A Lexicon for Privacy in Video Media Spaces

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 misuses, prompts ethical questions, and engenders grave concerns from both users and non-users. 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 much more deeply rooted than this, however. Privacy is an enormous concept and from it emerges an overwhelming torrent of interrelated words. In this article, we draw from resources in environmental psychology and CSCW to build a broadly- and deeply-rooted holistic description of this nebulous thing, privacy. Beyond this, we relate the vocabulary back to the real and hard problem of designing privacy preserving video media spaces. In doing so, we facilitate exploration and discussion of the privacy-design relationship.

Categories and Subject Descriptors: K.4.m [**Computers and Society**] Miscellaneous - *Privacy theory and models*; H.4.3 [**Information Systems Applications**] Communications Applications - *Video media spaces*; H.1.2 [**Information Systems Models and Principles**] User/Machine Systems - *Human factors*
General Terms: Privacy, Video media spaces, Social interaction, Computer Supported Cooperative Work (CSCW)
Additional Key Words and Phrases: Human-computer interaction, solitude, confidentiality, autonomy, environmental psychology, user interface design.

1. INTRODUCTION

Video media spaces (VMS) offer a small group of distance-separated collaborators with an always-on or always-available video channel that connects them. Through the video channel, people gain informal awareness of other's presence and their activities, which ultimately lead to 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., once or twice a minute [Dourish & 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 & Bly, 1991; Mantei et al, 1991, Gaver et al, 1992; Tang et al, 1994], as in Figure 1(a).

— Persistently open links between public spaces (personal offices, cafeterias, lounges), where the always-on camera is permanently broadcasting the information that is continuously displayed at distant sites [Fish et al, 1990; Jancke et al, 2001], as in Figure 1(b).

While video media spaces are a promising way to increase group interaction, the issue is that this technology is perceived by users and non-users alike to be privacy invasive and privacy insensitive. [e.g., Gaver et al, 1992; Bellotti & Sellen, 1993; Lee et al, 1997]. Researchers in the field [e.g., Bellotti, 1998] generally assume that *privacy problems in video media spaces arise because of the way they are designed, implemented, and deployed.* Although there is now a reasonable body of literature in Computer Supported

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

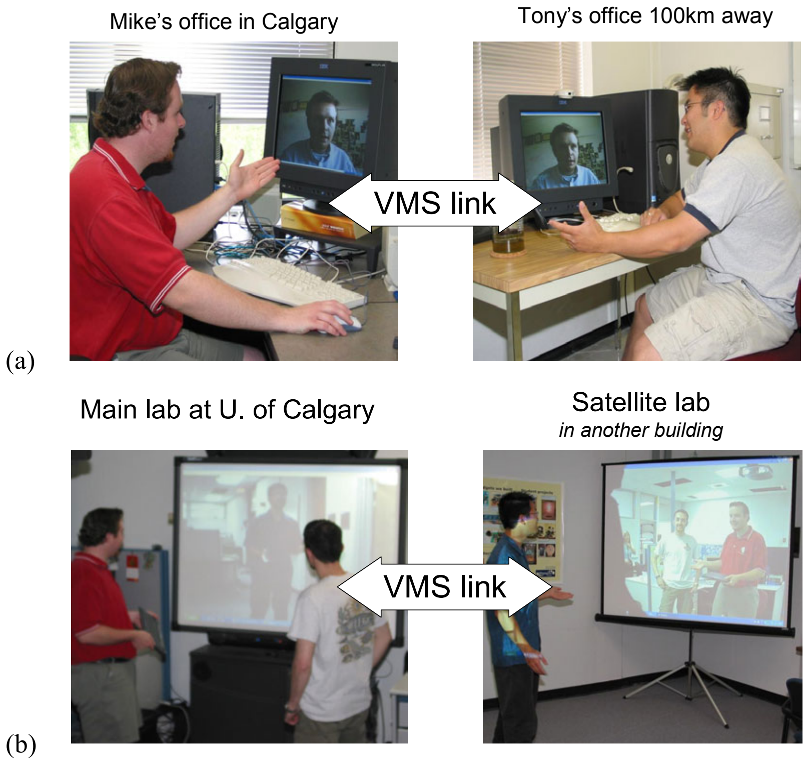

Figure 1. Two example configurations for a video media space. In (a), personal offices are connected using intermittent video links. In (b), public spaces are connected using persistently open video links.

Cooperative Work (CSCW)—and in particular video media spaces—that discusses the design problems found, the emphasis thus far has been on generalizing about the symptoms observed and then proposing specific countermeasures—point solutions—to offset specific symptoms. There has been some good discussion of the human and technical factors that may prompt privacy problems [e.g., Bellotti, 1998], but 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. Although this prior research provides a strong basis for our own work, it does not readily yield insight into how to diagnose privacy problems, predict when they will occur, or provide an intellectual foundation from which to generate new kinds of solutions.

We believe that privacy is a much larger and more interrelated concept than has been previously presented in the media space literature. We also believe that privacy problems in existing media spaces arise because *we as designers do not understand the real totality of privacy in a way that lets us see how our design choices will affect it.* To put this in context, the overall research objective of many researchers is to build a privacy-preserving media space. Yet, researchers cannot truly discuss privacy as applied to VMS design because even the basic vocabulary of privacy has become confused: different authors may use the same word to describe different phenomena, or the same author may use different words to describe the same phenomenon without relating the words to one another.

Thus, the goal of this paper is *to articulate a comprehensive lexicon for privacy in video media spaces*. We believe that this lexicon is needed to facilitate progress in designing privacy-aware video media spaces. We start building this lexicon by synthesising how privacy has been applied to VMS design by CSCW researchers (section 2). We broaden our lexicon in section 3 by looking at privacy from perspectives established outside the CSCW domain, such as anthropology, architecture, law, behavioural psychology, and sociology. Lastly, we deepen our lexicon in sections 4–6 by looking specifically at the theories of privacy developed in environmental psychology, and relate these theories to CSCW problems in designing a privacy-preserving video media space. While this article does not present specific solutions to privacy problems in video media spaces, it does satisfy our goal of creating a lexicon that will permit CSCW researchers to discuss privacy issues in video media space design in a much-needed holistic 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 started by *building* prototype systems and then *using* them. By building the technologies, they identified and overcame important technological roadblocks, but by living with the technology, researchers came to experience first-hand the privacy consequences of various design decisions and the symptoms of underlying problems. By carefully reflecting on their experiences, researchers came to intimately understand the relevant technological factors, individual human factors, and social factors. These symptoms (and their motivating factors) can be organised into three major themes.

— Deliberate privacy abuses are possible.

— Inadvertent privacy violations are possible.

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

We use the term "non-users" to refer to people who are present in the VMS but not necessarily users. For example, family members of a telecommuter using a VMS from home are non-users; also, we consider visitors and passers-by may appear in a VMS deployed in a personal office or open lab space—people who are not habitual occupants of the VMS—to be non-users.

In this section, we synthesise a broad compendium of VMS research to briefly review these three themes, the problems and design issues that each summarises, and the causal human and technical factors identified in prior work. While this synthesis is useful, we will reflect on why it is not sufficient for good media space design.

2.1. Deliberate Privacy Abuses

Even when participants in a video media space themselves may never willingly violate others' privacy the system affords the potential for such abuses. Worse, the systems afford the potential for *undiagnosed* abuse by non-participants. One example is surreptitious surveillance: e.g., a thief—or worse, a violent sex offender—intercepts the VMS video stream broadcast on the Internet, affording him the ability to monitor the presence and activity of others as he plots the perfect time to commit his crime. There is an implicit assumption in the literature 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.* The problem is that there is sensitive information in the media space and not everyone should have access to it.

2.1.1. Access control

One way to solve deliberate privacy abuses by "outsiders" is with *access control* i.e., put in place computer security and cryptographic measures to *deny* unauthorised individuals

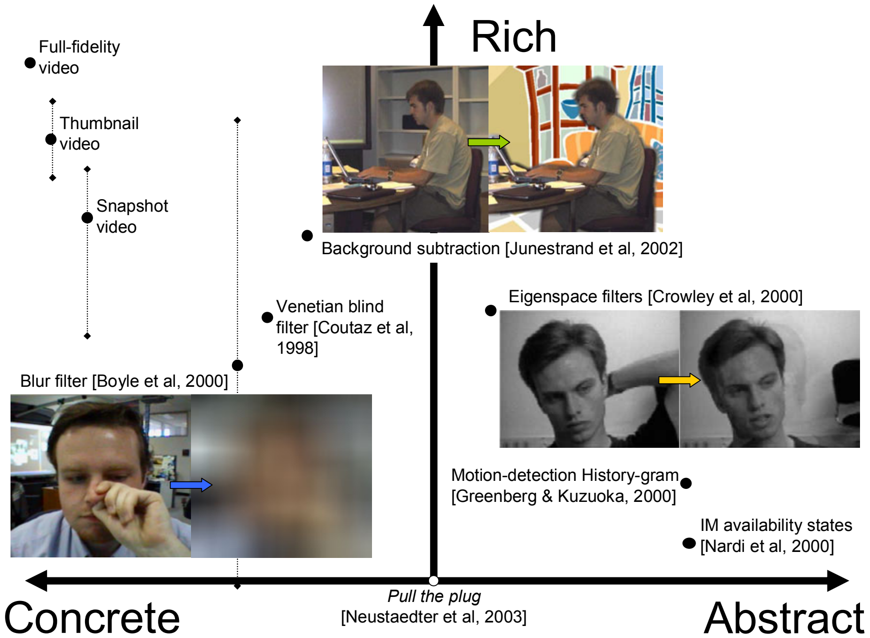
Rich
Full-fidelity video
Thumbnail video
Snapshot video
Background subtraction [Junestrand et al, 2002]
Venetian blind filter [Coutaz et al, 1998]
Eigenspace filters [Crowley et al, 2000]
Blur filter [Boyle et al, 2000]
Motion-detection History-gram [Greenberg & Kuzuoka, 2000]
IM availability states [Nardi et al, 2000]
Concrete
Pull the plug [Neustaedter et al, 2003]
Abstract

Figure 2. A design space showing some previously explored techniques for preserving privacy in video media spaces. The space is organised along two dimensions: presentation richness (the quantity of information content presented) and abstraction (how much of the original live video feed is presented).

access to sensitive information [Smith et al, 1995]. While access control is common on virtually all computers, those wishing to restrict access have faced a constant and unrelenting battle with those wishing to crack systems.

2.1.2. Content control

Another way to solve deliberate privacy abuses is to simply remove sensitive information from the media space, i.e., so there is nothing of worth for "outsiders" to access and so that little harm will result if access control measures are cracked. We call this technique c*ontent control*, but it is hard to put this technique into practice because the purpose of a media space is to *reveal* [Gaver et al, 1992]. More precisely, the function of a media space is to capture and present awareness information of others. Thus there is a fundamental trade-off between privacy and the utility of VMS for awareness: for one person in the media space to have richer awareness, others must have necessarily less privacy [Hudson & Smith, 1996]. Figure 2 shows several CSCW techniques for preserving privacy in video media spaces based on content control. *Distortion filters* (e.g., the blur filter in Figure 2) mask sensitive details in video while still providing a low-fidelity overview useful for awareness [Zhao & Stasko, 1998; Boyle et al, 2000]. *Publication filters* (e.g., the background subtraction filter in Figure 2) remove details considered unimportant for awareness information [Coutaz et al, 1998; Junestrand et al, 2001]. Also, one can remove potentially privacy-threatening details by abstracting away from the video altogether (e.g., the eigenspace filter in Figure 2) [Crowley et al, 2000].

2.1.3. Lightweight vs. fine-grained control

Both the above approaches involve control i.e., 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 is affected with 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. This kind privacy regulation demands very little cognitive or behavioural effort from the people involved and is usually an implicit activity realised as a natural consequence of the other activities.

There are few lightweight strategies for controlling a video media space. While unplugging the camera may be simple, effective, and appropriate when explicit action is needed to cope with dramatic changes in the environment for privacy, 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, configuring this information is usually heavyweight i.e., control is usually afforded through a complex panel of GUI widgets and if-then-else scripted access rules. As a result of this effortful process, people often do not make changes when appropriate, and often end up configuring the system to grant all others either full access at any time, or no access whatsoever. Unfortunately, these behaviours thwart the security of the system and open it up to deliberate privacy abuses.

2.1.4. Dissociation and the physics of social interactions

Dissociation, where one's actions become cut off (dissociated) from one's identity is another critical problem in video media space design [Bellotti, 1998]. Dissociation makes it very difficult for VMS users to determine *who* is accessing information about them even though they may be able to tell that it is being accessed. Thus deliberate privacy abuses are easy because one knows they can access information about others 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.

Moreover, the "physics" of social interaction in co-located settings have changed little in human history, and it is the predictability of the physical environment that permits control of social interactions [Dennet, 1995], such as with norms, taboos, and laws. Technology introduces new kinds of opportunities for multimedia information capture, use and abuse at a pace many times faster than the physical environment. Consequently, some deliberate privacy abuses are permitted because the social infrastructure needed to prevent them cannot keep pace [Langheinrich, 2001].

2.2. Inadvertent Privacy Violations

A fundamental premise of the cognitive sciences is that people are mostly rational [Simon, 1996]. It is largely assumed in the CSCW literature on privacy and video media spaces that the same principle applies to privacy: people will usually protect their own privacy and respect the privacy of others. Not all privacy violations are deliberate, and not all opportunities for deliberate privacy abuses are capitalised upon. Yet, accidental violations are known to happen from time to time. In this section, we will explore inadvertent privacy infractions because they are subtle, complex, and arise because *media space designs fit poorly with individual human and social factors* [Bellotti, 1998].

2.2.1. Disembodiment and VMS feedback cues for self-appropriation

Self-appropriation is a regulatory process where people modify their behaviour and appearance according to social norms and expectations [Bellotti, 1998]. Self-appropriation depends on contextual cues people sense from the environment: specifically, the culture of 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, i.e., to fit in. 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 perceptions of norms and expectations for social behaviour, and adapt their behaviour accordingly. The impoverished nature of a video media space means that people often do not appropriate themselves correctly for viewing by distant colleagues. Disembodiment—where a user becomes cut off from the (multiple) contexts of those people viewing him—confounds self-appropriate and leads to inadvertent privacy violations [Bellotti, 1998]. A person is entirely dependent on the VMS to feed context cues back to her in order to determine how she should behave. Yet the design of this feedback channel is fraught with technical factors that permit inadvertent privacy violations

It is hard to balance VMS feedback cue salience and distraction [Gaver et al, 1992; Hudson & Smith, 1996; Bellotti, 1998]. Feedback cues must be salient when needed (so people notice them) and unobtrusive when not (so people are not distracted) [Hudson & Smith, 1996]. 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—nullifying its salience and fostering disembodiment—or disable the VMS altogether, nullifying any benefit of its use.

It is hard to design and implement 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, yet these are hard to accomplish when designing a media space. By way of contrast, 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. Inadvertent privacy violations occur in video media spaces because social protocol for managing interactions is compromised. For example, using a telephone-style ring to signal a request for conversation generally doesn't reveal identity of the individual initiating the request. Even when different rings are used for different people, the association is not nearly as concrete as footsteps. Buxton's DoorCam [Buxton, 1997] situates the VMS camera and display at the display 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.

2.2.2. Places and culture for privacy

Place—its architecture and use—is an important feedback cue for self-appropriation (e.g., locker room versus boardroom). In the physically mediated environment, there is usually a physical transition when one moves between two places supporting distinct privacy cultures: a partition, a doorway, and even distance itself [Altman, 1975; Paylen & Dourish, 2003]. The time needed to navigate the transition affords opportunity to assess the resulting change in cultural expectations and make changes in appearance and behaviour as appropriate. On the other hand, m*edia spaces join places with differing privacy cultures and do so without such smoothing transitions*. 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 lack strategies for coping with the privacy problems this juxtaposition engenders. Second, without the transition, people are unaware of the juxtaposition and its impact on self-appropriation.

2.3. Apprehension

Non-users are often suspicious of the video media space and its handling of their privacy and even go out of their way to sabotage the system [Jancke et al, 2001]. Even users themselves are often leery about the system's handling of their privacy [Tang et al, 1994]. Although VMS design permits a number of ways that a user's privacy could be violated—inadvertently and deliberately—previous work on privacy in VMS design reveals other factors prompting this apprehension. Specifically, users are nervous about making "bad impressions" in the media space.

2.3.1. Impressions and surveillance

A fundamental premise of privacy research in VMS design is that people do not want to "look bad" in front of others—especially co-workers—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 or behaving in ways that others might judge unacceptable, e.g., caught on camera in an office media space while changing clothes after jogging during lunch caught on camera in a home office media space while spanking a disobedient child.

Users are apprehensive about making mistakes that make them "look bad" in the media space [Tang et al, 1994]. Since video media spaces permit detailed, surreptitious surveillance at *any* time, users must monitor their appearance, behaviour, and speech at *all* times [Lee et al, 1997]. Coping with surveillance requires *vigilant* self-monitoring, which can lead to errors [Reason, 1990], i.e., "looking bad." Worse, VMS technology affords new abilities for automated surveillance and *rigorous scrutiny*. In addition to making a bad impression with other people, people can now worry about making a bad impression with cold, socially-inept computer algorithms.

2.3.2. "Out of context"

When short segments of a conversation are viewed independent of its totality, listeners are forced to invent information needed to support its interpretation (context) and the invented context can make the speaker "look bad." Moreover, video media spaces afford new operations on multimedia information that permit taking others' speech and actions out of context.

Persistence and Retransmission. Technology makes speech and actions which 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. In other words, technology affords new degrees of temporal and spatial freedom for information access [Paylen & Dourish, 2003]. 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].

Modification. 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 e.g., computer security, cryptography, and the filtration methods described in section 2.1.2. Human factors, on the other hand, weigh heavily in problems related to inadvertent privacy violations, especially the interplay between human and technical factors. There are fewer concrete, generalised technological countermeasures for dealing with specific 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.

In our own research on designing a privacy-preserving video media space, our efforts came to a halt because we lacked an understanding of the psychology of privacy. Thus we could not consider the intersection of human behaviour, social behaviour, the environment, and technology and how this nebulous mass of complex phenomena interoperates. We even lacked a vocabulary to name the phenomenon and describe their interactions. These are complex concepts and hard problems, not often dealt with by computer scientists: they are problems that behaviour psychologists, sociologists, social psychologists, and environmental psychologists consider. Unfortunately, these other domains tend to focus on global problems related to human social communities and individual human development. Rarely are the practical aspects of building and designing computing systems given first-tier treatment in psychological theory of privacy. Somehow, we must integrate the two.

In the remainder of this article, we will focus on presenting a comprehensive lexicon for privacy in video media spaces that integrates that which we in the CSCW community have learned about the privacy-design problem with the rich understanding gained in the psychological sciences. Again we caution that we will not propose solutions to the hard problems related to deliberate and inadvertent privacy violations and apprehension. In fact, we will introduce even more problems into the fray. Meaningful solutions are still some distance away.

3. PERSPECTIVES ON PRIVACY

Many disciplines of study must deal with the notion of privacy: anthropology, architecture, behavioural psychology, law, sociology, and more recently computer science. The vocabulary we build for discussing the human factors relevant to the privacy-design link in media spaces draws from these varied areas, although we will pay special attention to Altman's [1975] theories of privacy established in environmental psychology. However, we begin with a broad overview of various themes in privacy research by drawing from Brierley-Newell's [1995] cross-disciplinary survey of privacy-related literature.

3.1. "Private" versus "Public"

"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 definition the most fundamental and broadly cross-cultural conceptualisation of privacy. There are nuances to this. Being apart is different from being alone, e.g., 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. This illustrates a tension between one wanting/needing/choosing/being private *vs.* 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 heighten 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.

3.2. Privacy as an Attribute of Places and People

In architecture, privacy is often seen to be contingent upon features of the design and construction of architectural space: e.g., the number of enclosing partitions, their height, the glazing that makes the space visually porous, and the intelligibility of human speech and loudness of other noises passing through walls and openings. Thus, privacy is a subjective phenomenon coupled to the person perceiving it. Treating privacy as an architectural attribute of space is important for VMS design. It permits construction of architectural metaphors [Greenberg & Roseman, 2003] for privacy safeguards, and it informs us that *some* aspects of privacy can be quantified as observable metrics.

Yet, there are other privacy metrics that are not so easily quantified. In particular, architecture not only defines a space, but it creates a social *place* full of social meaning [Harrison & Dourish, 1996] which in turn determines its privacy. For example, public toilets are not very private in construction but are nonetheless very private in the sociological experience of their use. This fact has definite implications for video media space designs designers: how people perceive privacy seems to be subtle, subjective, and social. Yet technology has, historically, had a hard time handling (i.e., observing and quantifying) phenomena that exhibit these properties. Furthermore, it is still 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 a Process

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

— *verbal*, e.g., telling someone across the VMS link to keep some information secret;

— *para-verbal*, i.e., non-verbal, e.g., pointing a VMS camera out the window; or,

— *social*, e.g., deciding, as a group, that it is taboo to turn on the VMS camera in the kitchen when the person who turned it off is still present (even though the system may allow it).

One perspective of privacy identified in Brierley-Newell's survey is that these behaviours 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: his body, thoughts and personality, and information about him. The negation occurs between the self and the *environment*: the physical environment and also the social environment i.e., the people immediately nearby and society at large.

Altman's privacy process is a *dialectic*. The actual level of privacy attained is decided through a process of negotiation between the self and the environment. This dialectic—back and forth—is *normative*. People's desired privacy starts out high and is subsequently constrained by the environment to socially accepted (normal) levels. Since one is typically involved in many groups simultaneously, there may be a number of norms that apply in a given situation. 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 or her own set of privacy *preferences* i.e., "personal norms," that determine one's initial desired privacy level and that subsequently influence the privacy dialect. The relationship between (group) norms and (individual) preferences seems complex and co-adaptive; in particular, group norms change in response to changes in group membership.

Altman's privacy process does not *deny* interactions between the self and the environment, rather, it *regulates* them. When one has too many interactions i.e., "too little privacy" these interactions are throttled: e.g., a person turns off the media space to get away from others. When the connections with others have been cut so much that one has "too much privacy" the privacy process opens access to the self so that a person gets the interactions he craves: e.g., 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 nebulous 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 will be difficult to design context-aware systems that can adapt to changes in the environment affecting the internalised privacy process. However, we can likely develop qualitative methods that allow us to observe this process, which in turn can help us evaluate the effectiveness of particular media space designs. To this end, Altman's theory holds potential heuristic value: because it has been specified so broadly, it can apply to many situations. Yet, Brierley-Newell speculates that this very broadness makes Altman's theory the most criticized. For example, some critics argue that social interactionalism may be better able to explain the privacy process, e.g., Fitzpatrick's Locales framework [Fitzpatrick, 1998] that applies social interactionalism principles to uncover and comprehend CSCW system design issues. However, this and other frameworks have yet to be applied to the design-privacy relationship in video media space design.

3.4. Privacy as a Need, Right, and Freedom

Researchers in behavioural psychology have studied individuals who routinely experience compromised privacy, such as the elderly and the mentally infirm living in institutions, and young children. They characterise the outcomes of failures in the privacy process that yield harmful effects, and a few of these effects are listed in Table I. These extreme effects don't apply to the general population, who are able to enjoy many benefits from a healthy amount of privacy. Some of these benefits are given in Table I. Perhaps because of these benefits, people place great value upon privacy in our society.

Consequently, privacy is often defined as a legal and moral right and as an inalienable freedom that no other person or organization may lawfully or morally unduly curtail. Privacy is thus legally enshrined in various laws to: discourage "peeping toms," prevent unjustified search, seizure, and confinement, punish slander and liable, and ensure contractual obligations to secrecy. This fact has relevance for science: Kelvin [1973] in particular discusses barriers to the scientific study of privacy, e.g., when so much value is placed upon privacy, the scientific manipulation of it for experimentation (needed to understand it) is seen as "morally suspect."

A privacy that is a right or freedom can be *violated*: others' actions may deny one this right or impair one's exercise of it. Specifically, it is a privacy violation when others' actions prevent one from obtaining the privacy he needs, he *normally* enjoys, and society deems that he *ought* to enjoy. Others' actions may prevent one from obtaining

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

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]).

desired privacy, but this itself may not necessarily be considered a privacy violation. Privacy violations have *outcomes*: for example, the effects of too much or too little privacy discussed in the previous sub-section. These outcomes vary in *severity,* a subjective measure of how "bad" this harm is. Although the environment may permit others actions that will lead to a privacy violation, they might not choose to invoke such actions; hence, privacy can be *threatened.* Privacy threat and privacy *risk* are used almost synonymously and seem to include both the *probability* that a violation will occur, and the severity of the harm it causes. Opportunities for violation are held in check by *policing*: providing punishments, taboos, etc., to discourage others from doing things that violate one's privacy. Some privacy violations are so severe that one is permitted actions to stop further harm and be awarded damages to offset harm already done. Given that privacy violations arising from the deliberate and inadvertent misuse of video media space technology may be inevitable, our designs could also support policing and recovery from violations in addition to providing safeguards to constrain misuse.

3.5. Privacy as a Balancing Act

People put their privacy at risk as they venture out and interact with the world. Aside from hermits and the like, people balance the benefits accrued from social interactions against the risks to privacy, engaging and withdrawing from others to satisfy both the need to be "apart" and the need to be "together." Even though there is risk, there may also be *reward*: i.e., benefits to having less privacy than may be possible. In most human activities reward exists commensurately with risk, yet many video media space designs ignore this relationship altogether.

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 so it is very questionable if they accrue any benefit from their participation [Neustaedter & Greenberg, 2003]. In many designs, people are together by the system in an indiscriminate way that disregards the need (or lack thereof) for social interaction [Fish et al, 1990; Fish et al, 1992; Greenberg & Rounding, 2001; Jancke et al, 2001]. Furthermore, many video media space designs permit some form of surreptitious *surveillance*: the ability to closely monitor the environment—usually the presence and activities of others—without revealing much about oneself. For example, in the CAVECAT media space [Mantei et al, 1991] a user could cover the camera lens to prevent others from seeing him and yet still see others. In both cases, the video media space design itself fosters or permits *disparity* between risk and reward such that reward does not accrue accordingly with risk or, conversely, risk does rise with reward.

Reciprocity [Root, 1988] is often enforced over video media space channels as a technological means for re-balancing this threat/benefit disparity. Yet, reciprocity does not always hold for the physical environment, and sometimes breaking the reciprocity rule is beneficial. For example, you can observe another to deduce his *availability* (willingness to engage in interaction) without disturbing him by moving quietly and peeking around the corner of an open office doorway. The RAVE media space [Gaver et al, 1992], for example, did not strictly enforce reciprocity. What this means is that while people strive for balance in the media space, the technology itself, the ways it can be subverted, and the awkwardness of the interface may hinder their ability to achieve it.

3.6. Summary: Focusing on a Process Model for Privacy

Our premise in this section is that in attempting to understand the complex human factors involved in privacy and the design of video media spaces, it is important to consider the different phenomenological perspectives on it that have been cultivated in disciplines such as anthropology, psychology, and sociology. What we have seen is that privacy is:

— a basic human need,

— a quality of people and places, and

— a behavioural process governing interactions that seeks to balance risks and rewards associated with social interactions.

These perspectives on privacy can be integrated. Privacy involves various aspects of the physical environment, human psychology, and social behaviour for in the maintenance of self and the regulation of social interactions. Of the perspectives offered, the one offered by 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 characterises privacy as a boundary-control process regulating access to the self. Specifically, he describes three genres of control, each of which is directly relevant and immediately applicable to understanding the problems we face in designing and building a privacy-preserving video media space.

— *Solitude* relates to understanding how people regulate social interactions. It applies since video media spaces are designed to support such interactions.

— *Confidentiality* relates to understanding how people manage access to information about themselves. It applies since deliberate and inadvertent privacy violations in video media spaces occur because of the way they handle such multimedia information.

— *Autonomy* relates to understanding how people choose to present themselves in social situations. It applies since VMS design confounds self-appropriation.

What are controls? Dennet gives a technical description: "*A* controls *B* if… *A* can drive *B* into whichever of *B*'s normal range of states *A* wants *B* to be in" [Dennet, 1995]. Gavison [cited in Brierley-Newell, 1995] points out two elements in control: the ability to make a choice (implying that a number of alternatives exist to select from) and the power to ensure the choice is respected.

Altman specifies that privacy controls—of which, solitude is an example—are *social*: privacy exists whenever human-human social relationships exist. As soon as social interactions (of casual or work topics) are made possible—be it by spatial propinquity or by technological mediation—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, the media space must also afford opportunities for people to "be apart" when they need it, to affect social relationships in intended ways. This intentionality is important: 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 participants may not have cameras) those participants without cameras are perceived to by spying on the community [Coutaz et al, 1998].

Privacy is not only a social phenomenon; it is also a *co-operative* one. A person will sometimes do things that help others respect his own privacy. While privacy violations occur regularly, gross privacy violations seem to not occur as often as the environment permits. Sometimes group members take advantage of opportunities to violate the privacy of other group members, but most often they do not. Given that group privacy is contingent upon the privacy of its individual members, it may be the case that in some groups, members take steps to *protect* the privacy of other group members and defend against outside intrusion. For example, even if I do not get a chance to close my office door before my lawyer calls me regarding a sensitive topic, my colleagues may sense my privacy needs and close my door for me. This cooperative view of privacy differs markedly from the competitive view common in computer science, which assumes that if opportunity for a privacy violation arises, it will necessarily be capitalised upon. This more extreme competitive view may be for theoretically 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 change her settings on her behalf, loosing out on opportunities to defend against inadvertent privacy violations.

In the next three sections we will delve deeply into each of these three genres of control—solitude, confidentiality and autonomy—to build our integrated lexicon for privacy. Each discussion starts broadly, with particular emphasis placed on the human behaviour, psychological and sociological concepts related within the genre 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.

4. SOLITUDE

Altman describes solitude as a control over interactions between the self and the environment, particularly other people. Solitude 'controls' help a person "be apart" from others and is involved in many behaviours that are vital to human development, e.g., self-evaluation and ego development [Altman, 1975]. We clarify that being apart is different from being alone: for example, two lovers can find solitude in each other's company, even in a crowded restaurant. "Togetherness" is thus a continuum of states, and the extremes present failure conditions that yield negative behavioural, psychological, and physiological responses. For example, *crowding* 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.

4.1. Attention and Distraction

To discuss other issues in video media spaces that closely relate to solitude, we generalise Altman's definition of solitude to include control over where one directs one's *attention* and how one controls *distraction*. Most video media spaces require people to expend extra effort if they are to attend the information within them, or they present this information ways that potentially distract or disrupt people. Both cases affect solitude, which is why we expand Altman's definition. This extension also helps to explain "camera shyness" problem in video media spaces [Lee et al, 1997]. In co-located settings, people track the focus of others' attention as an informal awareness cue that helps determine availability. In particular, a person notices if another is looking at her, i.e., that she is becoming the object of others' attention. This prompts her to reflexively focus her own attention back upon herself, to monitor self-appropriation and track others' impressions. This state of heightened self-awareness can cause discomfort if maintained for prolonged durations [Duval & Wicklund, 1972].

4.2. Verbal and Para-Verbal Solitude Controls

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

4.3. Westin's Four Privacy States

Westin, another noted privacy theorist, decomposed privacy into four "states" [Westin 1967] all of which are, in fact, states of Altman's privacy process that relate to the exercise of the solitude control.

— *Solitude* is a state of total isolation. (Note that Westin uses the word differently from Altman.)

— *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 co-present with others and yet not expect to be recognized by them and so 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.)

4.4. 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 solitude.

Conversely, when one craves social stimulation, one can go to places where others are. Place partially determines accessibility, i.e., the effort people must expend to engage

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

Table II. Example interpersonal distances and the modalities of interaction supported at each [Hall, 1966].

others for interaction [Harrison & Dourish, 1996]. People can control attributes of their physical environment as part of solitude regulation: architectural spaces can often be reconfigured to raise or lower their permeability to light, matter, and sound. In doing so, people control the affordance of architectural space for interactivity. For example, an office door can be closed to reduce visual and auditory distractions from the corridor and serve as a physical barrier to others' entry. Not surprisingly, doors have been used as metaphors for regulating solitude in video media spaces [e.g., Gaver et al, 1992; Buxton, 1997]. Moreover, doors can be anywhere from fully closed to slightly ajar to wide open, and that this becomes a social cue indicating one's solitude desires. In contrast, video media spaces generally provide only one modality for interactivity (an audio/video channel) and offer few ways to configure this channel to signal the desired level of engagement.

People can also capitalise upon the ambiguity inherent in some architectural changes to regulate solitude. For example, a closed door ambiguously symbolises both absence as well as a wish to be left undisturbed [Root, 1988]. Office lights left on at night may lead co-workers to consider that one is working late. People also capitalize on ambiguity in computer-mediated environments. For example, Nardi et al [2000] reports that people use the inaccuracies of the presence indicators as form of "plausible deniability," where they ignore requests for conversation from people because they know that the other person will be uncertain if they are really there.

4.5. Personal Space

We are coming to a picture of solitude in which space and social behaviour interoperate. *Personal space* refers to an invisible boundary in space "attached" to a person, separating him from others. Although the boundary's size and shape is never made explicit and also varies from moment to moment as part of the privacy dialectic, people show definite behavioural 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 artefacts and will be discussed in the next section on confidentiality.

Personal space regulates solitude by reducing sensory stimulation—and, therefore, attention—due to the presence of or interactions with others. 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 given in Table II. Each of these modes for interaction carries a social meaning as well: a pat on the back and a "thumbs up" gesture may both be used to pass along congratulations for a job well done, but have different social meanings. Because of this relationship between distance and interaction, distance itself becomes imbued with social meaning [Altman, 1975]. For example, consider when one person sits down at the same table as another: when the newcomer sits kitty-corner and out of direct eye contact, he sends a solitude-related message that differs markedly from when he chooses to sit directly across the person and in easy eye contact.

Personal space, as a tool for solitude regulation, depends on having a *range* of interpersonal distances at which people may space themselves. These distances define modalities for interaction that differ in both affordances for interaction and the level of engagement (attention) needed to sustain such interactions. These distances establish characteristics of social interactions, and are thus imbued with social meanings. Typically, in a video media space the camera position and display size dictates the visual distance between people; these are sometimes arbitrary and do not represent the desired social distance. For example, seeing a tightly cropped face shot on a large video monitor places someone visually close, but the actual mannerisms exhibited by that person may reflect actions of someone who is in fact quite far away. The concept of interpersonal distances 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: connected with everyone; connected with just one other person; and, disconnected from everyone. The limited choices for connectivity make the media space a crude tool for the selective expression of social interest for interactivity. Moreover, in physically co-located settings, adjusting distances is very lightweight and can be continuously adapted—just by moving around. In contrast, media spaces offer highly discrete (i.e., not continuous) choices selected using heavyweight GUIs and limit degrees of freedom, e.g., it is awkward to reposition the VMS camera because of limited cable lengths, lighting, shelf space, and similar factors.

4.6. Summary

Solitude—which we define as control over attention for social interactions—is an important concept for designers of tools to support casual interactions, particularly video media spaces. Yet, we have discussed several ways in which people's existing strategies for mediating solitude in physically co-located settings are confounded by VMS design. Typical VMS designs fail to afford a diverse range of modalities for interactivity and lightweight means to select and re-select from among them. As a result people are unable to appropriately signal interpersonal distance in social space. This "crippled social mobility" in this sense prompts psychological discomfort, distraction, and social awkwardness.

5. CONFIDENTIALITY

Confidentiality is the control of access to information about oneself, e.g., informal awareness cues, intentions, vital statistics, thoughts and feelings, medical history, criminal record. "Controlling access" is as much granting access as it is restricting it. *Secrecy* is similar to confidentiality but narrower because secrecy emphasizes that the information is *concealed* from certain people. Secrecy modulates the *disclosure* of information to others, but this is only one aspect of confidentiality. There is a kind of congruence between confidentiality and solitude. Confidentiality directly regulates the outward flow of information, whereas solitude indirectly regulates the inward flow of information. Similarly, confidentiality indirectly regulates the (inward) attention of others, whereas solitude directly regulates one's own (outward) attention.

5.1. 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 (similar to [Adams, 2000]). Others' impressions of a person are predicated upon their knowledge of her, and so

Parameter	Measurement	Description
Field of view	radians (degrees)	How much of the space is visible
Resolution	pixels/metre	The minimum size of object discernable
Frame rate	frames/second	How often frames are captured; determines motion smoothness
Codec quality	%, bits/pixel, kilobits/second	How much information is discarded during compression; determines precision and accuracy
Latency	seconds	Temporal delay between video capture on one side of the link and presentation on the other side
Jitter	seconds/frame	Variance in latency between successive frames in a motion video sequence

Table III. Parameters for video media space fidelity.

confidentiality is part of impression management [Goffman, 1959]. The harms that could arise from breeches of confidentiality include embarrassment, damage to ego and identity, and harms that arise from the loss of others' esteem (e.g., 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 compromising act.

5.2. Fidelity

Fidelity can be defined as 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 one places in the information, or the error in its perception). The same essential truth or description of circumstance may be perceived at a variety of fidelities. Also, people's perceptions of the fidelity of information about a person are situated in the context of the whole history of social interaction with that person. Information about oneself—the object of confidentiality—may be known by individuals—e.g., friends, colleagues, strangers—at different fidelities. We can broaden Altman's definition of confidentiality to address many common cases when we consider that *confidentiality includes control over fidelity*.

Video media spaces have several dimensions for video fidelity: field of view, resolution, frame rate, codec quality, latency and jitter. These are defined in Table III. The upper bound to the fidelity of most of these is limited by technology, and these upper 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 often fixed because the cameras typically lack pan/tilt/zoom capabilities or because they are difficult to move in practice.

As discussed in section 2.1.2, many video media space designs try to preserve confidentiality by discarding fidelity. They assume that it is the image details that are sensitive and therefore low fidelity "overviews" of the video pose less risk [Hudson & Smith, 1996]. For example, distortion filters such as the blur filter shown in Figure 3 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 it is masking the *sensitive* information. For example, Neustaedter et al [2003] questioned the effectiveness of a smoothing (i.e., blur) 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.

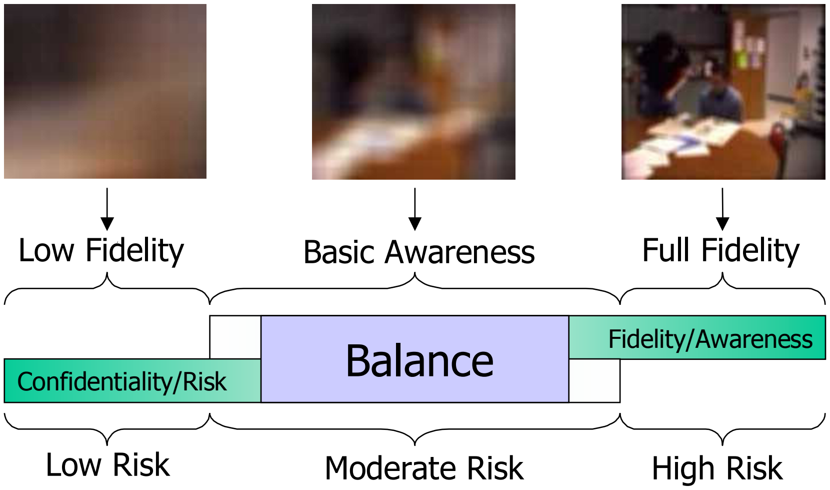

Figure 3. The blur distortion filter can operate at a variety of levels. Each level affects fidelity and risk, which in turn affect awareness and one's ability to control confidentiality.

The perceived fidelity of information is not static. In particular, it may change when it is transmitted between people, such as through oral or written statements. Hence, our broadened notion of confidentiality also involves the regulation of the fidelity of information that third parties may transmit about us. In addition to fidelity and sensitivity, we can also consider that information has properties such as ephemerality (persistency) and transitivity that are relevant to confidentiality. Persisting data that is otherwise fleeting increases its perceived fidelity. Receiving a data transmission may increase the perceived fidelity of information (especially if it was previously not known). The perceived fidelity of received information is influenced, for example, by the trust one places in the sender and the number of recipients. For example, imagine that Mary and Joe participate in a media space that archives the video streams and Mary *thinks* she saw Joe passionately kiss someone who is definitely not his wife. If the video was not archived, Mary would be left with lingering doubts, but archival changes the persistency of the information and permits scrutiny which yields a more accurate (i.e., higher fidelity) view of the event.

5.3. Direct Controls

Mechanisms for regulating confidentiality overlap greatly with those for solitude, emphasising their synergistic relationship. The principle means for confidentiality control involve keeping our bodies, possessions, and thoughts accessible to some (for communication) but inaccessible to others. We must consider possessions because things like diaries and driver's licenses reveal a great deal of sensitive information about a person. Territoriality and personal space (i.e., "going elsewhere") use distance to afford fine-grained control over others' access to our bodies and our things, e.g., the notepad example in section 2.1.3. 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. Curiously enough, the same technique is also used to preserve the solitude of others, e.g., we whisper at the cinema because we do not want to disturb others. Private vocabularies can be used to talk openly among others yet obscure what is being said: e.g., "pig latin," knowing looks between intimates, and hand signals in baseball.

Architecture also plays a vital role in the preservation of confidentiality (minimizing leaks out) as well as the preservation of solitude (minimizing leaks in). We surround our spaces with walls, which reduce access via visual and auditory channels. We draw window blinds and close doors to preserve visual confidentiality; we sound-proof walls and erect noise barriers to preserve aural confidentiality. In contrast to our subtle and lightweight confidentially controls in the real world, controlling confidentiality in the media space is awkward, and is usually achieved only through direct explicit controls, e.g., turning down microphone volume so as not to be overheard, or encoding information through cryptographic methods so others cannot eavesdrop (see next section). The filtration techniques discussed in section 2.1.2 also provide direct control of fidelity to minimise the risk of deliberate and inadvertent confidentiality breaches. Of course, "pulling the plug"—disconnecting from the media space by unplugging cameras and network cables—is a very effective, albeit crude, technique universally used by VMS users to guarantee confidentiality. Moreover, the unplugged cables are a reliable and easily understood cue as to how the system is treating one's confidentiality. Sadly, most VMS designs—nay, most hardware/software interface architecture designs—treat this kind of unexpected device disconnection or removal to be a fatal error condition, rather than a standard mode of use.

5.4. Computers and Confidentiality

Increasingly, computers are being used to store confidential information and *computer security* holistically addresses many aspects of confidentiality. *Authorization* is control not only over access, but also *use* i.e., a person's intention for using the system or the information it provides, or outcomes of access. *Data integrity* concerns ensuring that persisted information about oneself is not modified or transmitted information is not modified en-route. Both of these are obviously part of confidentiality. *Process integrity*, *availability, responsiveness,* and *reliability* concern perform their intended function when requested correctly and completely in an expected amount of time, and produce no undesired side-effects. Process integrity is an important component of confidentiality because, as stated in the introduction to this section, confidentiality includes granting people access and ensuring they have all the access which they've been granted.

Cryptographic methods, such as *encryption*, are used to provide access control, as described in section 2.1. These techniques can also be used to verify the identity of the receiver of information, or the identity of the sender of information and the integrity of the message (e.g., with *digital signatures*). Sadly, computer users often deliberately circumvent access controls or 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 they visit and what music tracks they play.

Computer systems afford defences which an individual may use to control his privacy; however, if the control is heavyweight, these privacy-preserving features may actually interfere with the normal privacy process. For example, security measures are often incomprehensible to set up and use, or they require great effort, or they do not supply sufficient feedback for people to know what is actually being transmitted [Balfanz & Simon, 2000]. Because this often stops people from doing their basic work, people often thwart computer security measures. Thus instead of (say) 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, we cannot yet predict how these script-like UIs will be used when deployed widely.

5.5. Indirect Controls

At times, people explicitly state (verbally or para-verbally) their confidentiality desires and perceptions on information sensitivity: e.g., one tells another person to "Keep this secret, okay?" Confidentiality is also preserved by simply not talking or by revealing less information than may be possible, i.e., exercising reserve in speech. In contract law, stiff penalties dissuade breeches of confidentiality. The law also enshrines confidentiality in certain relationships, e.g., doctor/patient and lawyer/client, such that desirable limits are placed on the judiciary's access to information obtained from questioning such confidants. Clearly, a host of social and psychological mechanisms control confidentiality indirectly. These controls probably will not map so cleanly onto VMS user interface features, but are nonetheless powerful and there is benefit to discussing them.

Telling a person that it is important to keep a piece of information secret doesn't prevent that person from revealing it to others. Yet, people can choose to—and sometimes do—keep others' secrets. In addition to explicit signals or requests, people can intuit others' sensitivity perceptions and from these infer self-imposed limits to behaviour. Information about others, including confidentiality preferences, are usually revealed over time i.e., one gets to know another better with each subsequent interaction. Access to information about a person accrues with the amount of social "work" invested to build and maintain the relationship with that person. Breaching a close colleague's confidentiality could foster distrust that might break down the relationship. Preserving privacy thus allows one to reap the *rewards* of social interaction; the denial of these rewards can act as a psychological mechanism for conforming to another's confidentiality desires. As mentioned before, there is a natural trade-off between risk and reward in social interaction and this is very apparent for risks to confidentiality: disclosure is an important part of building close relationships. Also, there are ways that VMS design can break the process of trading off risk for reward in social encounters.

Obvious caveats to these claims—e.g., "blabbermouths"—exist but these do not detract from their generality. While people *can* keep secrets or assess sensitivity, a particular individual may not keep a secret *well,* or may ultimately choose not to respect the apparent sensitivity. 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). We expect that technological safeguards against these kinds of confidentiality violations will never be perfectly effective, and so it is important to incorporate into the VMS design various awareness and interaction channels that can be used to diagnose, police, and reprimand these kinds of violations. Of course, peculiarities of the video media space may change the rules of engagement. For example, a VMS might record video/audio exchanges for later replay (this is sometimes seen as a good thing by those interested in meeting capture). People may not know or they may forget this, and consequently their request to another to keep information confidential is meaningless as neither has control over the information. Alternatively, the very fact that information is recorded means that one person cannot indirectly control one's confidentiality desires, since verbally telling the other does not preclude others from listening in later.

5.6. Summary

Thus, for a video media space to afford control over confidentiality, it must provide a choice of fidelities for information disclosure, where differs in sensitivity. It must also provide lightweight and fine-grained UI for adjusting fidelity, and feedback channels to ensure the fidelity accessed matches desires. Lastly, the environment must be stable enough to predict the social outcomes of one's confidentiality and information access choices. These requirements speak nothing of how *well* the VMS will support confidentiality, however.

6. AUTONOMY

Collectively, the freedom to choose how one interacts with the world (freedom of will) and the power to act in such a way are taken as the third privacy control: *autonomy.* Self-appropriation, described earlier, and autonomy point to the same basic control—control over one's own behaviour—yet, autonomy incorporates behaviours that facilitate *self-definition*, or, more broadly, *identity*. 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 systems' poor support for managing behaviour, identity and impressions. Thus, an understanding of autonomy—which regulates these things—is needed to design a privacy-preserving VMS.

6.1. 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. Also, one can communicate to others how important it is to us that he be allowed to do precisely as he wishes, explicitly through the content of speech, and implicitly through affect revealed in the structure of spoken language, facial expressions, and in posture. Generally, informal awareness cues simultaneously reveal one's autonomy desires.

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 own will—both autonomy and our legal entitlement to it take part in a dialectic based on group norms. Each may do as he wishes, so long as his actions conform to group expectations (e.g., do not tread upon the rights and freedoms of others). Indeed, as part of the normal regulation of autonomy, one routinely adjusts one's behaviour so that one may live cordially among others. Doing so ensures that long term plans come to fruition even if they are not done strictly as planned. This is the process of *self-appropriation*: one modulates one's behaviour and appearance to conform to group expectations of it. Thus, autonomy controls are *constrained* rather than compromised by group norms. Even so, 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. In other words, Saul's availability becomes a constraint that helps Mike regulate his autonomy and also Mike's solitude (since solitude is as much fostering interactions as it is denying them). This example underscores that in video media spaces, *privacy can be preserved by the judicious reveal of informal awareness cues.* This idea contradicts prior work that played privacy and awareness off each other in direct opposition [e.g., Boyle et al, 2000].

Other people can also constrain a person's autonomy. Some ways benefit the individual or society. For example, institutionalized people often incur great losses in autonomy [Altman, 1975 citing Goffman, 1959]: drugs or physical restraints are used to prevent injury to themselves, staff, or other residents; and, these patients often have no control over the scheduling of daily activities, such as when to awake, sleep, eat, bathe, use the toilet, etc. More generally, autonomy is constrained to enforce social protocol. Parents often restrict the autonomy of their young children to keep them safe and *socialise* them (teach them how to behave properly in society). Barriers are erected to restrict access to dangerous places or places where confidentiality is demanded or prohibit certain behaviours in communal spaces: e.g., no smoking in restaurants. Constraints to autonomy are the primary means for punishing bad behaviour: adults who commit crimes are incarcerated; children who disobey their parents are "grounded."

Video media spaces affect autonomy by taking away some constraints on behaviour, and changing some constraints so that differ importantly from those that exist in the physical world. For example, media spaces allow people to transcend geographic constraints on observation and interaction. As a result, media spaces provide rewarding opportunities for remote collaboration, but at the same time permit the problems relating to space and privacy that were discussed in section 3.2. Video media spaces often do not erect barriers to constrain users' autonomy so that they do not violate group norms. For example, a media space that connects home and corporate users is generally unable to switch its cameras off if the home worker appears in a bath robe (i.e., inappropriate attire). Furthermore, disembodiment design issues raised in previous research are linked to the modification of real-world constraints for behaviour. Embodiments provide feedback that tells a person how he appears in the environment, particularly relative to other people. Disembodiment, as discussed throughout sections 2 and 3, obscures the feedback that is a useful source of constraints. Placing a mirror next to the camera intends to remedy this problem by showing a person how she actually appears to others. Yet, this is only a partial solution because the mirror shows nothing about the norms that drive self-appropriation.

6.2. Autonomy-Confidentiality-Solitude Symbiosis

A second way in which autonomy is like the "muscle" of privacy regulation is that it provides the "power" that people have to enact their privacy choices, i.e., to control information access and direct attention for interactions. Thus, 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 opportunity to perform socially unacceptable acts (e.g., picking one's nose) where one "disobeys in private." During these times, one gains the strength needed to "obey in public" [Brierley-Newell, 1995]. Confidentiality is also needed to preserve autonomy, for example, when others can use privileged information to thwart one's short- and long-term plans. (In subsection 6.3, we will complete this discussion by describing how confidentiality is needed to protect identity.)

Because of the symbiotic relationship between solitude, confidentiality and autonomy, when a VMS design impairs the regulation of one kind of control, the other two may also be negatively affected. For example, when cameras are ubiquitously embedded into every corner of our physical space, their pervasiveness makes it difficult for people to find opportunities to be apart from others (i.e., regulate solitude) and thus limits choices for autonomy; they cannot do some desired behaviours because they are being watched.

6.3. Identity

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

Identity is highly relevant to VMS design. In previous research, dissociation has been cited as a major factor contributing to privacy problems. Dissociation relates to identity because the virtual embodiments of people—which signal presence and afford means to interact with others and access information about them—do not, unlike our corporeal bodies, reveal identity. Concepts from computer security also relate. *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 safeguards 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 his control over 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 invites 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 [Crowley et al, 2000] can make one person unintentionally appear as another.

6.4. Pseudonymity

A person is typically involved in a number of overlapping and non-overlapping social worlds. One maintains an identity for each such world, though we can recycle much of one identity for another if members between social worlds largely overlap. Keeping distinct identities separate is thus a core privacy task. *Pseuodnyms* are alternate identities which one creates and uses for interactions with the 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 an identity to the others. This is markedly different from impersonation, which involves using someone else's (pre-existing) identity. 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 part of one social world is also part of another or otherwise communicates with members of it. Also, some telecommunication technologies permit anonymity, i.e., allow one's interactions with the environment to proceed in a way that limits the reveal of genuinely identifying information of a person. Video media spaces are somewhat at odds with pseudonymity because so much identifying information is communicated in the video image of one's face and body. While video manipulation techniques could, conceivably replace a person's real visage with an artificial one, such algorithms are tricky to implement in practice, require considerable setup for creating replacement images for multiple identities, and likely reduce the value of the video channel for expressive communication.

6.5. 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 their children, a lover when with their mate, and a slob when alone at home. *Role conflict* [Adler & Adler, 1991] can result when previously non-overlapping social worlds collide and one is forced to assume two previously distinct roles simultaneously, exposing each to people whom one would rather not. The classical example of role conflict in the non-mediated environment is when parents go to visit their son at his college dormitory: the son must simultaneously play the role of a "child" and an "adult."

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 participants will likely inhabit quite different physical contexts. This is 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 they are alone at home and forget they are connected). By virtue of connecting two physically disjoint spaces (each embodying their own, possibly different sets of privacy norms) we create opportunities for role conflict. Furthermore, when the home worker is forced by situation to act as parent in the presence of office colleagues it fosters opportunity for an inadvertent privacy violation and contributes to the apprehension users and non-users feel towards the media space.

6.6. Summary

Autonomy is the process of regulating behaviour and expressing identity. It is symbiotically and synergistically related to solitude and confidentiality. Like all privacy controls, autonomy operates as part of a social, normative dialectic. The environment (particularly the social environment) affords norms that strengthen autonomy paradoxically by constraining it. Previous research has revealed dissociation and disembodiment design issues as key factors prompting the symptoms of privacy failures discussed in section 2. Both contribute to problems with self-appropriation, impression management, while more conventional computer security design issues concern identity theft. We also see that video media spaces encourage role conflicts by forcing a person to simultaneously exist in more than one social setting. Our vocabulary reveals that media spaces change constraints and these changes affect behaviour. Also, the VMS offers practically no facilities to police the space or reprimand violators: *the single user interface to a social technology (the video media space) eliminates social governance of its use.*

7. CONCLUSION

Our over-arching goal is to produce an implementation of a video media space that is proven to preserve privacy. Yet, we reached a point in the early stages of our work where we recognised that we lacked sufficient knowledge of privacy itself; thus, we could not really tackle the problem as we did not understand it. In particular, we lacked a sufficiently broad and deep vocabulary for articulating privacy concerns within VMS design in a clear and unambiguous way. To assemble this vocabulary, we have integrated theories of privacy developed in environmental psychology with ideas and observations of privacy and design developed in CSCW. We then applied this integrated lexicon to selected design problems that arise in video media spaces.

7.1. Summarising the Lexicon

Our lexicon for privacy in VMS design describes a process that intends to regulate the interactions between a person and her physical and social environment. The process consists of three kinds of controls.

— **Solitude:** control over social interactions, specifically control over the allocation of attention for interaction and engagement. *Related words from the lexicon*: distraction; "camera shyness;" intimacy; anonymity; refuge; place; space; culture; availability; accessibility; personal space; interpersonal distance; and, salience (of informal awareness cues).

— **Confidentiality:** control over information access, specifically control over the fidelity at which particular individuals access particular pieces of information about oneself. *Related words from the lexicon*: secrecy; reveal; sensitivity; access control; content control; persistency; transitivity; territoriality; computer security; authorisation; cryptography; trust; and surveillance.

— **Autonomy:** control over one's own behaviour and the expression of identity. *Related words from the lexicon*: self-appropriation; constraints; socialisation; impersonation; identity theft; authentication; pseudonymity; role conflict; dissociation; disembodiment; and, impression.

Our lexicon also includes terms that unify these three kinds of controls into a cohesive, integrated process that broadly considers not only aspects in behavioural and social psychology, but also architecture, law, and computer science. *Related words from the lexicon*: public; need; freedom; right; threat; severity; risk; reward; violation; harm; policing; repair; reprimand; control; choice; power; synergism; symbiosis; mutuality; reciprocity; apprehension; deliberate abuse; inadvertent misuse; norms; preferences; and, context.

7.2. Questions to Guide Us to Privacy-Preserving Video Media Space Designs

The specific goal of this article is to encapsulate and disseminate the understanding gained in assembling this lexicon. It contributes an important milestone towards guiding the design a privacy-preserving video media space because it exposes what we should evaluate 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 lexicon 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?

— Does automatic logging of VMS conversations violate confidentiality?

— 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 organisation) violate autonomy?

— What is the theoretical rationale for distortion filters or publication filters as privacy-preserving mechanisms?

— What insights will be gained when we apply the lexicon to comprehend privacy in other domains, such as email, instant messaging, data mining, surveillance, mobile computing, ubiquitous computing, context-aware computing, art, and online games?

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

First, we need to predict a design's effect on privacy at every stage of the iterative design cycle. To make predictions about privacy and design, we need to build a model describing the relationship. To build the model, we need to be able observe and track a design's effect on privacy while it is in limited use and after it has become widely used.

- What kind of effects do we need to track? Our lexicon gives some examples:
  - degrees of freedom for controlling solitude, confidentiality, and autonomy;
  - effort (time, cognitive energy, and physical energy) spent regulating privacy;
  - users' and non-users' perceptions;
  - violations permitted, their risk (probability, severity), conditions under which they arise, and their actual frequency of occurrence;
  - patterns of use, disuse, and misuse;
  - effects on social relationships, outcome of collaborative work (throughput, quality, enjoyment); and,
  - norms, taboos, and legalities of use that develop around its deployment.
- 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 ethical guidelines for large- and small-scale experiments for understanding privacy?

Next, we need to consider the relationship between design and privacy. We can conceptualise the set of all possible VMS designs as a multidimensional space, where we term each dimension a design factor. A design factor is a decision in the design process where we must make a choice from an enumerable set of options. We will need to construct hypotheses about the relationship between these design factors and their effects on privacy.

- What design factors are relevant to privacy? Our lexicon gives some examples: modalities for interactivity that vary in attention; fidelities for information access that vary in sensitivity; group interfaces to support policing; single-user interfaces to support lightweight, fine-grained control; and, communication and feedback channels to support dialectic negotiation of access to the self.
- What is the relationship between a given design factor and users' and non-users' 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?

This is only the beginning. We need to verify the privacy-design hypotheses we generate, and so we will need to design experimental methods and protocols for the controlled study of the design-privacy link, as well as techniques for field observation. The lexicon we presented in this article will help us facilitate careful articulation of hypotheses and the results of verification. We will also need prototypes of strategies and techniques we think might possibly preserve privacy. We can draw examples from existing literature, but we also hope that this lexicon—coupled with advances in context-aware computing—will yield new and better techniques. To realise them, we require toolkits for rapidly constructing these prototypes and iterating over their design quickly, such as our Collabrary toolkit for multimedia groupware [Boyle & Greenberg, 2002]. Our hypotheses, verified experimentally, will complete our model of privacy with axioms we can use to understand what "privacy-preserving" means in the first place. This model, when complete, will drive not only the design of a privacy-preserving video media space, but also establish that it does, in fact, preserve privacy.

8. ACKNOWLEDGEMENTS

This research was funded in part by NSERC and Microsoft Research. A special note of thanks is made to Anthony Tang and Stacey Scott for their insightful and invaluable feedback on drafts of this article. Thanks also to Gina Venolia for introducing us to Adler & Adler's work on role conflict, and the students of the Interactions Lab at the University of Calgary for their assistance producing figures for this article.

9. REFERENCES

ADAMS, A. 2000. Multimedia Information Changes the Whole Privacy Ballgame. In *Proceedings of Computers, Freedom, and Privacy* (CFP 2000, Toronto), ACM Press, New York, NY, 25–32.

ADLER, P., AND ADLER, P. 1991. *Backboards and Blackboards*. Columbia University Press, New York, NY.

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, Stamford, 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.

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), 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.

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 Press, New York, NY.

BRIERLEY-NEWELL, P. 1995. Perspectives on privacy. In *Journal of Environmental Psychology*, 15, Academic Press, New York, NY, 87–104.

BRIERLEY-NEWELL, P. 1998. A cross-cultural comparison of privacy definitions and functions: A systems approach. In *Journal of Environmental Psychology*, 18, Academic Press, New York, NY, 357–371.

BURKERT, H. 1998. Privacy-Enhancing Technologies: Typology, Critique, Vision. In *Technology and Privacy: The New Landscape*, P. Agre 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. In *The Information Society*, 10:2. Taylor and Francis, New York, NY, 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), ACM Press, New York, NY, 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. In *Communications of the ACM*, 43:3 (March), ACM Press, New York, NY, 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 Third European Conference on Computer-Supported Cooperative Work* (ECSCW'93, Milan), 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: The Journal of Collaborative Computing*, 5:1, Kluwer Academic Publishers, Dordrecht, 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), ACM Press, New York, NY, 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 Press, New York, NY, 541-547.

DUVAL, S., AND WICKLUND, R. 1972. *A theory of objective self-awareness*. Academic Press, New York, NY.

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), ACM Press, New York, NY, 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 Press, New York, NY, 37–48.

FITZPATRICK, G. 1998. *The Locales Framework: Understanding and designing for co-operative work*. Ph.D. thesis, 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), ACM Press, New York, NY, 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 Press, New York, NY, 27–34.

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. *Personal Technologies*, 4:1 (January). Elsevier.

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), ACM Press, New York, NY, 515–521.

HALL, E.T. 1966. *Distances in Man: The Hidden Dimension*. Double Day, Garden City, NY.

HARPER, R.H.R. 1996. Why People Do and Don't Wear Active Badges: A Case Study. In *Computer Supported Cooperative Work: The Journal of Collaborative Computing*, 4:4, Kluwer Academic Publishers, Dordrecht, 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). ACM Press, New York, NY, 67-76.

HIXON, R. 1987. *Privacy in a public society: Human rights in conflict*. Oxford University Press, New York, NY.

HUDSON, S.E., AND SMITH, I. 1996. Techniques for Addressing Fudamental Privacy and Disruption Tradeoffs in Awareness Support Systems. In *Proceedings of the Conference on Computer Supported Cooperative Work* (CSCW'96, Cambridge), ACM Press, New York, NY, 248-247.

JANCKE, G., VENOLIA, G.D., GRUDIN, J., CADIZ, JJ, 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), ACM Press, New York, NY, 530-537.

JUNESTRAND, S., KEIJER, U. AND TOLLMAR, K. 2001. Private and public digital domestic spaces. In *International Journal of Human-Computer Studies*, 54, 5 (May), Academic Press, New York, NY, 753–778.

KELVIN, P. 1973. A Social Psychological Examination of Privacy. In *British Journal of Social and Clinical Psychology*, 12, Swets & Zeitlinger, Lisse, 284-251.

KRAUT, R., EGIDIO, C., 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), Springer, New York, NY, 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), ACM Press, New York, NY, 385–394.

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., 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). ACM Press, New York, NY, 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, NY, 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), ACM Press, New York, NY, 129-137.

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

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), ACM Press, New York, NY, 25-38.

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 Third ACM International Conference on Multimedia* (Multimedia 95, San Francisco), ACM Press, New York, NY, 91-97.

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), ACM Press, New York, NY, 23–34.

TANG, C., MCEWAN, G., AND GREENBERG, S. 2003. A Taxonomy of Tasks and Visualisations for Casual Interation of Multimedia Histories. To appear in *Proceedings Graphics Interface 2003* (GI 2003, Halifax). Canadian Information Processing Society Missasauga, 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. In *International Journal of Human-Computer Studies*, 42:5 (May), Academic Press, New York NY, 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), ACM Press, New York, NY, 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), ACM Press, New York, NY, 11–18.