The Language of Privacy: Learning from Video Media Space Analysis and Design

MICHAEL BOYLE and SAUL GREENBERG
University of Calgary

Video media spaces are an excellent crucible for the study of privacy. Their design affords opportunities for misuse, prompts ethical questions, and engenders grave concerns from both users and nonusers. Despite considerable discussion of the privacy problems uncovered in prior work, questions remain as to how to design a privacy-preserving video media space and how to evaluate its effect on privacy. The problem is more deeply rooted than this, however. Privacy is an enormous concept from which a large vocabulary of terms emerges. Disambiguating the meanings of and relationships between these terms facilitates understanding of the link between privacy and design. In this article, we draw from resources in environmental psychology and computer-supported cooperative work (CSCW) to build a broadly and deeply rooted vocabulary for privacy. We relate the vocabulary back to the real and hard problem of designing privacy-preserving video media spaces. In doing so, we facilitate analysis of the privacy-design relationship.

Categories and Subject Descriptors: H.4.3 [**Information Systems Applications**]: Communications Applications—*Computer conferencing teleconferencing, videoconferencing*; H.1.2 [**Models and Principles**]: User/Machine Systems—*Human factors*; K.4.m [**Computers and Society**]: Miscellaneous—*Privacy theory and models*

General Terms: Theory, Human Factors, Design, Languages

Additional Key Words and Phrases: Human-computer interaction, solitude, confidentiality, autonomy, environmental psychology, user interface design, privacy, video media spaces, social interaction, computer-supported cooperative work (CSCW)

1. INTRODUCTION

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,

This research was supported by NSERC and Microsoft Research.

Authors' address: Department of Computer Science, The University of Calgary, Alberta T2N 1N4 Canda; email: boylem@cpsc.ucalgary.ca.

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).

1.1 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. 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 and Bly 1992; 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 and Bly 1991; Mantei et al. 1991; Gaver et al. 1992; 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 [Fish et al. 1990; 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 [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 nonusers alike to be privacy invasive and privacy insensitive [e.g., Gaver et al. 1992; Bellotti and 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 well 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].

1.2 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 generalizing 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. Technocentric 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 [2001] compares "bottom-up" versus "top-down" methods for exploring privacy-design issues. 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 and 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. We have found that 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.

1.3 The Present Article

It is in this progressive context that we present this article. This article is premised on the idea that the language used to discuss the link between privacy and technology design impacts the course of the development of scientific understanding of this link. We focus not so much on the words as how they are used to describe privacy in video media spaces. As the CSCW community draws inspiration and insight from its own empirical work and from theoretical analysis provided by social and behavioral sciences, it becomes increasingly important that the vocabulary used to discuss privacy serves to disambiguate its many facets yet still transmits a holistic perspective of it.

The objective of this article is to describe such a vocabulary. The vocabulary solves an important problem: facilitating the unambiguous discussion of privacy and the impact technology has on it. We attempt to do more than merely summarize the findings of others. We offer a comprehensive deconstruction of privacy not previously found in CSCW literature and establish concrete links to technology not found in social, psychological, or legal literature. Although the vocabulary we present here has been developed to address privacy issues in video media spaces, there are obvious extensions to other system areas that share common problems. Although we have strived to give this vocabulary the widest, most stable theoretical footing, we readily concede that the scientific understanding of privacy is still incomplete and so too is our vocabulary.

We start building this vocabulary in Section 2 by synthesizing CSCW observations of privacy in video media spaces. We broaden the vocabulary in Section 3 by looking at privacy from perspectives established outside the CSCW domain, such as anthropology, architecture, law, behavioral psychology, and sociology. Out of these, we adopt a framework for interpersonal operational privacy in Section 4. This framework is further elaborated in Sections 5, 6, and 7 in which we deconstruct privacy along three lines: solitude, confidentiality, and autonomy.

In the text, vocabulary terms are set in italic when first introduced and discussed. As we discuss these terms, we relate them to observations drawn from video media space design practice and use. While this article 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.

2. THE CSCW PERSPECTIVE

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 roadblocks, but by living with the technology, they came to experience firsthand 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 dichotomy of problems [Bellotti 1998] which consists of the following.

—Deliberate privacy abuses are possible.
—Inadvertent privacy violations are possible.
To these, we add a third problem.
—Users and nonusers 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.

2.1 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 and 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*. Worse, media spaces are not adequately designed to safeguard against malicious use arising from unauthorized access. Thus, they afford the potential for undiagnosed abuse by outsiders. One example is surreptitious surveillance: for example, 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.

2.1.1 *Methods for Controlling Media Space Access.* One way to solve deliberate privacy abuses is with *access control* which puts into place computer security and cryptographic measures to deny unauthorized individuals access to sensitive information [Smith and Hudson 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*. 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 and Smith 1996]. Figure 1 shows several techniques for preserving privacy in video media spaces based on content control. Distortion filters, such as the blur filter in Figure 1, mask

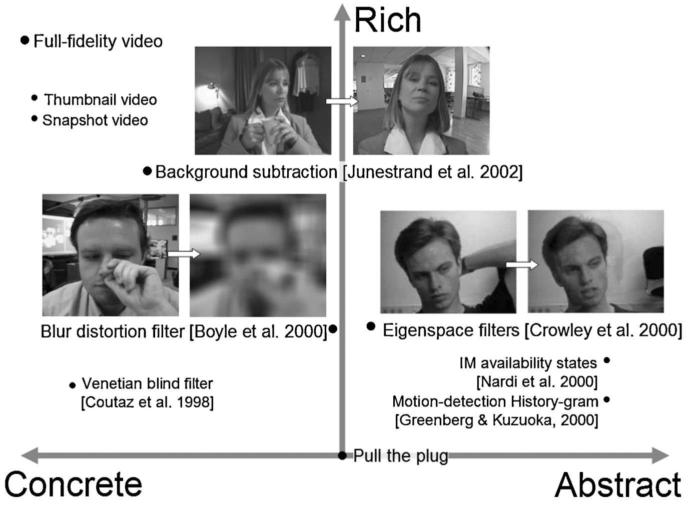
Rich
Full-fidelity video
Thumbnail video
Snapshot video
Background subtraction [Junestrand et al. 2002]
Blur distortion filter [Boyle et al. 2000]
Eigenspace filters [Crowley et al. 2000]
IM availability states [Nardi et al. 2000]
Venetian blind filter [Coutaz et al. 1998]
Motion-detection History-gram [Greenberg & Kuzuoka, 2000]
Pull the plug
Concrete
Abstract

Fig. 1. A design space showing some previously explored techniques for preserving privacy in video media spaces. The space is organized along two dimensions: presentation richness (the quantity of information content presented) and abstraction (how much of the original live video feed is presented).

sensitive details in video while still providing a low-fidelity overview useful for awareness [Zhao and Stasko 1998; Boyle et al. 2000]. Publication filters, such as the background subtraction filter in Figure 1, remove details considered unimportant for awareness information [Coutaz et al. 1998; Junestrand et al. 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 Figure 1 [Crowley et al. 2000].

2.1.2 *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* can be adjusted on a person-by-person, instance-by-instance basis. *Lightweight control* 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 and Heath 1998]. This kind of privacy regulation demands very little cognitive or behavioral effort from the people involved and is usually an implicit activity realized 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 behaviors 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*: 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*, 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 institutionalized 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].

2.2 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 capitalized upon. Accidental violations are known to happen from time to time. *Inadvertent privacy infractions* 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* [Suchman 1987]. Environmental constraints for interactivity keep interactions situated in a temporally and spatially localized *context*. Technology changes these constraints, causing actions and interactions to be desituated and *decontextualised* [Grudin 2001]. Inadvertent privacy violations occur because people are no longer operating in clearly situated contexts [Palen and Dourish 2003].

2.2.1 *Disembodiment Confounds Self-Appropriation. Self-appropriation* is a regulatory process where people modify their behavior and appearance according to social norms and expectations [Bellotti 1998]. Self-appropriation depends on cues for behavior sensed 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 behavior (*norms*) and adapt their behavior accordingly.

The impoverished nature of a video media space means that people often do not appropriate themselves correctly for viewing by distant colleagues. *Disembodiment*—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 the speech and actions of others, increasing chances for miscoordination and miscommunication. This makes it difficult for media space users to understand the privacy wishes of others, leading to inadvertent violations. More broadly, Palen and Dourish 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 [Palen and Dourish 2003].

2.2.2 *Presence in Multiple Places Forces Appropriation in Multiple Contexts.* Place—its architecture and use [Harrison and Dourish 1996]—is an important feedback cue for self-appropriation. Places differ with respect to privacy expectations for example, 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 and 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 behavior as appropriate.

Media spaces, though, join places with differing privacy cultures without such smoothing transitions, permitting weird intersections of privacy expectations [Bellotti 1998; Palen and 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.

2.2.3 *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 and 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 [1998] presents a framework for analyzing deliberate and inadvertent privacy problems in systems and evaluating solutions. 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* 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 signaling 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 deny quick and graceful transition into it [Tang et al. 1994].

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.

2.3 Apprehension

Privacy violations can be *aesthetic* (affecting appearances and impressions) or *strategic* (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. Nonusers 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* itself is a significant problem. Specifically, participants are apprehensive about making bad *impressions* in the media space and the aesthetic or strategic consequences of them.

2.3.1 *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, behavior, 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. [2003] 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. 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.

2.3.2 *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. [1997] mention that no one media space channel alone conveys the complete meaning of an event or utterance. In the hospital media space in their ethnographic study, neurosurgical operating room staff used humor to relieve the stress of a mentally demanding surgical task. The decontextualization of such humor 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 humor 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. [1997] 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 and 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.

2.4 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 in Section 2.1.1. 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 generalized 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 operationalize 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 1.2, this cyclical dependency has prompted many to draw theoretical frameworks for understanding privacy from other disciplines. In the next section, we look at some of the diverse conceptions of privacy that can inform the design of video media spaces.

3. PERSPECTIVES ON PRIVACY

Many disciplines of study must deal with the notion of privacy: anthropology, architecture, behavioral 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 cross-disciplinary survey of privacy-related literature [Brierley-Newell 1995]. She classified works discussing privacy as being "person-centered," "place-centered," or interested in person-environment interactions. By far, the majority 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.

3.1 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 [1998] found this to be the most fundamental and broadly cross-cultural conceptualization of privacy. 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 and Dourish [2003] call this the *disclosure boundary* tension: a tension between one wanting/needing/choosing/being private versus public. This tension carries over to VMS design. From an organizational 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 concept of privacy. The Greeks had their *idion* (private life) and *koinon* (public life) [Arendt 1958]. Goffman [1959] describes front and backstage performances. 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., Boyle et al. 2000]. Schwartz's macrosociological analysis of privacy characterizes it as a "highly institutionalised counterpattern of withdrawal" complementing a pattern of social interaction [Schwartz 1968].

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 stabilize institutions (societies) in which people are organized into status hierarchies. Horizontally, privacy stabilizes 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 stabilizes the hierarchy by reinforcing status divisions. It also allows high-status members of a hierarchy to conceal their flaws from low-status members, preserving idealized impressions that reinforce social superiority, authority, and obligatory relationships. Overall, privacy conceals deviant behavior which, if widely publicized, would destroy social order. Strict conformance is rarely possible and *deviance* 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.

3.2 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 [1968] notes that this kind of privacy can easily be changed by people. Treating privacy as an architectural attribute of space is useful for VMS design. It permits construction of architectural metaphors for privacy safeguards [Greenberg and 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 and 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 and Chemers 1985] or are culturally universal aspects of humanity [Brierley-Newell 1998].

3.3 Privacy as an Interpersonal Process

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

—verbal, for example, telling someone across the VMS link to keep some information secret;

—paraverbal or nonverbal, for example, pointing a VMS camera out the window; or,

—social, for example, 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 [1998] survey is that these behaviors are part of a *privacy process*. Altman in particular sees it as a boundary-regulation process which facilitates the negotiation of access to the self [Altman 1975]. The *self* 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*: the physical environment and also the social environment, that is, the people immediately nearby and society at large.

Altman's [1975] privacy process is a *dialectic*. The actual level of privacy attained is decided through a process of negotiation between the self and the environment. This dialectic is *normative*. 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* 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 [1975] 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* 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 changes in the environment affecting the internalized 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 [1975] theory holds potential heuristic value: because it has been specified so broadly, it can apply to many situations. Yet, Brierley-Newell [1995] speculates that this broadness also makes Altman's theory the most criticized. For example, some critics argue that social interactionalism may be better able to explain the privacy process. Within the CSCW community, Fitzpatrick's [1998] Locales framework applies social interactionalism principles to uncover and comprehend CSCW system design issues. Although it has yet to be done, it is conceivable that methods for analyzing privacy and design in video media spaces could be based on Locales.

3.4 Privacy as a Need, Right, and Freedom

Researchers in behavioral 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 I. 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 in Table I.

Perhaps because of these benefits, people place great *value* on 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

Table I. Negative Aspects of Insufficient Control Over Privacy, and Positive Aspects of Sufficient and Necessary Control Over Privacy (from Altman [1975]) and Brierley-Newell [1995])

Too Few Interactions *(Too Much Privacy)*	Too Many Interactions *(Too Little Privacy)*	Just Right
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 comparison)
	Rage and misbehaviour	Accountability and responsibility
	"Looping" i.e., role separation failures	Fantasy

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 on 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*. The actions of others may deny one this right or impair one's exercise of it. Thus, it is a privacy violation when the actions of others prevent one from obtaining the privacy he needs, he normally enjoys, and society deems that he ought to enjoy. The normalized definition of violation is important. For example, Schwartz [1968] calls surveillance an institutionalized form of privacy violation. The actions of others may prevent one from obtaining desired privacy, but this in 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 I. These outcomes vary in *severity* which is a subjective measure of how bad the harm due to the outcome is.

Although the environment may permit the actions of others that will lead to a privacy violation, one might not choose to invoke such actions. Hence, privacy can be threatened without necessarily being violated. Privacy *threat* and privacy *risk* are used almost synonymously and seem to include the *possibility* of a violation, the *probability* 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*: providing punishments, taboos, social consequences and so on, to discourage others from doing things that violate one's privacy. Schwartz [1968] cites Simmel's [1964] 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. Some privacy violations are so severe that one is permitted to take actions that stop further harm and to 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 behaviors groups employ to govern media space use.

3.5 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 and 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 self-maintenance. 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 them to satisfy both the need to be apart and the need to be together. Even though there is risk, there may also be *reward*: 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* 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 level of 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 concept 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 civilization 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 something or receive some service, emphasizing 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 and 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, 1992; Greenberg and 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, that is, close monitoring of 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* 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 [1992] that is broadly applicable to all genres of CSCW systems.

Reciprocity 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 rebalancing 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* (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 that users and nonusers 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. 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 make when they are presented out of their original context. Relevancy, however, can be just as difficult for computers to measure as sensitivity.

3.6 Summary: Focusing on a Interpersonal Process Model for Privacy

In this section, 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 institutionalized phenomenon,
—a state in which people find themselves,
—a quality of places, and
—a behavioral process governing interactions that seek 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 behavior for self-maintenance 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 article, but our deconstruction will reflect a fundamental assertion that privacy behaviors 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 [1975] broadly characterizes privacy as a boundary-control process regulating access to the self. This concept 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 and Bly 1992; Bellotti 1998]. Not surprisingly, it is also the foundation selected by Palen and Dourish [2003] in their deconstruction of privacy and technology design confluence.

Altman's [1975] 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 rationalizing 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 rationalizing about macrosociological privacy problems, or about problems that span wide temporal or spatial boundaries. While privacy is very important to people, their behaviors (e.g., 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 can 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 [1975] 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 internetwork, 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. An important part of the value of the work done by Palen and Dourish [2003] and by us in this article is to provide these concrete links.

4. FOUNDATIONS OF THE VOCABULARY

In addition to the specialization of Altman's [1975] 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 overarching 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 [1980] emphasizes the role of control in privacy management and points out 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 to underscore parallels to the problems faced by technologists. Gavison's decomposition of privacy yields a powerful vocabulary that we use to disambiguate the many interrelated meanings of privacy discussed by Altman [1975]. 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;

—*confidentiality*: control over other's access to information about oneself, specifically the fidelity of such accesses;

—*autonomy*: control over the observable manifestations of the self, such as action, appearance, impression, and identity.

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 by Palen and 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* is regulated mostly by confidentiality, but also by solitude. The *identity boundary* is regulated by autonomy. The *temporal boundary* 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 $A \ldots$ can drive B into whichever of B's normal range of states A wants B to be in". Gavison [1980] points out two elements of 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 [1975] theory. Such control is founded upon individual and social human behaviors such as those discussed by Altman and Chemers [1980] and Langheinrich [2001]. These behaviors are the low-level mechanical means by which control is exerted. An implication of a dialectic sort of control is that the processes are *satisficing* [Simon 1996]: there is no need for complete control in order to experience privacy.

All three modalities of control are negotiated concurrently. Behaviors 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 [1975] 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, according to Altman's [1975] theory, social. As soon as social interactions of casual or work topics are made possible—either 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 cooperative process. A person will sometimes do things to respect another's 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 capitalized on. 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 behaviors 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.

5. SOLITUDE

Altman [1975] 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 behaviors that are vital to human development, for example, self-evaluation and ego development. 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 a continuum of states, and the extremes present failure conditions that yield negative behavioral, psychological, and physiological responses. For example, *crowding* results when others are granted too much access to the self. *Isolation* results when one cannot interact with others to the degree they wish. Both conditions indicate failures in solitude control.

5.1 Attention and Distraction

To discuss other issues in video media spaces that closely relate to solitude, we generalize Altman's [1975] definition of solitude to include control over where one directs one's *attention* and how one controls *distraction*. 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. Presence and availability are regulated by solitude. Our broadened concept 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 in Section 7.6.

This extension also helps to explain "camera shyness" problems in video media spaces [Lee et al. 1997]. In colocated 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, that is, 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 and Wicklund 1972]. In our extended notion of solitude, technology can invade users' solitude (or permit users to invade another's solitude) by making it difficult for users to control how they direct their attention for self-work and interactions.

5.2 Verbal and Paraverbal Solitude Controls

A variety of individual and social behaviors are used to regulate solitude. Verbal and paraverbal mechanisms for controlling solitude usually involve signaling availability, for example, verbally telling another you wish to be left alone or hanging a "do not disturb" sign outside a hotel door. Desires can be signaled in both the content (the meaning of the words spoken) and the structure (pitch, duration, volume etc. of voice) of speech [Altman and Chemers 1980]. Paraverbal means for signaling one's desired solitude include postures facial expressions, and/or 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 artifacts of compression) mask subtle paraverbal cues for communicating availability. Because such desires must instead be communicated with speech, video media spaces can make the process of signaling solitude desires more explicit and heavyweight. These changes alter social interpretation of the expressed desires.

5.3 Westin's Four Privacy States

Westin [1967], another noted privacy theorist, decomposed privacy into four states.

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

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

—*Anonymity* is the state in which one is physically copresent with others and yet does not expect to be recognized by them and is free from interactions with them. It refers to a condition in which one can be "lost in a crowd."

—*Reserve* 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 the typical exercise of solitude.

5.4 Affordances of Space for Solitude

To regulate solitude, one can go someplace to be alone. These places of *refuge* 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 they can engage with others. Place partially determines *accessibility*, that is, the effort people must expend to engage others for interaction [Harrison and 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 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.

Table II. Example of Interpersonal Distances and the Modalities of Interaction Supported at Each [Hall 1966]

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

People can also capitalize on the ambiguity inherent in some architectural changes to regulate solitude. For example, a closed door ambiguously symbolizes both absence and a wish to be left undisturbed [Root 1988]. People also capitalize on ambiguity when it is possible in computer-mediated environments. For example, Nardi et al. [2001] reports that people use the inaccuracies of IM presence indicators as a 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.

5.5 Personal Space

Space and social behavior interoperate with respect to solitude. *Personal space* 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 behavioral and physiological responses when others physically enter their personal space. *Territory* is similar, but usually implies a recognizably fixed spatial or psychological location, even if it is defined relative to its owner. Territories are important for the regulation of workspace artifacts and confidentiality and will be discussed in Section 6.

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 [1966] describes four interpersonal zones each with differing modalities for social interaction; these are presented in Table II. Because of the 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 from 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 from the person and in easy eye contact.

Personal space, as a tool for solitude regulation, depends on having a range of *interpersonal distances* 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 dictate 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 generalized to include engagement and connectivity. In a typical VMS, only two or three such distances are offered: (1) full interconnectivity, (2) connected to just one other person, and, (3) 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 colocated 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, for example, it is awkward to reposition the VMS camera because of limited cable lengths, lighting, shelf space, and similar factors.

6. CONFIDENTIALITY

Confidentiality is the control of access to information about oneself, for example, 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 emphasizes 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 and 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 the attention of others, while 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 and Dourish [2003]). 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).

6.1 Sensitivity

Sensitivity 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]. Impressions of a person are predicated upon knowledge of the person, and so confidentiality is part of impression management [Goffman 1959]. The harm that could arise from breeches of confidentiality include embarrassment, damage to ego and identity, loss of 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.

6.2 Fidelity

Fidelity is a perception of how faithfully a piece of information represents some truth. It includes both *precision* (how detailed the information is perceived) and *accuracy* (the confidence or certainty placed on 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, 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 and Dourish's [2003] temporal boundary). Information about oneself—the object of confidentiality—may be known by different individuals at different fidelities. Our concept of confidentiality is broadened to address VMS design issues by considering that confidentiality includes control over fidelity. Confidentiality is breeched when a person is unable to control the fidelity with which others are able to access her/his information.

Video media spaces have several dimensions for video fidelity, for example, field of view, resolution, frame rate, codec quality, latency, jitter, and so on. 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 2). These techniques presume that sensitive information lays mostly in image details and so low-fidelity overviews of the video pose less risk [Hudson and Smith 1996]. The manipulation of fidelity (especially timeliness) introduces ambiguity that is incorporated into privacy control for example, "plausible deniability" [Nardi et al. 2000]. For example, distortion filters such as the blur filter shown in Figure 1 can operate at many levels, discarding a little or a lot of fidelity [Boyle et al. 2000]. Of course, while fidelity is reduced, there is no guarantee that these techniques mask the sensitive information. For example, Neustaedter et al. [2003] questioned the effectiveness of a blurring video filter in extremely risky home telecommuting scenarios. They 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* placed in the sender and the number of recipients. We also consider that

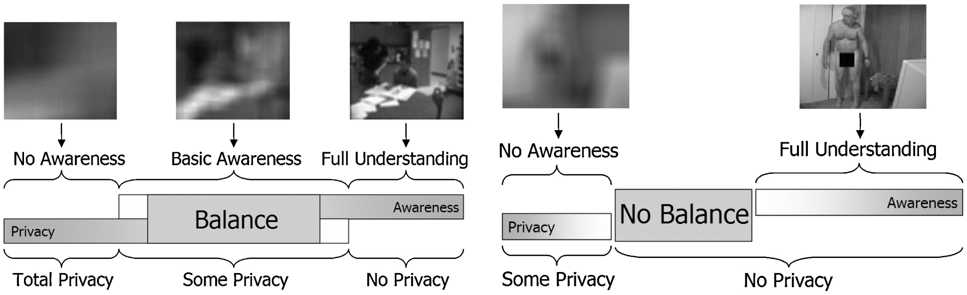

Fig. 2. The blur distortion filter can operate at a variety of levels. Each level affects fidelity and risk which, in turn, affect awareness and the ability to control confidentiality. The left part of the figure shows a mundane scene used in the Boyle et al. [2000] filter study in which privacy could be balanced with awareness. The right part of the figure shows a risky scene used in the Neustaedter et al. [2003] filter study in which privacy and awareness could not be balanced.

information has properties such as *persistency* and *transitivity* that are relevant to confidentiality. Information 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 [2001] and by Palen and Dourish [2003] 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 storage 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 and Dourish [2003] put it, a tension across temporal boundaries for performances. This results in the logical separation of the text of a performance and its context, that is, the audience. A common theme among theoretical frameworks for privacy is that typical privacy regulation behaviors assume people act in clearly situated local contexts and changes to this assumption can make the regulatory behaviors ineffectual.

6.3 Direct Controls

Mechanisms for regulating confidentiality overlap greatly with those for solitude, emphasizing their synergistic relationship. The principle means for confidentiality control involves 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 judge status and individuality [Schwartz 1968]. Territoriality and personal space use distance to afford fine-grained control over the access of others to our bodies and our things, for example, the notepad example in Section 2.1.2 [Luff and 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, yet obscure what is being said: for example, pig latin among children and hand signals in baseball.

Architecture also plays a vital role in the preservation of confidentiality (minimizing the leaks out) as well as the preservation of solitude (minimizing the 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, for example, 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 Section 2.1.1. These techniques, however, suffer from the feedback and control issues discussed in Sections 2.1.2 and 2.2.3.

6.4 Computers and Confidentiality

Increasingly, computers are being used to store or transmit confidential information and *computer security* holistically addresses many aspects of confidentiality. *Authorization* is control not only over access, but also use, that is, a person's intention for using the system or the information it provides, or outcomes of access. *Data integrity* concerns ensuring that persisted information about oneself is not modified, or that transmitted information is not modified enroute. Both of these are obviously part of confidentiality. *Process integrity*, 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 sideeffects. 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* (encryption) 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 rationalizing 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 defenses against confidentiality threats, but if the control is heavyweight or feedback inappropriate, these confidentiality-preserving features may actually interfere with usual processes and behaviors. 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 and 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.

6.5 Indirect Controls

People explicitly state (verbally or paraverbally) 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 other peoples’ secrets. People can intuit the sensitivity perceptions of others and from these infer self-imposed limits of behavior. In contract law, stiff penalties dissuade breeches of confidentiality. The law also enshrines confidentiality in certain relationships, for example, doctor/patient and lawyer/client, so 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 contact invested to build and maintain the relationship with that person. Palen and 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 and Dourish introduce the notion of *genres of disclosure* to capture not only institutional (socially constructed) expectations regarding confidentiality but also situational ones that change with the temporal boundary. The genres of disclosure are loosely defined, permitting feelings that privacy has been violated through the misappropriation or misuse of confidential information. With several case studies, they illustrate how even subtle situational differences can greatly change the genre of disclosure.

The risk/reward trade-off mentioned in Section 3.5 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—for example, "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 desirable 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. 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* is often used synonymously with autonomy. Self-appropriation, described earlier, and autonomy point to the same basic control—control over one's own behavior—yet, autonomy incorporates behaviors that facilitate *self-definition* and identity. Accordingly, Altman [1975] places great emphasis on the importance of self-definition and the role privacy plays in it. As suggested by Table I, autonomy and identity afford vital rewards for ego development. Many of the symptoms of privacy problems in video media spaces that were discussed in Section 2 can be blamed on the poor support of systems for managing behavior, identity, and impressions. Thus, an understanding of autonomy which regulates these things is needed to design a privacy-preserving VMS.

7.1 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 signaling 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 [1968] describes walls and doors as partitions that permit individuality. The violation of these barriers implies a loss of control over access to the self. 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 [1975] 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 organization's culture, in such cases, the social fabric of the organization 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 the 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 the worker of the opportunity to decide when to 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, as long as her/his actions conform to group expectations. Indeed, as part of the normal regulation of autonomy, one routinely adjusts their behavior to live cordially among others. Doing so ensures that longterm 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 revelation of informal awareness cues, contributing to the disclosure boundary tension [Palen and 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 [Boyle et al. 2000] 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, institutionalized 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 socialize 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 behaviors in communal spaces, for example, no smoking in restaurants. Constraints to autonomy are the primary means for punishing bad behavior: 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 the group's ability to govern 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 in Section 2.2.2. 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 bathrobe. 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 as she actually appears to others. Yet, this is only a partial solution because the mirror shows nothing about the norms that drive self-appropriation.

7.2 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, that is, 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 the confidentiality needed to perform socially unacceptable acts. Confidentiality is needed to preserve autonomy when others can use privileged information to thwart one's short and longterm 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 physical space, their pervasiveness makes it difficult for people to find opportunities to regulate solitude and thus limits the choices for autonomy to the point where they cannot do some desired behaviors because they are being watched.

Some important autonomy-related terms can be borrowed from Goffman's [1959] framework for self-presentation. People are *actors* who have *fronts* 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 nonverbal *signifiers*: *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 institutionalized and the audiences' expectations of a front become part of the front itself. Fronts are carefully constructed and maintained (e.g., by confidentiality) to ensure homogeneity between performances. The *back* is a secondary presentation of the self to the team only (for team fronts) or the individual her/himself. Here, deviance occurs and the self is maintained.

Bellotti's [1998] framework discussed in Section 2.2.3 focuses on the usability of video media space control and feedback affordances to support the kind of self-appropriation process developed by Goffman [1959], 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 Identity

We broaden our concept of autonomy to include control over *identity* and its expression, for example, 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 licenses, credit cards, and so forth are tangible artifacts that verify identity. These exist separately from physical presence and are held in a person's possession. Electronic equivalents include email addresses, personal Web pages, and network IDs. These make up part of the digital persona [Clarke 1994]. While there are legal safeguards to discourage others from mishandling conventional identity such as civil penalties for libel or unauthorized use of 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 their digital persona [Burkert 1998].

Identity is highly relevant to VMS design. Dissociation relates to identity because the virtual embodiments of people which signals presence and affords a means to interact with others and access information about them do not, unlike our physical presence, reveal identity. Computer security also relates to identity. *Impersonation* is the act of assuming the identity of another, usually without authority. *Identity theft* 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, minimizing 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].

7.4 Pseudonymity

A person is typically involved in a number of intersecting and disjointed social worlds. An identity is maintained for each world. Although we can recycle much of one identity for another, keeping distinct identities separate is a core privacy task. *Pseudnyms* 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 disclosure of identifying information. Video media spaces are at odds with pseudonymity because much identifying information is communicated in the video image of the 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.

7.5 Role Conflict

People often assume different *roles* 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* [Adler and Adler 1991] can result when previously nonoverlapping 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 nonmediated 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 (Section 2.2.1). Moreover, there is an analogue of role conflict for privacy norms: decontextualisation confuses which norms apply in a given circumstance [Palen and 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.

7.6 Focus and Nimbus

The tripartite concept 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 artifacts in the environment. Regulation of nimbi therefore roughly corresponds to confidentiality and autonomy. *Awareness*, which is a functional composition of focus and nimbus, is analogous to the dialectic negotiation of privacy boundaries.

Rodden [1996] 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 analyzing 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 nontrivial in part because the Rodden model does not account for some important topics:

—normalized, institutionalized 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.

8. CONCLUSION

This article 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 behavioral psychology. The vocabulary is grounded by empirical observations of privacy problems in video media spaces as a representative application domain where the design of social technology is known to raise a diverse spectrum of privacy concerns. Although we have tried to give this vocabulary a broad theoretical footing, scientific understanding of privacy in individual and social human life is still incomplete, and so too is our vocabulary.

8.1 Summarizing the Vocabulary

The vocabulary we have presented in this article is extensive and so it is difficult to find good ways to summarize it. Our vocabulary for privacy in VMS design describes a process that intends to regulate the interactions between a person and their physical and social environment. The process consists of three 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 with which others access information about oneself.

—Autonomy: control over one's own behavior 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.

Technology disrupts privacy regulation in a myriad of ways. Principally, technology lifts or changes environmental affordances and constraints for interactivity so that privacy regulating behaviors 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 the social self.

8.2 Questions to Guide Us to Privacy-Preserving Video Media Space Designs

The preceding summary listed only three words in a vocabulary consisting of dozens of highly interrelated terms. There is little benefit to be gained by repeating them all here as an enumerated list: such a list would show words out of context and ignore the deeper discussions of their subtleties and nuances. Rather, the specific goal of this article is to encapsulate and disseminate the understanding gained in assembling and discussing this vocabulary in depth. It contributes an important milestone towards guiding the design of privacy-preserving video media space because it exposes what could be evaluated in VMS design and implementation. Space does not permit us to address in this article many important design questions that could illustrate the utility of our vocabulary for deconstructing real-world privacy questions in VMS design, such as:

—Does the telephone model for establishing intermittent high-quality VMS links confound solitude management?

—Do *de facto* norms (stemming from slow, viral, grassroots adoption of the media space) violate autonomy? Do *de jure* norms (stemming from edicts passed by upper management in an organization) violate autonomy?

Rather than deal with specific issues, we have focused our contribution on providing the vocabulary to communicate the totality of privacy. To situate this vocabulary in the larger context of building a privacy-preserving video media space, we conclude this article by putting forth questions, the answers to which will help us develop tools and methods for evaluating support for privacy in a VMS design, outlining a course of future research in privacy and video media spaces.

First, there is an obvious need to be able to predict a design's effect on privacy at every stage of the iterative design cycle. Models describing the relationship between privacy and design are one way to aid prediction. These models are based not only on the theories that informed the vocabulary we presented, but also observations of a VMS design's effect on privacy. These effects need to be tracked while a VMS is in limited use and after it has become widely used.

—What effects could be or ought to be tracked?

—What observable metrics correspond to these effects?

—What tools (e.g., questionnaires) and methods (e.g., experimental protocols) can be used to elicit and measure these effects?

—What are the ethical guidelines for large- and small-scale experiments for understanding privacy?

The vocabulary we presented in this article yields some example effects and metrics that could be tracked.

- Degrees of freedom for controlling solitude, confidentiality, and autonomy.
- Effort (time, cognitive energy, and physical energy) spent regulating privacy.
- Violations permitted, their risk (probability, severity), conditions under which they arise, and their actual frequency of occurrence.
- Users' and nonusers' perceptions of these effects.
- Patterns of use, disuse, and misuse.
- Throughput, quality, enjoyment of social relationships and collaboration.
- Norms, taboos, and legalities of use that develop around deployment.

But considerably more work is needed to develop theoretically and empirically informed models of privacy and design. Specifically, the model must incorporate hypotheses about the relationship between design factors and privacy effects.

- What design factors are relevant to privacy?
- What is the relationship between a given design factor and users' and nonusers' capacities to control solitude, confidentiality, and autonomy?
- How does a given design factor affect rewards, risks, and violations?
- How does a given design factor affect society at large once the technology's use becomes a norm?

The vocabulary we presented facilitates the clear expression of such hypotheses and gives some examples of design factors that might be relevant.

- Modalities for interactivity that vary in demands to attention.
- Fidelities for information access that modulate sensitivity.
- Group interfaces to support policing.
- Single-user interfaces to support lightweight, fine-grained control.
- Communication and feedback channels to support dialectic negotiation of access to the self.

Of course, such hypotheses must be verified and so there will also be the need for experimental methods and protocols for the controlled study of the design-privacy link and techniques for field observation. The vocabulary we presented in this article facilitates unambiguous discussion of results and establishes a comprehensive theoretical background needed to interpret the results. However, verification requires more than just hypotheses and methods. The highly situational and subjective nature of privacy underscores the value of field observation complementing controlled experimentation. Such evaluations require VMS prototypes that incorporate designs and techniques that might possibly better support privacy. The vocabulary we presented is useful for articulating design ideas but is not yet useful for generating new ideas and does little to support their implementation. Privacy researchers have underscored the importance of iterative design which demands toolkits for rapidly constructing prototypes and iterating over their design quickly [e.g., Boyle and Greenberg 2002].

Although there has been a considerable body 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 supporting" means and 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 surely there is value in assembling a vocabulary to facilitate such discussion.

ACKNOWLEDGMENTS

We would like to thank Bonnie Nardi and our anonymous reviewers for their extremely helpful suggestions. Thanks also to Gina Venolia for introducing us to Adler and Adler's [1999] work and to Anthony Tang and Stacey Scott for their feedback on drafts of this article.

REFERENCES

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*.

Adams, A. 2000. Multimedia information changes the whole privacy ballgame. In *Proceedings of Computers, Freedom, and Privacy* (CFP 2000) (Toronto, Ont., Canada). ACM, New York, 25–32.

Adler, P. and Adler, P. 1991. *Backboards and Blackboards*. Columbia University Press, New York.

Altman, I. 1975. *The Environment and Social Behavior*. Brooks/Cole Publishing, Monteray, CA.

Altman, I. and Chemers, M. 1980. *Culture and Environment*. Wadsworth Publishing Company, Stanford, CT.

Angiolillo, J. S., Blanchard, H. E., Israelski, E. W., and 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, 51–74.

Arendt, M. 1958. *The Human Condition*, University of Chicago Press, Chicago, IL.

Balfanz, D. and Simon, D. 2000. WindowBox: A simple security model for the connected desktop. In *Proceedings of the 4th USENIX Windows Systems Symposium* (Seattle, WA), Advanced Computing Systems Association, 37–48.

Bellotti, V. 1998. Design for privacy in multimedia computing and communication environments. In Technology *and Privacy: The New Landscape*, P. Agre and M. Rottenberg, Eds. MIT Press, Cambridge, MA, 63–98.

Benford, S., Greenhalgh, C., Bowers, J., Snowdon, D., and 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, CO). ACM, New York, 242–249.

Boyle, M., Edwards, C., and Greenberg, S. 2000. The effects of filtered video on awareness and privacy. In *Proceedings of the Conference on Computer Supported Cooperative Work* (CSCW'00 Philadelphia, PA). ACM Press, New York, NY, 1–10.

Boyle, M. and Greenberg, S. 2002. GroupLab Collabrary: A Toolkit for Multimedia Groupware. In *ACM CSCW 2002 Workshop on Network Services for Groupware*, J. Patterson Ed. ACM, New York.

Brierley-Newell, P. 1995. Perspectives on privacy. *J. Environ. Psych. 15*, 87–104.

Brierley-Newell, P. 1998. A cross-cultural comparison of privacy definitions and functions: A systems approach. *J. Environ. Psych. 18*, 357–371.

Burkert, H. 1998. Privacy-enhancing technologies: Typology, critique, vision. In *Technology and Privacy: The New Landscape*, P. Agre and M. Rottenberg, Eds. MIT Press, Cambridge, MA, 125–142.

Buxton, W. A. S. 1997. Living in augmented reality: Ubiquitous media and reactive environments. In *Video-Mediated Communication*, K. Finn, A. Sellen, and S. Wilbur, Eds. Lawrence Erlbaum Associates Publishers, Mahwah, NJ, 363–385.

Clarke, R. 1994. The digital persona and its application to data surveillance. *The Inf. Soc. 10*, 2, 77–92.

Cool, C., Fish, R. S., Kraut, R. E., and Lowery, C. M. 1992. Iterative design of video communication systems. In *Proceedings of the Conference on Computer Supported Cooperative Work* (CSCW'92) (Toronto, Ont., Canada). ACM, New York, 25–32.

Coutaz, J., Bérard, F., Carraux, E., and Crowley, J. 1998. Early experience with the mediaspace CoMedi. In *Proceedings of the IFIP Working Conference on Engineering for Human-Computer Interaction* (EHCI98) (Heraklion). Kluwer Academic Publishers, Dordrecht, 57–72.

Crowley, J. L., Coutaz, J., and Bérard, F. 2000. Things that see. *Commun. ACM, 42*, 3 (Mar.), 54–64.

Dennet, D. 1995. *Elbow Room: The Varieties of Free Will Worth Wanting*. MIT Press, Cambridge, MA.

Dourish, P. 1993. Culture and control in a media space. In *Proceedings of the 3rd European Conference on Computer-Supported Cooperative Work* (ECSCW'93) (Milan, Italy). Kluwer Academic Publishers, Dordrecht, 125–138.

Dourish, P., Adler, A., Bellotti, V., and Henderson, A. 1996. Your place or mine? Learning from long-term use of audio-video communication. In *Computer Supported Cooperative Work: J. Collab. Comput. 5*, 1, 33–62.

Dourish, P. and Bellotti, V. 1992. Awareness and coordination in shared workspaces. In *Proceedings of the Conference on Computer Supported Cooperative Work* (CSCW'92) (Toronto, Ont., Canada). ACM, New York, 107–114.

Dourish, P. and 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) (Monteray). ACM, New York, 541–547.

Duval, S. and Wicklund, R. 1972. *A Theory of Objective Self-Awareness*. Academic Press, New York.

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*, J. Galegher, R. Kraut, and C. Egido, Eds. Lawrence Erlbaum Associates Publishers, Hillsdale, NJ, 351–371.

Fish, R. S., Kraut, R. E., and Chalfonte, B. L. 1990. The VideoWindow system in informal communications. In *Proceedings of the Conference on Computer Supported Cooperative Work* (CSCW'90) (Los Angeles, CA). ACM, New York, 1–11.

Fish, R. S., Kraut, R. E., Rice, R. E., and Root, R. W. 1992. Evaluating video as a technology for informal communication. In *Proceedings of the ACM/SIGCHI Conference on Human Factors in Computing Systems* (CHI'92) (Monteray). ACM, New York, 37–48.

Fitzpatrick, G. 1998. The locales framework: Understanding and designing for co-operative work. Ph.D. dissertation. The University of Queensland.

Gaver, W. 1992. The affordances of media spaces for collaboration. In *Proceedings of the Conference on Computer Supported Cooperative Work* (CSCW'92) (Toronto, Ont., Canada). ACM, New York, 17–24.

Gaver, W., Moran, T., MacLean, A., Lövstrand, L., Dourish, P., Carter, K., and 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) (Monteray). ACM, New York, 27–34.

Gavison, R. 1980. Privacy and the limits of law. In *Yale Law Journal 89*, 3 (Jan.), 421–471.

Goffman, E. 1959. *The Presentation of Self in Everyday Life*. Doubleday Publishers, Garden City, NY.

Greenberg S. and Kuzuoka, H. 2000. Using digital but physical surrogates to mediate awareness, communication and privacy in media spaces. *Pers. Tech. 4*, 1 (Jan.).

GREENBERG, S. AND ROSEMAN, M. 2003. Using a room metaphor to ease transitions in groupware. In *Sharing Expertise: Beyond Knowledge Management*. M. Ackerman, V. Pipek, and V. Wulf, Eds. MIT Press, Cambridge, MA, 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, WA). ACM, New York, 515–521.

GRUDIN, J. 2001. Desituating action: Digital representation of context. *Hum.-Comput. Interact. 16*, 2–4, 269–286.

GRUDIN, J. 1992. Groupware and cooperative work: Problems and prospects. In *Readings in Computer Supported Cooperative Work*, R. Baecker Ed. Morgan Kaufmann Publishers, 97–106.

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. *Computer Supported Coop. Work: J. Collab. Comput. 4*, 4, 297–318.

HARRISON, S. AND 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, MA). ACM, New York, 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. *ACM Trans. Internet Tech. 2*, 4, 276–306.

HUDSON, S. E. AND SMITH, I. 1996. Techniques for addressing fundamental privacy and disruption tradeoffs in awareness support systems. In *Proceedings of the Conference on Computer Supported Cooperative Work* (CSCW'96) (Cambridge, MA). ACM, New York, 248–247.

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

JUNESTRAND, S., KEIJER, U., AND TOLLMAR, K. 2001. Private and public digital domestic spaces. *Internat. J. Human-Comput. Stud. 54*, 5 (May), 753–778.

KELVIN, P. 1973. A social psychological examination of privacy. *Brit. J. Soc. Clin. Psych. 12*, 284–251.

KRAUT, R., EGIDIO, C., AND GALEGHER, J. 1990. Patterns of contact and communication in scientific research collaboration. In *Intellectual Teamwork: Social and Technological Foundations of Cooperative Work*, J. Galegher, R. Kraut, and C. Egido, Eds. Lawrence Erlbaum Associates Publishers, Hillsdale, NJ, 149–171.

LANGHEINRICH, M. 2001. Privacy by design—Principles of privacy-aware ubiquitous systems. In *Proceedings of UbiComp 2001* (Atlanta, CA). Springer, New York, 273–297.

LEE, A., GIRGENSOHN, A., AND SCHLUETER, K. 1997. NYNEX Portholes: Initial user reactions and redesign implications. In *Proceedings of the ACM/SIGGROUP Conference on Groupware* (GROUP'97) (Phoenix, AZ). ACM, New York, 385–394.

LUFF. P. AND HEATH, C. 1998. Mobility in collaboration. In *Proceedings of CSCW'98*. ACM, New York, 305–314.

MANTEI, M. M., BAECKER, R. M., SELLEN, A. J., BUXTON, W. A. S., MILLIGAN, T., AND 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). ACM Press, New York, NY, 203–208.

MOORE, G. 1997. Sharing faces, places, and spaces: The ontario telepresence project field studies. In *Video-Mediated Communication*, K. Finn, A. Sellen, and S. Wilbur, Eds. Lawrence Erlbaum Associates Publishers, Mahwah, NJ, 301–322.

NARDI, B. A., KUCHINSKY, A., WHITTAKER, S., LEICHNER, R., AND 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, 487–517.

NARDI, B., WHITTAKER, S., AND BRADNER, E. 2000. Interaction and outeraction: Instant messaging in action. In *Proceedings of the Conference on Computer Supported Cooperative Work* (CSCW 2000) (Philadelphia, PA). ACM Press, New York, 91–97.

NEUSTAEDTER, C. AND GREENBERG, S. 2003. The design of a context-aware home media space for balancing privacy and awareness. Report 2003-722-25, Department of Computer Science¸ University of Calgary.

NEUSTAEDTER, C., GREENBERG, S., AND BOYLE, M. 2003. Balancing privacy and awareness for telecommuters using blur filtration. Report 2003-719-22, Department of Computer Science, University of Calgary.

OLSON, M. H. AND BLY, S. A. 1991. The Portland experience: A report on a distributed research group. In *Computer-Supported Cooperative Work and Groupware*, S. Greenberg, Ed., Academic Press, New York, 81–98.

PALEN, L. AND DOURISH, P. 2003. Unpacking privacy for a networked world. In *Proceedings of the Conference on Human Factors in Computing Systems* (CHI 2003) (Ft. Lauderdale, FL). ACM, New York, 129–137.

POSNER, R. A. 1981. The economics of privacy. *The Amer. Econom. Rev. 71*, 2, 405–409.

REASON, J. 1990. *Human Error*. Cambridge University Press, New York.

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, MA.). ACM, New York, 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). ACM, New York, 25–38.

SAMARAJIVA, R. 1998. Interactivity as though privacy matters. In *Technology and Privacy: The New Landscape*, P. Agre and M. Rottenberg, Eds. MIT Press, Cambridge, MA.

SAMUELSON, P. 2000. Privacy as intellectual property? In *Stanford Law Review*, vol. 52. Stanford Univ. School of Law, Stanford CA, 1125–1174.

SCHWARTZ, B. 1968. The social psychology of privacy. In *Amer. J. Soc. 73*, 6, 741–752.

SIMMEL, G. 1964. The secret and the secret society. In *The Sociology of Georg Simmel*, K. Wolff Ed. Free Press, New York, 334.

SIMON, H. A. 1996. *The Sciences of the Artificial (3rd Ed.)*. MIT Press, Cambridge, MA.

SMITH, I. AND HUDSON, S. E. 1995. Low disturbance audio for awareness and privacy in media space applications. In *Proceedings of the 3rd ACM International Conference on Multimedia* (Multimedia 95) (San Francisco, CA). ACM, New York, 91–97.

SPIEKERMANN, S., GROSSKLAGS, J., AND 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, FL). ACM, New York, 38–47.

SUCHMAN, L. 1987. Plans and situated actions: The problem of human-machine communication. Cambridge University Press, Cambridge, MA.

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

TANG, C., MCEWAN, G., AND GREENBERG, S. 2003. A taxonomy of tasks and visualisations for casual interaction of multimedia histories. In *Proceedings of the Graphics Interface 2003* (GI 2003) (Halifax). Canadian Information Processing Society Mississauga, ON, and A K Peters Limited, Natick, MA.

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. *Int. J. Human-Comput. Stud. 42*, 5 (May), 501–530.

WHITTAKER, S., FROHLICH, D., AND 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, MA). ACM, New York, 131–137.

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

Received May 2003; revised February 2004; accepted July 2004 by Bonnie Nardi