to these cognitive phenomena, much fewer of the **environmental phenomena** which support privacy discussed in my descriptive theory are questioned or assumed in Hong's framework. Hong's framework is limited to analysis of **choices** available for regulating privacy and control feed-through cues as part of the risk management model. These represent explicit mechanisms for privacy regulation. The vocabulary used in my analysis enumerates many other mechanisms: for example, **situated action, transitions, constraints, and embodiments**. These features of the environment are used as part of more implicit mechanisms for regulating privacy. These are also features that are extremely important for mediated interaction, and consequently, are extremely important to privacy in video media spaces. Their exclusion from Hong's framework may indicate they are not sufficiently general to all ubiquitous computing systems and scenarios to merit inclusion in a framework like Hong's. This may suggest that the basic framework put for by Hong could be tailored for more narrow categories of systems with questions and concepts more specific to each.

As stated in the first partition step of this analysis, Hong's analysis framework and method deal with issues peculiar to **computers and privacy**. First, Hong's framework does not provide an ontology of support methods, asking the analyst to classify each countermeasure therein. Instead, the risk analysis model gets the analyst to determine how—given all of the support methods in place in the system—unwanted information disclosure occurs. My analysis vocabulary offers a first-cut of such an ontology and despite any caveats or limitations of it, even a first-cut ontology becomes a useful aid when performing the kinds of analysis directed in Hong's framework. The support methods deconstruction in my vocabulary serves as a checklist of sorts for determining typically suspect aspects of a system when

unintended disclosure occurs. Thus, this deconstruction aids the analyst in answering Hong's question about how unwanted disclosure takes place.

Taking things one step further, Hong's framework and method itself may be considered a tool for supporting privacy in computers because it may be used to identify privacy risks and facilitate systematic decision making about **countermeasures** to risks found. Thus, we see an omission from the vocabulary used in my analysis method: in the deconstruction of various computer privacy support methods, frameworks like Hong's (or even my own method used in this analysis) are not explicitly mentioned, yet ought to be. Frameworks may be broad enough to include toolkits, guidelines, and methods for system inquiry, analysis, evaluation, and decision making.

Hong's framework is specifically geared to examine one class of important privacy problems with computers: unwanted disclosure of personal information. Hong mentions three root causes for this problem: accidents, i.e., slip-type errors (Reason, 1980), misunderstandings, and maliciousness. Comparing this to my analysis vocabulary, we see that Hong's risk management model is geared towards identifying specific kinds of inadvertent confidentiality infractions and specific kinds of deliberate confidentiality abuses. There are also four more critical problem areas for computers and privacy that Hong's framework addresses.

**Fitness of default settings for configurable system behaviour**—this issue applies generally, even to VMS. Although is mentioned in my descriptive theory of privacy, it is omitted as an oversight in my privacy vocabulary.

**Cost of developing and deploying countermeasures**—this issue applies generally, even to VMS. Yet, both my theory and vocabulary consider only cost in terms of the cognitive and physical effort in using a countermeasure, not the more holistic notion of cost put forth here by Hong.

**Control and feedback for data sharers**—the characteristics they specifically mention with which control and feedback may be examined are simplicity and understandability. My privacy vocabulary deconstructs these attributes into believability, social naturalness, utility, effort, and granularity. All of the attributes given in my deconstruction apply generally.

**Privacy policy**—specifically the contents thereof and the means by which it is disseminated to data sharers. Privacy policies are important in a broad class of ubiquitous computing systems in which a second or third party is an unfamiliar enterprise, or when such natural constraints promoting cooperative preservation of privacy (e.g., **trust**, **mutual harms**, or **reciprocal obligations**) are missing. It is not so important for privacy policies to be explicitly formulated or communicated between intimate collaborators. Consequently, they are largely irrelevant to a VMS, and are unsurprisingly omitted from my analysis vocabulary. Yet, privacy policies are important tools for supporting privacy in the broader class of privacy-sensitive systems and it would be worthwhile to include them in the support methods section of my analysis vocabulary.

Hong's framework omits two groups of problems specifically mentioned in my analysis vocabulary. First, users' and non-users' apprehension of the system's handling of privacy and resentment towards the system and its designers and advocates is single out in my privacy theory and vocabulary because these have been demonstrated consistently to be important privacy problems in video media spaces. Although they seem very general, I do not know if **apprehension** and **resentment** are as important privacy issues for general ubicomp systems as they are for VMS. The second group of problems omitted in Hong's framework include role conflict and the "four 'D's'". These problems are more peculiar to autonomy and remote **presence**, **embodiment**, and **interactivity**; Hong's framework does not itself deal with these areas.

### 10.2.3 Reveal

The principle assumption in Hong's framework is not hidden at all. It is a deliberate, explicitly stated constraint on the scope of privacy limiting it to **confidentiality**. There is an additional narrowing of domain to that of ubiquitous computing applications. Both these influence the output of analysis, specifically the kinds of problems and risks identified and the kinds of countermeasures considered. Similar comments may be made about my analysis vocabulary, which gives weight to problems specific to video media spaces but not necessarily endemic to all kinds of systems. Also, my analysis vocabulary omits some things that are irrelevant to video media space use between intimate collaborators but are highly relevant to ubiquitous

data sharing systems used between strangers. Thus, instead of revealing hidden assumptions particular to Hong's framework, in this subsection I wish to relate Hong's analysis method and my own analysis method so as to get an understanding of how each may improve upon the other. My concern is how can we arrive at a better analysis method and materials. To understand this, I will look at in this subsection the differences between what kinds of analysis can be done well and easily with Hong's framework and of these which can or cannot be currently done with my vocabulary based approach, and vice versa.

Hong's query-based analysis method permits the analyst to identify opportunities in which undesirable information disclosure may occur and how these incidents happen. Hong's framework also includes a risk management model and query-based method for making design choices about which confidentiality risks to counter first and how to counter them. Like Hong's analysis method, my vocabulary-based analysis method also permits description of confidentiality risks and the means by which they arise. I have nothing comparable to Hong's risk management model or method, although there is a definite need for them.

Hong's analysis method will identify cases of unwanted disclosure that involve:

— access by unintended individuals (persons with deliberate malicious intent or those who incidentally acquire access to the information as a consequence of the way it is captured or transmitted)

-or-

— access by an intended individual in an deliberate yet unintended manner, such as scrutiny

-and-

— the unwanted disclosure is clearly peripheral to the value proposition for disclosure in the first place.

My vocabulary-based method finds, by demonstration in the case studies given in the chapters in this act, unwanted disclosures in *at least* these conditions. In addition to these circumstances, my vocabulary-based method leads analysts to identify privacy risks that occur even when the expected users of the system use it in the expected manner, such as inattention slips in self-appropriation. Furthermore, because my vocabulary-based analysis method

assumes a much broader, more holistic notion definition of privacy, it permits the identification and description of solitude and autonomy risks, in addition to those that involve confidentiality. For example, my privacy vocabulary will identify as a problem a case when the system fails to support disinformation as a means for ameliorating role conflict in desituated interaction. As another example, imagine an in-hospital video-based patient monitoring system that surveils the patient in his or her room and archives the video for a period of no more than 18 hours. Physicians and nurses use the video archive to ensure that patients get out of bed and move about as per physicians' instructions. Such activity is often a critical component of speedy recovery. My vocabulary-based method, when applied to such a system, would identify apprehension, self-appropriation and refuge issues even though notice was given and consent collected, and even if capture and transmission are well fortified against access by unauthorised individuals.

This relates to an important observation about the relationship between Hong's query-based analysis method and my vocabulary-based analysis method. Mine offers words as anchor-points for discussion; Hong's offers questions. Of course, each term in my vocabulary may be easily worked into a question. For example, a depth-first descent of the vocabulary tree gives terms "solitude", "physical dimensions", "interpersonal distance", and "isolation to crowding". We could phrase this as a question: "At what ranges of interpersonal distances, from isolation to crowding, does the system simulate or approximate and permit people to arrange themselves?" In fact, when using the vocabulary, the analyst will often formulate these questions on the fly during analysis as a means of utilising the vocabulary.

Conversely, the questions in Hong's framework can be reduced to the core vocabulary used in each. Take, for example, the first group of questions asked by Hong. These may be reduced to the following vocabulary: social and organizational context—users—data sharers, data observers. The process may be repeated for the entire set of questions in Hong's framework. Given that the basis of the two methods—questions or vocabularies—are mutually interchangeable, are the only differences between the two frameworks, therefore, in the content of the materials used? Certainly, as explained before, scope is one issue: privacy in my vocabulary includes a broader set of phenomena and so naturally the vocabulary is larger. Although a lot of questions would need to be added, it is possible to extend Hong's framework

to cover as many dimensions of privacy. (It's also important to point out that there are a handful of concepts in Hong's framework that would need to be added to my vocabulary in order to achieve true parity, and there are several natural extensions to these concepts presently absent from either framework that should be added, as well.) Granularity is also an issue: in this case study I've described places where monolithic terms used in Hong's framework are better decomposed in my privacy vocabulary.

Finally, the dimensions along which terms are decomposed become an issue. My descriptive theory of privacy asserts that privacy is a union of many different modalities of control, social behaviours, cognitive phenomena, and environmental factors. They are bound up into the one thing privacy, even though I tease this one thing apart into solitude, confidentiality, and autonomy. Hong's framework splits things out at the top level into social and organisation context and technology. Yet, an alternate organisation might split things out into "users and their relationships", "motivations and motives", "mechanics of data collection and transfer", "data characteristics". Indeed, how analysis items are organised influences the output of analysis, as seen earlier in this case study and there is no single decomposition that is universally satisfactory for all analysts and all objects of analysis.

There are two subtle and important differences between the methods. First, the query-based approach constrains the analysis process more than the vocabulary based approach. In a sense, the vocabulary terms are like extremely open-ended questions. The way questions are phrased determines both the answer the analyst finds and the process by which she or he finds it. Constraints guide analysis and this guidance is especially appreciated when first starting analysis. A sheet of paper filled with vocabulary terms and little idea on how they should be used can be quite daunting at the start. Later on, however, as the analyst is more familiar with the vocabulary terms and the object of analysis, these constraints can hinder broad and multifaceted description of the object of analysis.

The second subtle difference concerns the kind of output that each method generates. My vocabulary-based analysis method seeks to output a concise and systematic description of the cases in which a system may be expected to support privacy. Hong's method seeks to enumerate the normal-use situations in which confidentiality cannot be appropriately regulated. Although the stances are complimentary, I would argue that the vocabulary-based approach is

superior because it is relatively trivial to imagine situations in which confidentiality cannot be successfully regulated given the description produced with my method, but it is extremely difficult to generalise about the successes given only the description of specific circumstances of failure produced with Hong's method.

#### 10.2.4 Summarise

In this case study I have analysed Hong's privacy risk models method for identifying and prioritising privacy risks in ubicomp systems. I have shown that the model Hong et al develop constrains notions of privacy to only confidentiality as a modality for privacy regulation. I have shown that their model further constrains risk to unwanted disclosure involving system access by unintended people or unintended use by expected users.

I have also shown that these are limitations in the materials used, not the method of analysis. I have shown that the query-based analysis method is interchangeable with my vocabulary-based method. I have explained how the materials used in privacy risks models method can be enriched with concepts taken from my vocabulary to encompass broader notions of privacy and risk. I have listed the few places where my descriptive theory and privacy vocabulary may be enriched as well by borrowing from Hong's framework. I have shown how my privacy theory and analysis method lacks an equivalent to Hong's privacy risk management method for prioritising risks to be countered, and how Hong's method may be enriched with concepts taken from my descriptive theory.

It is clear by this point how the two analysis methods may be combined to arrive at a better method. First, we utilise the comprehensive descriptive theory of privacy I develop in this thesis and the vocabulary that emerges from it as the basis for analysis. Second, we utilise the structure afforded by Hong's query-based method as a means of easing into the analysis, and allow the analyst the freedom to move fluidly between the query-based and vocabulary-based methods as fits the object of analysis. This revised method will better aid identification of a greater variety of privacy risks in a wider selection of subjects (not just systems, but techniques, evaluations, and competing frameworks) than would otherwise be obtained with either one of the incorporated analysis methods alone. Finally, we augment this kind of analysis

with the privacy risk management model developed by Hong et al to facilitate prioritisation of the risks identified and chart courses for future iterative refinement or exploration.

### 10.3 Conclusion to the 2<sup>nd</sup> set of case studies

This chapter presented the final two case studies that illustrate the application of the analysis/description method developed in this act and in doing so brings the act to a close.

**Case study #4** looked outside this thesis, to Neustaedter & Greenberg's HOME MEDIA SPACE. This case study's analysis revealed that the autonomous capture control features in HMS are designed to reduce the risk of inattention slips in self-appropriation. The theory of privacy developed in Act II describes many other kinds of privacy problems not addressed in HMS.

**Case study #5** examined Hong's privacy risk models framework for analysing privacy problems in ubiquitous computing systems. This case study revealed that Hong's framework identifies cases of unwanted disclosure when information is accessed by unintended individuals or in unintended ways, or when the disclosed information is clearly peripheral to the value proposition of the system.

### 10.4 Reflecting back on all of Act III

The analysis method described in Chapter 8 was motivated by the observation that different kinds of knowledge inform design in very different ways. A descriptive theory supports description. Systematic description identifies those aspects of privacy are supported, but also draws the analyst's attention to what has been omitted and, with further reflection, what hidden assumptions have been made about the nature of privacy and what constitutes a 'privacy problem.' The method uses the vocabulary developed in Act II as anchor-points in each of the four analysis steps: partition, describe, reveal, and summarise.

Chapter 9 presents the first three case studies that illustrate the use of the analysis method to revisit work that I presented in Act I and see it in the light of the knowledge gathered in order to produce Act II. Chapter 10 presented two additional case studies that

show the method and the concepts in the vocabulary being applied to work others have done independent of

The case studies in Chapter 9 and 10 demonstrated that a method to systematically analyse and describe privacy phenomena supports comprehension and communication about our systems' effects on privacy. Basing the method on a comprehensive descriptive theory of privacy makes it possible for analysts to shed light on hidden assumptions and omissions in their notions of privacy that significantly affect their designs. In bringing this out, this act has illustrated the true power of a descriptive theory of privacy for informing the design of privacy-preserving video media spaces.

## Chapter 11—Conclusions

In this thesis, I have explored the notion of privacy in the context of the design of video media spaces. The over-arching goal of this thesis has been to inform the design and implementation of video media spaces that support privacy and enrich distributed collaboration. Out of this larger goal, I identified four specific goals to address in this thesis.

In this concluding chapter, I will review the progress I have made on these four problems/goals/deliverables (Section 11.1). I will identify the major and minor contributions and draw particular attention to the key lessons learned from my research and the impact it will have on the way privacy is conceptualised within the context of video media space design (Section 11.2). I will describe new ways of applying the results of this research, including potential alternate methods for using the results and how they can apply to domains other than video media spaces (Section 11.3).

Finally, I go back to the original research questions #1 and #2 and project a plan of further research beyond this thesis to provide theoretical knowledge and applied tools for designing, building, and evaluating video media spaces for privacy and distributed collaboration (Section 11.4).

## 11.1 Progress on thesis problems, goals, and deliverables

### 11.1.1 Problem #1: Rapid prototyping toolkit

**Thesis Problem #1:** It is hard to rapidly develop video media spaces because the programmatic interfaces for multimedia are complicated and require considerable programmer effort and expertise.

**Thesis Goal #1:** Develop a toolkit to support the rapid prototyping of video media spaces and the distortion filtration method for preserving privacy therein.

**Status of Goal #1:** Completed. See the description of the Collabrary toolkit in Chapter 3.

I chose to attack this problem—concerning technology for implementing privacy-preserving video media spaces—because much of what had already been learned about video media spaces has been by researchers designing, implementing, and living with the technology. I wanted to be able to rapidly prototype privacy preservation features that utilise video analysis and manipulation within a working video media space system because I felt this was the richest and most valid way of evaluating these features.

To achieve goal #1, I produced the COLLABRARY toolkit, which I described in Chapter 3. In various examples presented in that chapter, I illustrated how the COLLABRARY provides a lightweight and accessible API for multimedia capture, transmission, presentation, analysis, and manipulation. This API was packaged as a Microsoft COM object library that can be used with existing popular rapid application development environments (e.g., Microsoft Visual Basic, C#) to construct a working  $n$ -way video media space systems that afford reliability and performance sufficient for everyday use by resilient users in robust situations. This toolkit has seen active use by myself and other researchers and has been a critical technology underlying numerous published research systems and MSc theses. It has also been used to produce the materials used in the distortion filtration studies presented in Chapter 4 of this thesis.

### 11.1.2 Problem #2: Distortion filtration evaluation

**Thesis Problem #2:** It is widely suspected that distortion filtration may be useful for mitigating privacy issues in video media spaces but its usefulness has not been rigorously evaluated and there is no guidance as to how much filtration is ideal.

**Thesis Goal #2:** Determine if it is possible to use the distortion filtration technique to strike a balance between awareness and privacy in a video media space. If it is possible, determine at which levels a balance can be reached.

**Status of Goal #2:** Completed. See Chapter 4's descriptions of the evaluation studies performed in collaboration with Chris Edwards, Carman Neustaedter, and Saul Greenberg. The results of these studies clearly show that these techniques are unreliable in high-risk use cases.

I chose to attack this problem because at the time distortion filtration looked like a promising way to balance privacy and awareness—conceived of, at the time, as opposing interests—and I was searching for guidance on how best to implement them in a video media space prototype. The techniques were widely suggested as possible means for preserving privacy and I was unsatisfied with initial explorations of the technique (Zhao & Stasko, 1998). The innovation I chose in my examination concerned the relationship between image fidelity and awareness and privacy.

To achieve goal #2, I performed a first study of the blur and pixelize distortion filters in collaboration with Chris Edwards and Saul Greenberg. This study showed that the blur filter sufficed for mundane scenarios which posed little privacy risk. A second study conducted in collaboration with Carman Neustaedter and Saul Greenberg showed that the filter does not provide an adequate balance between awareness and privacy in the high-risk home-based telecommuting scenarios tested.

These studies clearly show that this technique will suffice for office colleagues in low-risk scenarios but will not be robust enough to preserve privacy for home-based colleagues in high-risk scenarios.

### 11.1.3 Problem #3: Descriptive theory of privacy in video media space design

**Thesis Problem #3:** There is no comprehensive vocabulary of privacy terms—one that integrates conceptions and theories of privacy from many disciplines—to support unambiguous description of how privacy is affected by video media space design and use.

**Thesis Goal #3:** Integrate privacy theories and observations from many disciplines of scientific inquiry to produce a vocabulary for describing privacy and a video media space's effect on it in an unambiguous and comprehensive manner, accounting for at least the privacy issues reported in previous literature.

**Status of Goal #3:** Completed. See the privacy theory described in Chapters 5~7 (Act II).

I chose to attack this problem because the theoretical knowledge of the nature of privacy available to me at the time was insufficient: it could not be applied to the design of privacy-preserving video media spaces. The perspective of assembling a descriptive theory of privacy came only after I immersed myself in social, psychological, and CSCW literature on privacy. It proved hard to apply these theories to video media spaces because the scope of privacy and the language used to talk about it differs with each researcher and discipline.

To achieve goal #3, I synthesised the observations of numerous privacy researchers in CSCW and added to this rich starting point perspectives of privacy borrowed from other disciplines. My main sources were Altman (1975); Bellotti (1998); Gavison (1980); Goffman, (1959); Grudin, (2001); Nardi et al (1997); Palen & Dourish (2003); and Schwartz (1968). I integrated and expanded upon these authors' theories and developed a tri-partite theory of privacy which decomposes it into solitude, confidentiality and autonomy modalities of control.

Out of these core themes, a powerful language emerges that can be used to disambiguate the many interrelated and subtle meanings of “privacy.”

#### 11.1.4 Problem #4: Methods for describing privacy

**Thesis Problem #4:** There is no systematic method for applying the concepts in the privacy vocabulary to understand and inform the design of privacy-preserving video media spaces.

**Thesis Goal #4:** Develop a systematic method of applying the terms in the privacy vocabulary to describe and analyse the effect of a video media space’s design and use on privacy.

**Status of Goal #4:** Completed. See the description of the method in Chapter 8, and example applications in Chapters 9 and 10.

I chose to attack this problem when it became clear that I needed to show how the descriptive theory of privacy I produced informs design. I had hoped that the deeper understanding of privacy gained through the privacy vocabulary would guide directly both to the design of privacy safeguards in a video media space, and to metrics for the user-centred evaluation of such safeguards. I came to see that this is the role of prescriptive theory. While the development of the prescriptive theory of privacy in video media spaces is part of the long-term goal of building a privacy-preserving video media space, such a theory is beyond the scope of this thesis.

To achieve goal #4, I focused my attention on exploring and illustrating how the descriptive theory of privacy informs—but does not prescribe—the design of privacy-preserving video media spaces. In Chapters 8~10, I provided a method based on my descriptive theory’s vocabulary that supports systematic analysis and discourse about the merits and demerits of the privacy support in a media space. I followed with example illustrations of the method, which themselves reveal important observations about privacy as it has already been conceived in various aspects of video media space design.

## 11.2 Thesis contributions

In this thesis I contribute a significant body of theoretical and analytical knowledge concerning privacy in the context of the design and use of video media spaces for intimate collaborators. This knowledge forms the basis of the answers to research questions #1 and #2 that have guided this thesis. Below I separate out the major and minor components of this contribution.

### 11.2.1 Major contributions

- A **comprehensive descriptive theory of privacy with a disambiguating vocabulary** that informs and sustains holistic analysis and discourse about privacy in the context of video media spaces (thesis deliverable #3/research question #2). It teaches us—researchers and practitioners—that the many facets of privacy—e.g., distraction, interpersonal distance, ambiguity, access control, territoriality, liberty, self-appropriation, and pseudonymity, to name a few—are all synergistically interrelated. This holistic perspective **fundamentally changes our understanding of the nature of privacy** and its role in people’s daily lives. This contribution, I feel, will have the most significant impact on CSCW research of all the contributions offered in this thesis.
- A **method to systematically analyse and describe privacy phenomena** in the context of video media spaces (thesis deliverable #4/research question #2). This supports our efforts to comprehend our systems’ effects on privacy and to share our understanding with others. In particular, this method **illuminates hidden assumptions and omissions in our notions of privacy** that significantly affect privacy in our designs. I feel this contribution has the potential to significantly impact the way CSCW systems are designed, analysed, and evaluated.
- A **toolkit for rapidly prototyping working system video media spaces** that incorporate multimedia analysis and manipulation techniques for preserving privacy (thesis deliverable #1/research question #1). This toolkit gives people a platform for exploring video media spaces and different effects to preserve privacy. It makes the arduous task of multimedia groupware programming so effortless that it **changes our ideas about what we as CSCW researchers can accomplish** when we are building media spaces. This contribution is

shorter-lived than the previous two but still yields a significant—albeit considerably different—impact.

- A **method for evaluating the blur and pixelize distortion filters and results** which conclusively show that they are not robust enough for use in high-risk scenarios (thesis deliverable #2/research question #1). This contribution **prompts reconsideration of the efficacy of existing approaches**. While I feel this is the narrowest of the four major contributions I offer in this thesis, I still feel that it has far-ranging impact because of the pervasiveness with which these techniques are employed.

### 11.2.2 Minor contributions

- A comprehensive survey of casual interactions and informal awareness, technology to support them, and the lessons that have been learned from prior video media space design and use.

Also, a variety of contributions stem from the case studies of Act II.

- An analysis of the distortion filtration technique that identifies the essential merit of the technique and four critical conditions that govern the successful realisation of this merit.
- An analysis of the method used in the evaluation studies (Chapter 4) that identifies notable omissions in the aspects of privacy evaluated.
- An analysis of the features of the COLLABRARY toolkit that identifies opportunities to expand it beyond its present rudimentary level of support for prototyping privacy preservation techniques.
- An analysis of Neustaedter & Greenberg's HOME MEDIA SPACE prototype that identifies the kind of privacy problems the features in the prototype will counter and the circumstances in which they will be successfully countered.
- An analysis of Hong et al (2004) privacy risk models that shows how it and my vocabulary-based analysis method can be integrated to gain the benefits of both.

## 11.3 Reapplying the results

In this thesis I have used video media spaces for intimate collaborators as the major constraint on the scope of my examination of privacy. However, the perspective adopted in this thesis is that privacy phenomena necessarily pervade all aspects of human social life and so it is not surprising that most of what has been discussed about privacy within the context of video media spaces appears to naturally apply to other kinds of collaborative computing systems and scenarios. There are several application domains that stand out in particular as being places where my descriptive theory of privacy may be readily applied, perhaps amended further with domain-specific elaborations.

### 11.3.1 Group collaboration tools

Key examples: Instant messengers, email archive analysis and visualisation, meeting capture.

Group communication tools—of which video media spaces are an example—are an application domain which seems broadly amenable to analysis with my descriptive theory of privacy.

Instant messengers (IM) are a good example, in particular, because plausible deniability is an important concept in the privacy theory that comes directly out of observations of IM use (Nardi, Whittaker & Bradner, 2000). Instant messenger clients are increasingly adopting video media space-like functionality, such as display pictures and webcam integration, prompting self-appropriation issues. Instant messenger software is slowly gaining uptake in enterprises, further exacerbating self-appropriation issues, but also prompting confidentiality and personal liberty concerns. My descriptive theory of privacy seems readily applicable to describing opportunities for privacy concerns in these kinds of scenarios.

Email is another natural area for reapplication of the theory. There are now emerging a new class of tools for analysing large archives of email messages. For example, consider Google gmail: the contents of one's email is analysed on Google's powerful supercomputing cluster to display targeted advertisements. This raises both confidentiality and solitude concerns. New tools to visualise relationships through email make explicit highly sensitive information that was unknown or implicit, prompting confidentiality concerns. For example,

consider a tool such as ContactMap (Whittaker et al, 2004): it derives an “overall measure of importance” algorithmically from email exchanges. Now consider a manager looking at the display that includes the members of the team she manages. This display may reveal misleading statistics (apparently high fidelity information) about the relationships she has with her team and this in turn raises confidentiality and autonomy concerns for everyone in the team. My descriptive theory of privacy may be useful for disentangling these kinds of privacy problems and helping designers design in preparation for these kinds of effects.

Finally, meeting capture systems—particularly those that capture and archive spontaneous and serendipitous casual interactions—raise numerous confidentiality and self-appropriation concerns. Offline access to meeting transcripts results in problems like decontextualisation and desituation of action that can be unambiguously described using my privacy theory. These may point designers towards specific use scenarios in which their system’s privacy features need to be evaluated.

### 11.3.2 Pervasive computing infrastructure and applications

Key examples: Targeted advertising, Intel Personal Server, automated security threat analysis (profiling).

It seems like there are even more opportunities for privacy violations to occur in other kinds of ubiquitous computing systems than there are in video media spaces. There are natural synergies between my privacy theory and work in the broader ubicomp domain, as deeply underscored in analysis of Hong’s privacy risk models framework targeted ubicomp (Chapter 10).

By definition, there is a tension between solitude and ubiquity. For example, consider the hypothesised pervasive use of retinal scanning for targeted advertising (a kind of context-aware service) depicted in the feature film *Minority Report* (20<sup>th</sup> Century Fox & Dreamworks Pictures, 2002). My descriptive theory could be used to deconstruct and describe the distraction issues prevalent in this kind of dystopian future. In facilitating this discourse, my privacy theory might serve as the starting point for exploration of compensatory measures to, say, better balance risk and reward.

There is a natural tension with confidentiality because many ubicomp systems rely on coordination and communication between personal computing systems and shared infrastructure. For example, consider the Personal Server (Pering et al, 2003): a small data storage device that relies on displays situated in the nearby environment for much of its user interface. My descriptive theory of privacy might point out to the designers of such a system the subtle role of territoriality in ascertaining the fidelity of information to be displayed via the borrowed equipment.

There is a natural tension with autonomy because some of the proposed uses for ubiquitous computing include surveillance and monitoring. Consider, for example, the use of facial recognition technology to help automate screening of attendees at large spectator events (Woodward, 2001). Inaccuracies in the system may cause an individual considerable public grief by repeatedly misclassifying him as a criminal or terrorist. There are also numerous confidentiality concerns regarding which branches of government and which levels of authority have access to such identifying information. My descriptive theory privacy permits designers to talk about both kinds of issues in conjunction with one another using a single unified theory that reveals the strong interrelationships between the two.

### 11.3.3 Regulatory measures for protecting individual rights and freedoms

Key examples: Corporate policies regarding distributed teamwork, regulations to protect consumer privacy.

On numerous occasions I have had the opportunity to discuss my research with people from outside computer science who have an interest in privacy. Because I have chosen to give my descriptive theory of privacy in video media spaces such a broad foundation, it appeals to many people, from museum curators concerned about the confidentiality of their benefactors, to public officers responsible for evaluating matters of public procedure and policy.

One natural area to which my work on privacy can be extended is on the development of corporate policies regarding distributed teamwork. For example, to help support confidentiality, a company may make it policy that one IM client be used for informal communications within the company while a second client (with, perhaps, mandatory archival, spelling/grammar/trade secret checks) be used for more formal communications with others

outside the company. As another example, to help support autonomy, a company may set as its policy a rule that any email not marked high priority sent after the close of the business day is withheld and not delivered until the start of the next business day. My descriptive theory of privacy might help identify areas like these where an organisation may want to establish or alter norms of technology use to achieve its goals.

Another area where the privacy theory might be informative is the development of regulatory legislation to protect people's rights to privacy. For example, an important message embedded in my privacy theory is that it is important to question the efficacy of notice and consent as means of supporting confidentiality. This is because people may not have the knowledge they need to make sound privacy decisions at the time notice is given and consent obtained. This observation could have enormous potential impact on the design of policy and law regulating corporate use of consumer data.

An important caveat to the reapplication of my research to the problem of design regulatory policy for privacy is that while my privacy theory recognises that there are macrosociological privacy problems that are entirely different from the interpersonal privacy issues that arise in video media spaces, my theory does not delve into this other set of problems at all. The value of my theory to this domain is that it provides a vocabulary that embodies the important interrelationships between solitude, confidentiality, and autonomy that must nonetheless be recognised and respected in social policy. This vocabulary can serve as the foundation of a language for the design of such policies, as well as the design of video media spaces.

## 11.4 Future work

In this section, I describe ways of expanding upon the work presented in this thesis. I have identified two near term action items which directly improve upon the work in this thesis. These improvements come directly from the thesis itself. I have also identified a programme of long term research that lead towards the development of software technology, prescriptive design theory, and empirical evaluation methods geared towards the design, development, and evaluation of privacy preserving video media spaces.

### 11.4.1 Near term: Improving on the results

- **COLLABRARY improvements.** In Chapters 3 and 9 I enumerated ways in which the COLLABRARY can be improved to better support rapid prototyping of video media spaces including greater scalability and reliability, integrated encryption, access control, better learning materials, etc. Future work has already begun on these objectives.
- **Integrating privacy risk models and my vocabulary-based analysis method.** In Chapter 10, I outlined steps to enrich my privacy theory and the vocabulary-based analysis method by integrating them with Hong's privacy risk models method. In particular, Hong's method offers good deconstructions of users, accountability, persistence, transmission, authenticity, and temporal *vs.* content precision/accuracy. These concepts can be readily integrated into my existing privacy theory. Over a slightly longer term, the methods themselves can be integrated and guidance offered about how and when to make the transition between the query-centred and vocabulary-centred perspectives.
- **Disseminating the theory and analysis method.** The analysis method given in Act III has been used by exactly one person: me. Important questions about the generalisability of the theory and the analysis method discussed in the introduction to Chapter 8 demand that the theory and method be refined based on others' experiences learning and applying them. As analysis procedure, of course, it is inescapable that the results obtained will vary according to the insightfulness of the analyst applying the method. This is a methodology caveat shared with, for example, discount usability inspection methods like Heuristic Evaluation (Niesen, 1993). To this end, it may be possible to borrow coping strategies employed with the Heuristic Evaluation method such as using multiple analysts working independently in parallel or by combining the analysis with observations in field testing and user testing, when appropriate. As future work, it will be necessary to get other practitioners to use the method and provide feedback on it to subsequently refine it, perhaps provide better instruction or more detailed guidance in the Reveal step, for instance.
- **Refining the vocabulary list.** I assembled the list of vocabulary terms that guide the analysis method in Act III by proceeding sequentially through the theory chapters in Act II, noting down keywords that I had set in boldface type, plus those that I felt were extremely

important. I subsequently “massaged” the list, weeding out duplicates and organising the keywords into compact hierarchies. This approach makes use of my expert understanding of the theory to make judgements about which terms to include. Others, reading the theory chapters, might very well come up with lists of their own which differ in content or organisation. This leaves open the question of if this is an ideal set of keywords or organisation. It can also be debated if a single list of vocabulary terms is ideal. One strategy proposed by Dr. Sheelagh Carpendale, University of Calgary (personal communication) is to split the vocabulary out into several smaller lists. One list might contain be vocabulary universally applicable across all applications and domains to be applied first. The remaining lists could be mutually independent sub-modules of vocabulary specialised to particular applications and domains or lists that further elaborate on the results found using the initial universal vocabulary list.

#### 11.4.2 Far term: Extending the results

In Chapter 1, I decomposed the larger goal of this thesis—informing the design of privacy-preserving video media spaces—into two research questions.

**Research Question #1:** What low-level technological factors need to be considered when building a privacy preserving video media space?

**Research Question #2:** What high-level social-psychological factors need to be considered when designing a privacy preserving video media space?

I stated that the full answer to these questions marks out a research programme that extends well beyond this thesis. In this section, I describe the kinds of research that I think will build upon the work in this thesis and bring out more of the answers to these questions. It is important to point out that I do not describe a sequential plan for future work. Instead, I describe areas for iterative exploration which will often overlap and synergistically inform each other.

- **Mechanics of privacy.** (Research question #2) In Chapter 8, I identified a large list of potential mechanics of privacy. Further expansion and condensation of this list is needed to identify and catalogue the core set of low-level behaviours and environmental properties that people use to manage privacy every day, e.g., embodied interaction, ambiguity in speech, change permeability of space to interactions, eye contact, reciprocal obligations, shared risks and rewards. These are analogous to the Mechanics of Collaboration (Gutwin & Greenberg, 2000) and may likewise serve as the foundation for heuristics for privacy evaluation (Baker, Greenberg & Gutwin, 2001).
- **Privacy effects and observable metrics for evaluation criteria.** (Research question #2) There is a need to determine which privacy phenomena to track and for these identify observable metrics which can be used as criteria for evaluations of the efficacy of the privacy safeguards in a video media space.
- **User interface software technology for privacy.** (Research question #1) My privacy vocabulary enumerates phenomena that are relevant to privacy and can serve as a checklist for the development of sensing technology and inferencing algorithms to expose these concepts to designers. Borrowing from the identification of privacy effects and observable metrics for evaluation criteria, there is a need to further develop computer tracking of privacy desires (which must also include formulae for automated hypothesis of privacy desires). My privacy theory also indicates that it will be important to develop a rich variety of modalities for interactivity that vary in attentional demands: these will be important for supporting solitude regulation. Similarly, there is a need to develop algorithms to modify the fidelity of information access that directly affect sensitivity: these will be important for supporting confidentiality regulation beyond simple access control. Finally, I expect that the development of technology to support autonomy regulation—beyond group interfaces for policing media spaces and diagnosing/preventing privacy problems as they occur—will be much more uncertain.
- **Catalogue of typical privacy problems and their remedies.** (Research question #1) As the development of privacy-preserving video media spaces progresses, it is important that solutions found successful or unsuccessful be tabulated. Such a catalogue is the foundation of practical prescriptive guidance.

- **Privacy evaluation methods and materials.** (Research questions #1 and #2) The aforementioned effects and metrics for evaluation combined with systems and safeguards worth evaluating will lead next to the iterative development of methods and materials for evaluating privacy in video media spaces. This includes observational techniques, a suite of canonical privacy scenarios/problems to be tested, example questions, heuristics, and useful analysis techniques. Much can be drawn from parallel work in ubiquitous computing and regulatory policy.
- **Ethical guidelines for designing and evaluating privacy in video media spaces.** One of the perspectives of privacy put forth in Chapter 6 is that privacy is a right. Scientific handling of it—in particular, user-centred evaluation of it or experimentation with it—fraught with ethical concerns. There is a need to develop principles that govern and guide the formation and implementation of ethically sound methods for assessing privacy and evaluating systems' effects on privacy. There are established guidelines on the ethical handling of study participants' confidential and personally identifying information. These guidelines provide a solid base to build upon and must also be expanded to deal with other privacy problems identified by descriptive theory, such as distraction, lapses in self-appropriation, and unanticipated pressures on individual autonomy.
- **Prescriptive guidance about how to design for privacy in video media spaces.** (Research questions #1 and #2) This is the culmination of all of work towards privacy preserving video media spaces described above. As an integrated, holistic discipline, it will include: comprehensive theory of privacy and design to ground designers; practical examples of good and bad privacy design to inspire designers; guidelines for good design to constrain the problem; privacy heuristics that can be applied at every stage of design so they can check their progress; ethical discount evaluation methods and field observation methods so that designers, users, and non-users can confidently understand what in a media space will support or threaten privacy.

## 11.5 Final words

I started this thesis seeking to answer the question:

### **How does one build a privacy preserving video media space?**

I took a bottom-up technology-driven approach described in Act I. The dissatisfaction I eventually came to feel with this approach prompted me to recognise that I needed to answer a different question first:

### **What does the word ‘privacy’ mean?**

I took a top-down theory-driven approach—drawing knowledge from many sources and tailored it suit the phenomena seen in video media spaces—that culminated in an original theory of privacy described in Act II. This theory transforms and elevates one’s understanding of what privacy is by shaping the way we talk about it. The question to be considered next is:

### **What does the phrase ‘privacy preserving’ mean?**

The systematic analysis and discourse method described in Act III begins the long process of unravelling this question. Though there is still much of the problem remaining to explore, I no longer feel any dissatisfaction: only excitement as to what questions this tiny thread of research will prompt next.

# Bibliography

- ACKERMAN M.S. (2000), The Intellectual Challenge of CSCW: The Gap Between Social Requirements and Technical Feasibility, in *Human-Computer Interaction*, Lawrence Erlbaum Associates, vol. 15.
- ACQUISTI, A. (2002), Protecting Privacy with Economics: Economic Incentives for Preventive Technologies in Ubiquitous Computing Environments. In *Workshop on Socially-informed Design of Privacy-enhancing Solutions in Ubiquitous Computing* at UbiComp 2002.
- ADAMS, A. (2000), Multimedia Information Changes the Whole Privacy Ballgame, in *Proceedings of Computers, Freedom, and Privacy 2000: Challenging the Assumptions*, ACM Press, pp. 25-32.
- ADLER, P., & ADLER, P. (1991) *Backboards and Blackboards*. Columbia University Press, New York, NY.
- ALLEN, R.E., (1985) *The Oxford Dictionary of Current English*. Oxford University Press.
- ALTMAN, I. (1975), *The Environment and Social Behavior: Privacy, Personal Space, Territory, Crowding*, Wadsworth Publishing Company.
- ALTMAN, I., & CHEMERS, M. (1980), *Culture and Environment*. Wadsworth Publishing Company, Stanford, CT.
- ANGIOLILLO, J.S., BLANCHARD, H.E., ISRAELSKI, E.W., & MANÉ, A. (1997), Technology Constraints of Video-Mediated Communication. In *Video-Mediated Communication*, K. Finn, A. Sellen, and S. Wilbur Eds. Lawrence Erlbaum Associates Publishers, Mahwah, NJ, pp. 51–74.
- ARENDT, M. (1958), *The Human Condition*, University of Chicago Press, Chicago, IL.
- ARGYLE, M. (1972), Non-verbal communication in human social interaction in *Non-verbal communication*, Hinde ed., Cambridge University Press, pp. 243-269.
- BAKER, K., GREENBERG, S. & GUTWIN, C. (2001) Heuristic Evaluation of Groupware Based on the Mechanics of Collaboration. In M.R. Little and L. Nigay (Eds) *Engineering for Human-Computer Interaction (8th IFIP International Conference, EHCI 2001, Toronto, Canada, May)*, *Lecture Notes in Computer Science*, Springer-Verlag, v. 2254, pp. 123-139.
- BALFANZ, D. & SIMON, D. (2000), WindowBox: A Simple Security Model for the Connected Desktop. In *Proceedings of the 4th USENIX Windows Systems Symposium* (Seattle), Advanced Computing Systems Association, pp. 37–48.
- BELLOTTI, V. (1996), What you don't know can hurt you: Privacy in collaborative computing, in *Proceedings of the HCI'96 Conference on People and Computers*, Springer-Verlaag Publishers, vol. 9, pp. 241-261.

- BELLOTTI, V. (1998), Design for Privacy in Multimedia Computing and Communications Environments, in *Technology and Privacy: The New Landscape*, Agre and Rotenberg eds., MIT Press, pp. 63-98.
- BELLOTTI, V., & SELLEN, A. (1993), Design for Privacy in Ubiquitous Computing Environments, in *Proceedings of the Third European Conference on Computer-Supported Cooperative Work (ECSCW'93)*, Kluwer Academic Publishers, Milan, pp. 77-92.
- BENFORD, S., GREENHALGH, C., BOWERS, J., SNOWDON, D., & FAHLÉN, L.E. (1995), User Embodiment in Collaborative Virtual Environments, in *Proceedings of the ACM/SIGCHI Conference on Human Factors in Computing Systems (CHI'95)*, Denver, pp. 242-249.
- BLY, S.A., HARRISON, S.R., & IRWIN, S. (1993), Media Spaces: Bringing People Together in a Video, Audio, and Computing Environment, in *Communications of the ACM*, ACM Press, vol. 3, no. 1, pp. 28-47.
- BORNING, A., & TRAVERS, M. (1991), Two approaches to casual interaction over computer and video networks, in *Proceedings of the ACM/SIGCHI Conference on Human Factors in Computing Systems*, New Orleans, pp. 13-19.
- BRADSKI, G.R. (1998), Computer Video Face Tracking for use in a Perceptual User Interface. In *Intel Technology Journal Q2'98*.
- BRIERLEY-NEWELL, P. (1995), Perspectives on privacy. In *Journal of Environmental Psychology*, 15, Academic Press, New York, NY, pp. 87-104.
- BRIERLEY-NEWELL, P. (1998), A cross-cultural comparison of privacy definitions and functions: A systems approach. In *Journal of Environmental Psychology*, 18, Academic Press, New York, NY, 357-371.
- BURKERT, H. (1998). Privacy-Enhancing Technologies: Typology, Critique, Vision. In *Technology and Privacy: The New Landscape*, P. Agre & M. Rottenberg, Eds. MIT Press, Cambridge, MA, pp. 125-142.
- BURRIDGE, R. (2004) *Shared Data Toolkit for Java Technology User Guide*, Sun Microsystems JavaSoft.
- BUXTON, W.A.S. (1997), Living in Augmented Reality: Ubiquitous Media and Reactive Environments, in *Video Mediated Communication*, Finn, Sellen, and Wilbur eds., Lawrence Erlbaum Associates, pp. 363-384.
- CADIZ, J.J, VENOLIA, G., JANCKE, G, & GUPTA, A. (2002) All ways aware: Designing and deploying an information awareness interface. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW 2002)*, New Orleans, pp. 314-323.
- CARRIERO, N. & GELERNTER, D. (1989), Linda in context. In *Communications of the ACM*, 32:4, pp. 444-458.
- CARY, M.S. (1978), The Role of Gaze in the Initiation of Conversation, in *Social Psychology*, vol. 41, no. 3, pp. 269-271.
- CLARK, H.H. AND BRENNAN, S.E. (1992), Grounding in Communication, in *Readings in computer supported cooperative work*, Baecker ed., Morgan Kaufmann Publishers, pp. 222-234.

- CLARKE, R. (1994), The digital persona and its application to data surveillance. In *The Information Society*, 10:2. Taylor and Francis, New York, NY, pp. 77–92.
- COCKBURN, A., & GREENBERG, S. (1993), Making contact: Getting the group communicating with groupware, in *Proceedings of the ACM/SIGOIS Conference on Organizational Computing Systems (COOCS'93)*, Milpitas, CA, pp. 31-40.
- COOL, C., FISH, R.S., KRAUT, R.E., & LOWERY, C.M. (1992), Iterative Design of Video Communication Systems, in *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW'92)*, Toronto, pp. 25-32.
- COUTAZ, J., BÉRARD, F., CARRAUX, E., & CROWLEY, J. (1998), Early Experiences with the mediaspace CoMedi, in *IFIP Working Conference on Engineering for Human-Computer Interaction (EHCI'98)*, Heraklion, Greece.
- COUTAZ, J., CROWLEY, J.L., & BÉRARD, F. (1997), Eigen-Space Coding as a Means to Support Privacy in Computer Mediated Communication, in *Proceedings of the IFIP Conference on Human Computer Interaction (INTERACT'97)*, Kluwer Academic Publishers, Sydney.
- CROWLEY, J.L., COUTAZ, J., & BÉRARD, F. (2000), Things That See, in *Communications of the ACM*, ACM Press, vol. 43, no. 3, pp. 54-64.
- DAHLEY, A., WISNESKI, C., & ISHII, H. (1998), Water Lamp and Pinwheels: Ambient projection of digital information inot architectural space, in *CHI 98 conference summary on Human factors in computing systems*, pp. 269-270.
- DENNET, D. (1995), *Elbow Room: The varieties of free will worth wanting*. MIT Press, Cambridge, MA.
- DERICHE, R. (1990), Fast algorithms for low-level vision, in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, no. 1, pp. 78-87.
- DIX, ALAN, FINLAY, ABOWD, & BEALE EDS. (1998) *Human Computer Interaction*, Second Edition. Prentice Hall International.
- DOURISH, P. (1993), Culture and Control in a Media Space, in *Proceedings of the Third European Conference on Computer-Supported Cooperative Work (ECSCW'93)*, Kluwer Academic Publishers, Milan, pp. 125-138.
- DOURISH, P. (2001), *Where the Action Is: The Foundations of Embodied Interaction*. The MIT Press, Cambridge, MA.
- DOURISH, P., ADLER, A., BELLOTTI, V., & HENDERSON, H. (1996), Your Place or Mine? Learning from Long-Term Use of Audio-Video Communication, in *Computer Supported Cooperative Work: The Journal of Collaborative Computing*, Kluwer Academic Publishers, vol. 5, no. 1, pp. 33-62.
- DOURISH, P., & BELLOTTI, V. (1992), Awareness and Coordination in Shared Workspaces, in *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW'92)*, Toronto, pp. 107-114.

- DOURISH, P., & BLY, S. (1992), Portholes: Supporting Awareness in a Distributed Work Group, in *Proceedings of the ACM/SIGCHI Conference on Human Factors in Computing Systems (CHI'92)*, Monterey, CA, pp. 541-547.
- DUVAL, S., & WICKLUND, R. (1972). *A theory of objective self-awareness*. Academic Press, New York, NY.
- EGIDO, C. (1988), Video Conferencing as a Technology to Support Group Work: A Review of its Failures, in *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW'88)*, Portland, OR, pp. 13-24.
- EGIDO, C. (1990), Teleconferencing as a Technology to Support Cooperative Works: Its Possibilities and Limitations, in *Intellectual Teamwork: Social and Technological Foundations of Cooperative Work*, Galegher, Kraut, and Egidio eds., Lawrence Erlbaum Associates, pp. 351-371.
- FISH, R.S., KRAUT, R.E., & CHALFONTE, B.L. (1990), The VideoWindow System in Informal Communications, in *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW'90)*, Los Angeles, pp. 1-11.
- FISH, R.S., KRAUT, R.E., RICE, R.E., & ROOT, R.W. (1992), Evaluating Video as a Technology for Information Communication, in *Proceedings of the ACM/SIGCHI Conference on Human Factors in Computing Systems (CHI'92)*, Monterey, CA, pp. 37-48.
- FISH, R.S., KRAUT, R.E., RICE, R.E., & ROOT, R.W. (1993), Video as a Technology for Informal Communication, in *Communications of the ACM*, ACM Press, vol. 36, no. 1, pp. 48-61.
- FITZPATRICK, G. (1998). *The Locales Framework: Understanding and designing for co-operative work*. Ph.D. thesis, The University of Queensland.
- FITZPATRICK, G., KAPLAN, S., MANSFIELD, T., DAVID, A., & SEGALL, B. (2002). Supporting Public Availability and Accessibility with Elvin: Experiences and Reflections, In *Computer Supported Cooperative Work*, v.11 n.3, pp. 447-474.
- FRIGO, M., & JOHNSON, S. (1998), FFTW: An Adaptive Software Architecture for the FFT, in *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP'98)*, Seattle, vol 3., pp. 1381-1384.
- GAVER, W. (1992), The Affordances of Media Spaces for Collaboration, in *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW'92)*, Toronto, pp. 17-24.
- GAVER, W., MORAN, T., MACLEAN, A., LÖVSTRAND, L., DOURISH, P., CARTER, K., & BUXTON, W. (1992), Realizing a Video Environment: EuroPARC's RAVE System, in *Proceedings of the ACM/SIGCHI Conference on Human Factors in Computing Systems (CHI'92)*, Monterey, CA, pp. 27-34.
- GAVISON, R. (1980), Privacy and the Limits of Law. In *Yale Law Journal*, 89:3 (January), The Yale Law Journal Company, New Haven, CT, pp. 421-471.
- GOFFMAN, E. (1959), *The Presentation of Self in Everyday Life*. Doubleday Publishers, Garden City, NY.

- GREENBERG, S. (1996), Peepholes: Low Cost Awareness of One's Community, in *Companion Proceedings of the ACM/SIGCHI Conference on Human Factors in Computing Systems (CHI'96)*, Vancouver, pp. 206-207.
- GREENBERG, S. (2004, In press), Toolkits and Interface Creativity. Special Issue on Groupware, *Multimedia Tools and Applications*, Kluwer Academic Publishers.
- GREENBERG S., & KUZUOKA, H. (2000), Using Digital but Physical Surrogates to Mediate Awareness, Communication and Privacy in Media Spaces. *Personal Technologies*, 4:1 (January). Elsevier.
- GREENBERG, S., & ROSMAN, M. (2003), Using a Room Metaphor to Ease Transitions in Groupware. In *Sharing Expertise: Beyond Knowledge Management*. M. Ackerman, V. Pipek, & V. Wulf, Eds. MIT Press, Cambridge, MA, pp. 203–256.
- GREENBERG, S. AND ROUNDING, M. (2001), The Notification Collage: Posting Information to Public and Personal Displays, in *Proceedings of the ACM/SIGCHI Conference on Human Factors in Computing Systems (CHI 2001)*, Seattle, pp. 515-521.
- GRUDIN, J., (1988) Why CSCW Applications Fail: Problems in the Design and Evaluation of Organizational Interfaces, In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW'88)*, Portland, OR, pp. 85-93.
- GRUDIN, J. (2001), Desituating Action: Digital Representation of Context. In *Human-Computer Interaction*, 16:2–4, Lawrence Erlbaum Associates, Hillsdale, NJ, pp. 269–286.
- GUTWIN, C. (1998), *Workspace Awareness in Real-Time Distributed Groupware*, Ph.D. thesis, Department of Computer Science, University of Calgary.
- GUTWIN, C. & PENNER, R. (2002). Improving interpretation of remote gestures with telepointer traces. In *Proceedings of the 2002 ACM Conference on Computer Supported Cooperative Work (CSCW 2002)*, New Orleans, pp. 49-57.
- HACKMAN, J.R. (1985), Doing research that makes a difference, in *Doing research that is useful for theory and practice*, Lawler, Mohrman, Mohrman, Ledford, Cummings and Associates eds., Jossey-Bass publisher.
- HALL, E.T. (1966), *Distances in Man: The Hidden Dimension*. Double Day, Garden City, NY.
- HARPER, R.H.R. (1996), Why People Do and Don't Wear Active Badges: A Case Study, in *Computer Supported Cooperative Work: The Journal of Collaborative Computing*, Kluwer Academic Publishers, vol. 4, no. 4, pp. 297-318.
- HARRISON, S., & DOURISH, P. (1996), Re-place-ing Space: The Roles of Place and Space and Collaborative Systems. In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW'96, Cambridge)*. ACM Press, New York, NY, pp. 67–76.
- HIXON, R. (1987), *Privacy in a public society: Human rights in conflict*. Oxford University Press, New York.
- HOCHHEISER, H. (2002), The platform for privacy preference as a social protocol: An examination within the U.S. policy context. In *ACM Transactions on Internet Technology (TOIT)*, 2:4, ACM Press, New York, NY, pp. 276– 306.

- HONG, J. BORRIELLO, G., LANDAY, J., McDONALD, D., SCHILIT, B., & TYGAR, D. (2003) Privacy and Security in the Location-enhanced World Wide Web. In *Proceedings of Ubicomp 2003*, Seattle.
- HONG, J.I., NG, J.D., LEDERER, S., & LANDAY, J.A. (2004), Privacy Risk Models for Designing Privacy-Sensitive Ubiquitous Computing Systems. In *Proceedings of the Conference on Designing Interactive Systems (DIS2004, Cambridge)*, ACM Press, pp. 91-100.
- HORN, D. (2001), *Seeing is believing: Video quality and lie detection*, Ph.D. Dissertation, University of Michigan.
- HUDSON, S.E., & SMITH, I. (1996), Techniques for Addressing Fundamental Privacy and Disruption Tradeoffs in Awareness Support Systems, in *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW'96)*, Cambridge, MA, pp. 248-247.
- ISSACS, E.A., TANG, J.C., & MORRIS, T. (1996), Piazza: A Desktop Environment Supporting Impromptu and Planned Interactions, in *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW'96)*, Cambridge, MA, pp. 315-324.
- JANCKE, G., VENOLIA, G.D., GRUDIN, J., CADIZ, JJ, & GUPTA, A. (2001), Linking Public Spaces: Technical and Social Issues, in *Proceedings of the ACM/SIGCHI Conference on Human Factors in Computing Systems (CHI 2001)*, Seattle, pp 530-537.
- JOHNSON, B., & GREENBERG, S. (1999), Judging People's Availability for Interaction from Video Snapshots, in *Proceedings of the Hawai'i International Conference On System Sciences*, Maui.
- JUNESTRAND, S., KEIJER, U. & TOLLMAR, K. (2001), Private and public digital domestic spaces. In *International Journal of Human-Computer Studies*, 54, 5 (May), Academic Press, New York, NY, pp. 753–778.
- KELVIN, P. (1973), A Social-Psychological Examination of Privacy, in *British Journal of Social and Clinical Psychology*, British Psychological Society, vol. 12, pp. 248-261.
- KNUDTZON, K., THOMAS, C., & SHNEIDERMAN, B. (2002). *Theories in Computer Human Interaction*. <http://www.cs.umd.edu/class/fall2002/cmssc838s/tichi/printer/intro.html>.
- KRAUT, R., EGIDO, C., AND GALEGHER, J. (1988), Patterns of Contact and Communication in Scientific Research Collaboration, in *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW'88)*, New York, pp. 1-12.
- KRAUT, R.E., FISH, R.S., ROOT, R.W., & CHALFONTE, B.L. (1990), Informal Communication in Organizations: Form, Function, and Technology, in *Peoples Reactions To Technology*, Oskamp and Spacapan eds., Sage Publications, pp. 145-199.
- LANGHEINRICH, M. (2001), Privacy by Design—Principles of Privacy-Aware Ubiquitous Systems, in *Proceedings of Ubicomp 2001*, Atlanta.
- LEE, A., GIRGENSOHN, A., & SCHLUETER, K. (1997), NYNEX Portholes: Initial user reactions and redesign implications, in *Proceedings of the ACM/SIGGROUP Conference on Groupware (GROUP'97)*, pp. 385-394.
- LUFF, P. & HEATH, C. (1998), Mobility in Collaboration, In *Proceedings of CSCW'98*. ACM Press, New York, NY. pp. 305-314.

- LÖVSTRAND, L. (1991), Being Selectively Aware with the Khronica System, in *Proceedings of the Second European Conference on Computer-Supported Cooperative Work (ECSCW'91)*, Amsterdam, pp. 265-277.
- MANSFIELD, J. & KAPLAN, S. (2001), Designing for co-evolution in information systems. In *Conference on Complex and Dynamic Systems Architecture*, Brisbane.
- MANTEI, M.M., BAECKER, R.M., SELLEN, A.J., BUXTON, W.A.S., MILLIGAN, T., & WELLMAN, B. (1991), Experiences in the Use of a Media Space, in *Proceedings of the ACM/SIGCHI Conference on Human Factors in Computing Systems (CHI'91)*, New Orleans, pp. 203-208.
- MCCANNE, S., BREWER, E., KATZ, R., ROWE, L., AMIR, E., CHAWATHE, Y., COOPERSMITH, A., MAYER-PATEL, K., RAMAN, S., SCHUETT, A., SIMPSON, D., SWAN, A., TUNG, T.L., WU, D., & SMITH, B. (1997), Toward a Common Infrastructure for Multimedia-Networking Middleware, in *Proceedings of the 7th International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV'97)*, St. Louis.
- MCCANNE, S., & JACOBSON, V. (1995), vic: A Flexible Framework for Packet Video, in *Proceedings of the third ACM international conference on Multimedia*, San Francisco.
- MCEWAN, G. & GREENBERG, S. (2005), Community Bar: Designing for Awareness and Interaction. In *ACM CHI 2005 Workshop on Awareness systems: Known Results, Theory, Concepts and Future Challenges*. Organized by P. Markopoulos, B. de Ruyter, and W. Mackay.
- MOORE, G. (1997), Sharing Faces, Places, and Spaces: The Ontario Telepresence Project Field Studies, in *Video Mediated Communication*, Finn, Sellen, and Wilbur eds., pp. 301-322.
- NARDI, B.A., KUCHINSKY, A., WHITTAKER, S., LEICHTNER, R., & SCHWARZ, H. (1997). Video-as-Data: Technical and Social Aspects of a Collaborative Multimedia Application. In *Video-Mediated Communication*, K. Finn, A. Sellen, and S. Wilbur Eds. Lawrence Erlbaum Associates Publishers, Mahwah, NJ, pp. 487-517.
- NARDI, B.A., WHITTAKER, S., & BRADNER, E. (2000), Interaction and Outeraction: Instant Messaging in Action, in *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW'00)*, Philadelphia, pp. 79-89.
- NIELSEN, J. (1993). *Usability Engineering*, Academic Press.
- NEUSTAEDTER, C. (2003), *Balancing Privacy and Awareness in a Home Media Space*. MSc Thesis, Department of Computer Science, University of Calgary.
- NEUSTAEDTER, C. & GREENBERG, S. (2003) The Design of a Context-Aware Home Media Space. *Proceedings of UBIComp 2003 Fifth International Conference on Ubiquitous Computing*. LNCS Vol 2864, Springer-Verlag, pp. 297-314,
- NEUSTAEDTER, C., GREENBERG, S., & BOYLE, M. (2005, in press), Blur Filtration Fails to Preserve Privacy for Home-Based Video Conferencing. To appear in *ACM Transactions on Computer Human Interactions (TOCHI)*.
- OECD (1980) *Guidelines on the Protection of Privacy and Transborder Flows of Personal Data*. OECD, Paris.

- OLSON, M.H., & BLY, S.A. (1991), The Portland Experience: a report on a distributed research group, in *Computer-supported Cooperative Work and Groupware*, Greenberg ed., Academic Press, pp. 81-98.
- PAGANI, D.S., & MACKAY, W.E. (1993), Bringing Media Spaces into the Real World, in *Proceedings of the Third European Conference on Computer-Supported Cooperative Work (ECSCW'93)*, Kluwer Academic Publishers, Milan, pp. 341-356.
- PERING, T., LIGHT, J., SUNDAR, M., HAYES, G.R., RAGHUNATHAN, V., PATTISON, E., & WANT, R. (2003), The Personal Server: Personal Content for Situated Displays. In *Extended Abstracts of UbiComp 2003*, October 12-15, Seattle, WA, USA, pp. 97-99.
- PALEN, L., & DOURISH, P. (2003) Unpacking Privacy for a Networked World. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI 2003, Ft. Lauderdale)*, ACM Press, New York, NY, pp. 129-137.
- POSNER, R.A. (1981), The Economics of Privacy. In *The American Economic Review*, 71:2, American Economic Association, Nashville, TN, pp. 405-409.
- REASON, J. (1990), *Human Error*. Cambridge University Press, New York, NY.
- RODDEN, T. (1996), Populating the Application: A Model of Awareness for Cooperative Applications. In *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW'96, Cambridge)*, ACM Press, New York, NY. pp. 87-96.
- ROOT, R.W. (1988), Design of a Multi-Media Vehicle for Social Browsing, in *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW'88)*, Portland, OR, pp. 25-38.
- ROSEMAN, M., & GREENBERG, S. (1996), Building Real Time Groupware with GroupKit, in *ACM Transactions on Computer Human Interaction*, ACM Press, vol. 3, no. 1, pp. 66-106.
- ROUNDING, M. (2004), *Informal Awareness and Casual Interaction with the Notification Collage*. MSc Thesis, Department of Computer Science, University of Calgary, Calgary.
- ROUSSEL, N. (2001), Exploring New Uses of Video with VideoSpace, in *Proceedings of the Eighth IFIP Working Conference on Engineering for Human-Computer Interaction (EHCI'01)*, Springer-Verlag Publishers.
- SAMARAJIVA, R. (1997), Interactivity as Though Privacy Matters. In *Technology and Privacy: The New Landscape*, P. Agre & M. Rottenberg, Eds. MIT Press, Cambridge, MA.
- SAMUELSON, P. (2000), Privacy as Intellectual Property? In *Stanford Law Review*, vol. 52, Stanford University School of Law, Stanford CA, pp. 1125-1174.
- SCHWARTZ, B. (1968), The Social Psychology of Privacy. In *American Journal of Sociology*, 73:6, University of Chicago Press, Chicago, IL, pp. 741-752.
- SIMON, H. A. (1996). *The sciences of the artificial* (3rd ed.). MIT Press, Cambridge, MA.
- SIMONS, D.J., & LEVIN, D.T. (1997), Change blindness, in *Trends in Cognitive Sciences*, Elsevier Science Ltd., vol. 1, no. 7, pp. 261-267.

- SMITH, I., & HUDSON, S.E. (1995), Low Disturbance Audio for Awareness and Privacy in Media Space Applications, in *Proceedings of the third ACM international conference on Multimedia*, San Francisco, pp. 91-97.
- SMITH, I.E., HUDSON, S.E., MYNATT, E.D., & SELBIE, J.R. (1995), Applying Cryptographic Techniques to Problems in Media Space Security, in *Proceedings of the ACM/SIGOIS Conference on Organizational Computing Systems (COOCS'95)*, Milpitas, CA, pp. 190-196.
- SPIEKERMANN, S., GROSSKLAGS, J., & BERENDT, B. (2001), E-privacy in 2nd Generation E-Commerce: Privacy Preferences versus actual Behavior. In *Proceedings of the 3rd ACM Conference on Electronic Commerce (EC'01, Tampa)*, ACM Press, New York, NY, 38-47.
- STEPHENSON, G.M., AYLING, K., & RUTTER, R. (1976), The role of visual communication in social exchange, in *British Journal of Social and Clinical Psychology*, British Psychological Society, vol. 15, pp. 113-120.
- SUCHMAN, L. (1987), *Plans and Situated Actions: The Problem of Human-Machine Communication*. Cambridge University Press.
- TANG, A.H.T. (2005), *Embodiments in Mixed Presence Groupware*. MSc Thesis, Department of Computer Science, University of Calgary, Calgary.
- TANG, C. (2003), *Capturing and Visualizing Histories of Multimedia-based Casual Interactions*. M.Sc. Thesis, Department of Computer Science.
- TANG, J.C., & ISSACS, E. (1993), Why Do Users Like Video?, in *Computer Supported Cooperative Work (CSCW)*, Kluwer Academic Publishers, vol. 1, no. 3, pp. 163-196.
- TANG, J.C., ISAACS, E.A., & RUA, M. (1994), Supporting Distributed Groups with a Montage of Lightweight Interactions, in *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW'94)*, Chapel Hill, NC, pp. 23-34.
- TANG, J.C., & RUA, M. (1994), Montage: Providing Teleproximity for Distributed Groups, in *Proceedings of the ACM/SIGCHI Conference on Human Factors in Computing Systems (CHI'94)*, Boston, pp. 37-43.
- TANG, J.C., YANKOLOVICH, N., BEGOLE, J., VAN KLEEK, M., LI, F., & BHALODIA, J. (2001), ConNexus to Awarenex: Extending awareness to mobile users, in *Proceedings of the ACM/SIGCHI Conference on Human Factors in Computing Systems (CHI 2001)*, Seattle, pp. 221-228.
- TWENTIETH CENTURY FOX & DREAMWORKS PICTURES (2002) *Minority Report*. Feature film.
- WANT, R., HOPPER, A., FALCÃO, V., & GIBBONS, J. (1992), The Active Badge Location System, in *ACM Transactions on Information Systems*, ACM Press, vol. 10, no. 1, pp. 91-102.
- WESTIN, A. (1967), *Privacy and Freedom*. Atheneum, New York, NY.
- WHITTAKER, S. (1995), Rethinking video as a technology for interpersonal communications: theory and design implications, in *International Journal of Human-Computer Studies*, Academic Press, vol. 42, no. 5, pp. 501-530.
- WHITTAKER, S., FROHLICH, D., & DALY-JONES, O. (1994), Informal workplace communication: What is it like and how might we support it?, in *Proceedings of the*

- ACM/SIGCHI Conference on Human Factors in Computing Systems (CHI'94)*, Boston, pp. 131-137.
- WHITTAKER, S., & O'CONNELL, B. (1997), The Role of Vision in Face-to-Face and Mediated Communication, in *Video Mediated Communication*, Finn, Sellen, and Wilbur eds., Lawrence Erlbaum Associates Inc., pp. 23-50.
- WHITTAKER, S., JONES, Q., NARDI, B., CREECH, M., TERVEEN, L., ISAACS, E., & HAINSWORTH, J. (2004), ContactMap: Organizing communication in a social desktop. In *ACM Transactions on Computer-Human Interaction (TOCHI)*, ACM Press, 11:4, pp. 445-471.
- WINOGRAD, T. & FLORES, C.F. (1986), *Understanding Computers and Cognition*, Ablex Publishing.
- WOODWARD JR., J.D. (2001) Super Bowl Surveillance: Facing Up to Biometrics. RAND Issue Paper IP-20, <http://www.rand.org/publications/IP/IP209>.
- ZHAO, Q.A., & STASKO, J.T. (1998), Evaluating Image Filtering Based Techniques in Media Space Applications, in *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW'98)*, Seattle, pp. 11-18.

## Appendix A. Co-author permission



February 23, 2005

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

I, Saul Greenberg, give Michael John Boyle permission to use co-authored work from our papers “The Effects of Filtered Video on Awareness and Privacy”, “The Language of Privacy: Learning from Video Media Space Analysis and Design”, and “Blur Filtration Fails to Preserve Privacy for Home-Based Video Conferencing” for Chapters 3~7 of his thesis and to have this work microfilmed.

Sincerely,

Saul Greenberg



UNIVERSITY OF  
CALGARY

---

February 23, 2005

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

I, Christopher Edwards, give Michael John Boyle permission to use co-authored work from our paper "The Effects of Filtered Video on Awareness and Privacy" for Chapter 4 of his thesis and to have this work microfilmed.

Sincerely,

Christopher Edwards



UNIVERSITY OF  
CALGARY

---

February 23, 2005

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

I, Carman Neustaedter, give Michael John Boyle permission to use co-authored work from our paper “Blur Filtration Fails to Preserve Privacy for Home-Based Video Conferencing” for Chapters 4 of his thesis and to have this work microfilmed.

Sincerely,

Carman Neustaedter

## Appendix B. Glossary

This glossary lists vocabulary in the privacy theory which may be used ways that are specialisations of their generally accepted “dictionary” definitions which may deviate from conventionally understood meanings. Where indicated, definitions given below have been taken from the Oxford English Dictionary of Current English. If a word comes from the theories of a particular privacy researcher, the researcher is indicated in the definition. My goal with this glossary is not to communicate what I think the definition of a word ought to be used, but rather equip the reader of Acts II and III with knowledge about *how I have actually used* them.

**absence:** in the dictionary sense; being away; cf. presence.

**access control:** computer supported confidentiality technique in which owners of systems, networks, or data expressly grant or deny permission for others to read or write data, or use the systems and networks; cf. authentication, authorisation.

**accessibility:** the degree to which a person may be reached for communication or interaction, the channels by which the person may be reached; cf. availability.

**accountability:** in Dourish’s theory, the degree to which a system allows users to know what is going on inside the system, particularly how the system is handling their sensitive information.

**accuracy:** the degree to which some piece of information conforms to the truth, assuming that there is an accepted truth; typically refers to content accuracy but may also refer to temporal accuracy (timeliness); cf. fidelity, precision.

**action:** in the dictionary sense, the doing of an actor, as seen from a lower-level mechanical perspective; cf. activity, interaction.

**activity:** in the dictionary sense, the coordinated doings of an actor, as seen from a higher-level intentional perspective; cf. action, interaction.

**actor:** a person as a social entity.

**aesthetic harm:** a undesirable outcome of a privacy violation in which an actor looks “bad”; cf. strategic harm.

**affect:** unstated emotion, attitude, and disposition independent of logic or reason, and stated content.

**aggregate:** in the dictionary sense, to compile together small pieces of information.

**ambient:** surrounding, enveloping; cf. ephemeral.

**ambiguity:** in the dictionary sense, the degree to which more than one meaning is thought to be given to something; cf. certainty.

**analysis:** in the dictionary sense, detailed examination and combination of information.

**anonymity:** in the sense used by Altman, the condition of being unnoticed; also, the condition of being unrecognised; also, the condition of being without a name; cf. pseudonymity.

**appearance:** in the dictionary sense, a person’s visible form; also, implicitly, the social meanings attached to their postures, movements, and costume.

**apprehension:** in the dictionary sense, fearful anticipation.

**archival:** in the dictionary sense, the act of storing information for later use.

**attained privacy:** in Altman’s theory, the actual level of privacy that a person experiences, the result of dialectic negotiation

in which the environment constraints desired privacy.

**attention:** in the dictionary sense, the application of a person's mind assuming Norman's model of peripheral and focal working memory.

**aural:** in the dictionary sense, relating to sound or hearing; cf. visual.

**authentication:** in the computer security sense, the process of verifying the correct identity of a person such as by checking a name/password pair; cf. access control.

**authorisation:** in the computer security sense, the process of determining what access rights a person has for accessing systems, networks or information; cf. access control.

**autonomy:** in my descriptive theory, control over the observable manifestations of a person's social self, such as appearance, behaviours, actions, and utterances.

**availability:** the degree to which a person is willing to receive communication or engage in interaction; cf. accessibility.

**back:** in Goffman's theory, the "stage" on which a secondary presentation not made to the audience is performed; may contradict that offered on the "front stage"; cf. front.

**believability:** in the dictionary sense, the degree to which something can be accepted as true.

**boundary:** in Altman's theory, an understood separation between things; boundaries may be purely mental phenomena, or may have physical manifestations; they are porous, leaky, often ambiguously defined, and are sometimes likened to fields (as in Fitzpatrick's notion of centres and peripheries).

**capture:** in the dictionary sense, the initial sensing of environmental data, e.g., video.

**certainty:** in the dictionary sense, the degree to which the meaning of something is understood without doubt; cf. ambiguity.

**choice:** in the dictionary sense, a variety of options and the power to select one for effect.

**coarse-grained control:** control in which the choices are few and the differences between each choice are great.

**cognition:** in the dictionary sense, knowing, perceiving, or conceiving in the mind as a faculty distinct from emotion and volition.

**computer security:** a holistic discipline of practice that seeks to ensure that computer systems, networks, and information are used only by those for whom it is intended, available for use when intended, are used only for intended purposes, and operate in intended manners.

**confidentiality:** in my descriptive theory, control over the fidelity at which information is accessed, which covers not only the fidelity of the content, but also the operations that may be performed on it and the purposes motivating such operations.

**conformance:** in the dictionary sense, fit to external expectations; cf. deviance.

**constrain:** in the dictionary sense, to alter, limit, guide or otherwise influence to course of an action.

**content control:** computer supported confidentiality technique in which information presented is altered according to the person accessing it; cf. access control.

**context:** in the dictionary sense, unstated aspects of a scenario which fix the scenario's meaning.

**control:** in Reason's theory, the ability and power of a person to put something into which ever of its normal states they desire.

**cooperation:** in the dictionary sense, the act of working together or in favour of another.

**costume:** in Goffman's theory, vestments and adornments e.g., clothing, hairstyle, makeup, jewellery, tattoos.

**crowding:** in Altman's theory, a condition when a person has too many interactions with others, as typical arises when many people occupy a limited space; cf. isolation.

**cryptography:** a discipline which seeks to find ways alternate, usually secret, codes for information.

**cue:** in the dictionary sense, a signal or hint.

**data integrity:** the degree to which information is accurate and complete.

**decontextualisation:** in Bellotti's theory, the condition in which the context for an action is not communicated along with the action.

**deliberate abuse:** wilful circumvention of or interference in another's privacy control.

**desituated action:** in Bellotti's theory, action which does not happen in a clearly situated context; cf. decontextualisation.

**deviance:** in the dictionary sense, divergence from external expectations; cf. conformance.

**dialectic:** in Altman's theory, a back-and-forth negotiation process between dyads.

**digital persona:** in Burkert's theory, digital information that is related to the identity and embodiment of an individual.

**disclosure:** in the dictionary sense, to reveal some information to make it known by another.

**disclosure boundary tension:** in Palen & Dourish's theory, the tension that exists between a person's need for non-disclosure and the need to disclose information in order to be sociable and act in a social setting.

**disembodiment:** in Bellotti's theory, a condition in which an actor is logically or physically separated from their embodiment.

**disinformation:** intentionally inaccurate information designed to obscure the truth.

**dissociation:** in Bellotti's theory, a condition in which the actions of an actor are separated from the actor in a way that makes it hard to tell which actor performed an action.

**distortion filtration:** a content control technique which alters a representation of information to reduce the fidelity at which the information is represented.

**distraction:** in the dictionary sense, the condition in which something causes a person's attention to be directed away from intended targets, usually unexpectedly and undesirably.

**dynamic:** in the dictionary sense, perpetually changing.

**edit:** the modification of information.

**effort:** in the dictionary sense, the application of mental or physical energy.

**embodiment:** in Dourish's theory, the observable and knowable representation of a person or entity in a world; in the physical world, this is a person's corporeal body; in an instant messenger, it could be made up of a conversation window to that person.

**encounter:** in the dictionary sense, a meeting of people.

**environment:** in my theory, the physical and social setting for a performance.

**ephemeral:** ambient and lasting only a short while; cf. ambient.

**expectation:** in the dictionary sense, what is thought likely to happen or deemed necessary, desired, or fit to happen.

**explicit:** in the dictionary sense, stated outright.

**fantasy:** in the dictionary sense, fanciful and engrossing mental imagery.

**feedback:** in the dictionary sense, the return of information about an action to the actor.

**feed-through:** in groupware, the transmission of cues signalling an action as it is in progress.

**fidelity:** the accuracy and precision of information or the operations available for access.

**fine-grained control:** control in which there are many choices and although the difference between two choices may small, there are enough choices that there is great diversity among the collection.

**flaw:** in the dictionary sense, an imperfection; cf. deviance.

**focus:** the object of one's attention.

**freedom:** in the dictionary sense, the condition of being unrestricted.

**front:** in Goffman's theory, the primary performance for an audience; cf. back.

**genre of disclosure:** in Palen & Dourish's theory, the emerging pattern of disclosure between people taking into account situational and institutional factors and evolving over time.

**harm:** in the dictionary sense, a negative outcome.

**idealisation:** in the dictionary sense, the regard of something as ideal and the disregard of apparent flaws.

**identity:** in the dictionary sense, that totality of things which specifies an individual.

**identity theft:** a computer-related crime in which one person falsely and without permission acts on behalf of another, usually for personal profit.

**impersonation:** in the dictionary sense, the momentary assuming of another person's identity so that a person's actions are treated

as though they were made by someone else; cf. identity theft.

**implicit:** in the dictionary sense, unstated but often understood.

**impression:** in the dictionary sense, the effect on mind or feelings that a person makes on others.

**inadvertent privacy infractions:** unintentional circumvention or interference in a person's privacy control, often their own control.

**inference:** in the dictionary sense, a statement derived from supporting data.

**information about the self:** information about a person's appearance, actions, activities, utterances, encounters, etc.

**institutional:** in the dictionary sense, relating to or decided by an institution; concerning organised groups like societies and enterprises; cf. situational.

**intention:** in the dictionary sense, the aim or purpose for an action or activity.

**interaction:** in the dictionary sense, the coordinated and interspersed actions of two or more people.

**interpersonal distance:** a measure of spatial and psychological closeness.

**interpersonal privacy:** the parts of privacy affected by interpersonal relationships; e.g., stumbling across someone in a compromising situation; cf. macrosociological privacy.

**interpolation:** in the dictionary sense, the deduction of intermediate values.

**intimacy:** the feeling of being close with or closely acquainted with another person.

**isolation:** in Altman's theory, the condition of being separated spatially and socially from others; cf. crowding.

**liberty:** in the dictionary sense, the ability and permission to do as one pleases.

**lightweight:** not requiring much physical or cognitive effort.

**location:** in the dictionary sense, the absolute point or region in a spatial world a person occupies and often the social interpretation of it; cf. position.

**macrosociological privacy:** the parts of privacy that are affected by collective decisions affecting societies; e.g., justice courts revoking

or upholding doctor-patient confidentiality; cf. interpersonal privacy.

**manner:** in the dictionary sense, the way or style in which an actor delivers the performance; e.g., in an annoyed manner.

**medium:** in the dictionary sense, the intermediary substance through which signals travel.

**misappropriate:** in the dictionary sense, to take wrongly for a person's own use; cf. misuse.

**misinformation:** unintentionally inaccurate information; cf. disinformation.

**misuse:** in the dictionary sense, to wrongly apply to a wrong purpose; cf. misappropriate.

**nimbus:** in Rodden's theory, the subspace in which a person's presence is projected; a person's embodiment and the observable traces it leaves behind; cf. focus.

**normative:** in the dictionary sense, defined by or under the direct influence of norms.

**norm:** in the dictionary sense, a shared understanding of the way something is customarily done.

**obligation:** in the dictionary sense, a binding agreement, duty, or responsibility.

**performance:** in Goffman's theory, an actor's self-presentation in a social context.

**periphery:** that which lies just beyond the focus, where people can sense what is going on but they are usually not distracted by it; cf. focus, attention.

**persistence:** storing information temporarily or indefinitely, usually on computer disks; cf. ephemeral.

**personal space:** in Hall's theory, a invisible region around a person's body; it can be psychological uncomfortable for a person when others act in their personal space; cf. territory.

**personally identifying information:** information that can be used in conjunction with other personally identifying information, to establish with certainty the identity of a person; e.g., full legal name, date of birth, citizenship, postal address, place of employment.

**place:** a part of space occupied by actors and artefacts; usually refers not just to the Cartesian coordinates which bound the

subspace if applicable, but also the social expectations about what actors and artefacts will be in it and what they do there; often places are described by common nouns e.g., a coffee house, or proper noun e.g., Bubba's Espresso Emporium; cf. space.

**plausible deniability:** in Nardi's theory, an affordance for solitude/confidentiality/autonomy in which a particular source of information about a person is commonly understood to be habitually inaccurate and imprecise; this ambiguity affords the person a choice of how to behave; e.g., in instant messengers, a person's availability status is inferred by the system from keyboard/mouse idle input time and since it is often incorrect a person can act as though they are absent, ignoring incoming messages, even though they in fact present.

**police:** in the dictionary sense, to keep in order and enforce conformance to norms by routinely monitoring a place for violations and deviance, and punishing those responsible.

**possibility:** in the dictionary sense, the condition that something is capable of happening, independent of its likelihood of happening; cf. probability.

**posture:** in the dictionary sense, the relative positions of a person's body parts, often used to signal disposition.

**position:** in the dictionary sense, the point or region in space a person occupies relative to other people or landmarks; cf. location.

**power:** in the dictionary sense, the mechanical and physical capacity to carry out an act and the social freedom to carry it out without let or hindrance.

**precision:** in the dictionary sense, a part of fidelity that concerns the degree of refinement in a measurement; often split out into content precision, precision in the typical scientific numerical analysis sense, and temporal precision, which is the frequency at which a continuous variable is sampled; cf. accuracy.

**preference:** in my privacy theory, the choices normally available for control and in particular the habitual or situational selection of a few choices in favour of others by an individual person; typically a range of states are deemed normal, and preferences are a subset of the

norms which an individual person considers more favourable.

**presence:** in the dictionary sense, the existence of an embodiment for a person in a space; cf. availability.

**privacy-enhancing technology:** in Burkert's theory, technology which removes personally identifying information.

**probability:** in the dictionary sense, the likelihood that something will occur; cf. possibility.

**process integrity:** in computer security, the assurance that computer services function correctly according to specified requirements.

**processing:** in the dictionary sense, the treatment of information; may include capture, analysis, modification, transmission, or archival.

**prop:** in Goffman's theory, an object (not a person) that is used as a part of a performance; artefact.

**pseudonymity:** in my privacy theory, a condition in which a person has several distinct identities, each customarily used in a different social world.

**publication filtration:** computer supported confidentiality technique in which the personal state variables true values are mapped to representational values, where the transformational mapping used may differ according to, e.g., the viewer; cf. distortion filtration.

**reciprocity:** a property of a world such that when person A may perform an operation on person B (e.g., see, hear, touch) it is always possible for the same operation to be performed by person B on person A.

**reflexive interpretability of action:** in Bellotti's theory, the ability of a person to understand from their own embodiment how others will perceive their action.

**refuge:** in the dictionary sense, shelter from the mental or physical stresses of the environment.

**regulation:** in the dictionary sense, to adapt to requirements; dynamically negotiated control.

**reliability:** in computer security, the degree to which a computer system or process is assured to function correctly according to specification.

**reprimand:** in the dictionary sense, official acknowledgement of someone's wrong-doing, usually accompanied with scolding, punishments, forced restitution.

**resentment:** to retain bad feelings about something.

**reserve:** in Altman's theory, wilful non-disclosure for confidentiality and separation from the environment for solitude.

**reward:** in the dictionary sense, an enjoyable thing to be gained.

**rich:** a property of a communication system in which it conveys many modes of information at high fidelity that foster satisfying experiences.

**right:** in the dictionary sense, something for which there is a moral and normal obligation for society to ensure that every person is never denied.

**risk:** in my privacy theory, the combined probability and severity of a harm arising from a violation of privacy.

**risk/reward disparity:** a condition in which the reward a person gains from using a computer system does not rise proportionally to the risk the system puts to their privacy.

**risk/reward trade-off:** the economic problem behind the disclosure boundary tension; in order to gain some reward, a person must necessarily tolerate increased risk to privacy; cf. disclosure boundary.

**role:** a shared understanding of the performances a person is expected to give in a situation.

**role conflict:** in Adler & Adler's theory, a condition in which a person must satisfy two distinct roles concurrently; e.g., when the parents of college students visit them at their dormitories, the students must simultaneously play the role of a responsible child (in front of their parents) and a fashionable adult (in front of their peers).

**sample:** in the dictionary sense, to take a discrete measurement of a continuous environmental variable.

**satisfice:** in the dictionary sense, to obtain an outcome that is good enough.

**scenery:** in the dictionary sense, the furnishings in the environment that indicate the place for a performance.

**scrutiny:** in the dictionary sense, close, critical inspection and examination.

**secrecy:** wilful non-disclosure to some, but not all, people.

**self:** in the environmental psychological sense, the essence and totality of a person.

**self-appropriation:** in Bellotti's theory, a process by which people tailor their performances to conform to norms fitting the situation.

**self-definition:** in the environmental psychological sense, the act of establishing a separate identity for a person.

**sensitivity:** in my privacy theory, the importance ascribed to retaining careful control over confidentiality with respect to a particular piece of information.

**severity:** the degree of harshness of a harm arising from a privacy violation.

**signifier:** in Goffman's theory, a quality about some thing which can be observed to infer the meaning of it.

**situated action:** in Suchman's theory, the consideration of a person's actions in their original context.

**situational:** in the dictionary sense, related to or defined by the immediate local circumstance; cf. institutional.

**social acceptability:** in the dictionary sense, the degree to which a part of a performance will be generally considered conforming to situationally modified institutional expectations.

**social environment:** the part of the environment which concerns interpersonal relationships, institutional rules, customs, obligations; cf. physical environment.

**social setting:** in Goffman's theory, the social environment and in particular the aspects of the place for a performance which factor significantly in the social environment.

**solitude:** in my privacy theory, control over attention for interactions.

**space:** in the dictionary sense, a bounded expanse in which actors and artefacts (or more precisely their embodiments) exist; cf. place.

**stability:** in the dictionary sense, the degree to which a computer system remains available for use.

**status divisions:** in Schwartz's theory, logical separations of people into higher and lower classes; people in lower classes are generally at a disadvantage in terms of opportunities for advancement, material quality of life, health, and ability to enjoy rights and freedoms.

**strategic harms:** harms which affect a person's or enterprise's execution of plans or material welfare; cf. aesthetic harms.

**surreptitious surveillance:** the act of monitoring a space, the people in it, and their activities, utterances, and interactions in such a way that it is not immediately known to the people that they are being watched.

**territory:** a part of space over which a person has some sort of governance or control; particular, control over which other people have access to the space, the artefacts or information in it; unlike personal space, territory includes parts of space not immediately around a person's body, e.g., their office desk drawers; cf. personal space.

**threat:** a possible privacy violation; cf. risk.

**topic:** in the dictionary sense, theme of some piece of information.

**transition:** in the dictionary sense, a passage or change from one state to another, possibly crossing a boundary.

**transitivity:** in the dictionary sense, a characteristic of information in some system or medium in which it can be transmitted to

other people, through possibly other systems or media.

**trust:** in the dictionary sense, the degree to which one person expects another person to conform to shared expectations for behaviour.

**usability:** in the dictionary sense, the degree to which it is easy and natural for a person to use a computer system; cf. utility.

**use:** in the dictionary sense, the aim or purpose of an action which incorporates some artefact or piece of information.

**utility:** in the dictionary sense, the degree to which something is useful; cf. usability.

**utterance:** in the dictionary sense, verbal communication, usually as spoken words but could also be text in email or hand signals such as sign language.

**value:** in the dictionary sense, the worth of something.

**violate:** in my privacy theory, to hinder a person's privacy regulation, such as by deviating from normal practice, eliminating preference, reducing their power to select a choice, or overriding the choice.

**whereabouts:** the position of a person as a continuous variable.

**withdrawal:** in Altman's theory, the spatial and psychological separation of a person from others to reduce interactions with them; cf. solitude.