# Unpacking "Privacy" for a Networked World

**Leysia Palen**
Department of Computer Science
University of Colorado, Boulder
Boulder, CO 80309
palen@cs.colorado.edu

**Paul Dourish**
School of Information & Computer Science
University of California, Irvine
Irvine, CA 92697
jpd@ics.uci.edu

## ABSTRACT
Although privacy is broadly recognized as a dominant concern for the development of novel interactive technologies, our ability to reason analytically about privacy in real settings is limited. A lack of conceptual interpretive frameworks makes it difficult to unpack interrelated privacy issues in settings where information technology is also present. Building on theory developed by social psychologist Irwin Altman, we outline a model of privacy as a dynamic, dialectic process. We discuss three tensions that govern interpersonal privacy management in everyday life, and use these to explore select technology case studies drawn from the research literature. These suggest new ways for thinking about privacy in socio-technical environments as a practical matter.

### Keywords
Privacy, surveillance, monitoring, access regulation, boundary management, disclosure, social psychology

## INTRODUCTION
In an increasingly networked world, privacy protection is an ever-present concern. New technologies and infrastructures, from pervasive Internet to mobile computing, are being rapidly introduced and woven into the fabric of daily life; devices and information appliances, from electronic picture frames to digital video recorders, carry with them new possibilities for information access, and so for privacy management. Human-Computer Interaction (HCI) researchers have long acknowledged the implications their designs have for personal privacy. Indeed, the synergy between the technologists and social scientists who belong to that community, and the related computer-supported cooperative work (CSCW) community in particular, has led to mutual appreciation of the interdependent relationship between technology and situations of use. This, in turn, has heightened awareness of the privacy concerns that novel technologies introduce.

However, despite broad concern, there have been few analytic or systematic attempts to help us better understand the relationship between privacy and technology. We recognize when our systems introduce "privacy issues," but we have few tools for understanding exactly what those issues are. Privacy regulation is complicated, and has a range of functions, from maintaining comfortable personal spaces to protecting personal data from surreptitious capture. Both social and design studies of technology often unknowingly conflate these functions, and consequently fail to provide satisfying analytical treatment.

We hope to provide researchers and practitioners with a better understanding of "privacy" by unpacking the concept so that more specific statements can be made vis-à-vis technology. We do this by building upon privacy regulation theory developed by social psychologist Irwin Altman [6, 7]. Altman is primarily concerned with how people manage face-to-face interactions; we extend it to consider the lessons for information technology analysis and design. We then apply these concepts in case analyses.

While traditional approaches understand privacy as a state of social withdrawal, Altman instead sees it as a *dialectic* and *dynamic boundary regulation process* [6]. As a *dialectic* process, privacy regulation is conditioned by our own expectations and experiences, and by those of others with whom we interact. As a *dynamic* process, privacy is understood to be under continuous negotiation and management, with the *boundary* that distinguishes privacy and publicity refined according to circumstance. Privacy management is a process of give and take between and among technical and social entities—from individuals to groups to institutions—in ever-present and natural tension with the simultaneous need for publicity. Our central concern is with how this process is conducted in the presence of information technology.

### Common Concerns
Technology and personal information is haunted by the specter of Big Brother, with its implications of invasive and subversive action.[1] When privacy is discussed abstractly, concerns about surveillance and personal identity theft are among the most prominent. Certainly, these are important and pressing concerns. However, in studies of



*CHI 2003*, April 5–10, 2003, Ft. Lauderdale, Florida, USA.



[1] The idea of "Big Brother," drawn from Orwell's dystopian vision in *1984*, is often used to refer to the idea of pervasive monitoring and recording of activity, often by some central authority. Two things are important to note, however. First, in *1984*, the actual threat is of *potential* monitoring; "there was of course no way of knowing whether you were being watched at any moment" [22:6]. This is also true of Bentham's Panopticon, a metaphor Foucault and others have used. Second, the threat lies in the culture of pervasive mutual monitoring, rather than centralized surveillance; Winston Smith's friend Parsons is proud to be turned in not by his telescreen but by his children.

information technology in mundane and pervasive activity like video conferencing, shared calendar management and instant messaging communications, concerns most salient to users include minimizing embarrassment, protecting turf (territoriality) and staying in control of one's time. Although Big Brother actions may threaten life and liberty, it is interpersonal privacy matters that figure primarily in decisions about technology use on an everyday basis.

Our most familiar ways of managing privacy depend fundamentally on features of the spatial world and of the built environment, whether that be the inaudibility of conversation at a distance, or our inability to see through closed doors. We can also rely on others to honor behavioral norms around physical touch, eye contact, maintenance of interpersonal space, and so on [6].

With information technology, our ability to rely on these same physical, psychological and social mechanisms for regulating privacy is changed and often reduced. In virtual settings created by information technologies, audiences are no longer circumscribed by physical space; they can be large, unknown and distant. Additionally, the recordability and subsequent persistence of information, especially that which was once ephemeral, means that audiences can exist not only in the present, but in the future as well. Furthermore, information technology can create intersections of multiple physical and virtual spaces, each with potentially differing behavioral requirements. Finally, in such settings, our existence is understood through representations of the information we contribute explicitly and implicitly, within and without our direct control. These concepts of *disclosure*, *identity* and the shifting expressions and implications of these in *time* are central to our analysis, and we will return to them later in the paper. Information technology has shifted and complicated privacy regulation practices by creating numerous possible consequences from our computer-mediated interactions.

### Related Research

Our treatment builds upon the work of those who have walked this path before. A number of HCI researchers have turned their attention to privacy concerns in modern information environments. In the domain of ubiquitous computing, Bellotti and Sellen [8] reflected on user experiences in a pervasive digital environment that combined computer, audio and video networking with individual tracking and control technologies. They identified common problems that arise when the site of someone's activity and the site of its effect are separated, as can often happen in these environments. Grudin has suggested that threats to privacy as a result of these technologies might be more fundamentally explained as the "steady erosion of clearly situated action," and that our control over how disclosed information is interpreted in different contexts and times is diminished or absent [18]. Dourish [12] also investigated questions of privacy that arose in a range of media space environments [9] and pointed in particular to the organizational situatedness of appropriate solutions. Clement [11] broadly explored the privacy concerns raised by these technologies, paying particular attention to the institutional norms that govern forms of participation and control.

Agre [2, 3, 4, 5] has written extensively on privacy concerns and new technologies. In particular, he has critically examined technical discourse surrounding privacy and information technology, in an effort to uncover the assumptions and analytic approaches at work (e.g. [2, 3]). He has advocated an institutional approach that casts privacy as an issue not simply of individual needs and specific technologies, but one that arises from recurrent patterns of social roles and relationships [5].

Technology may be able to help as well as hinder. Writing in the context of the W3C's Platform for Privacy Preferences [26], Ackerman and Cranor [1] propose privacy critics – agents that will provide users with feedback on the potential privacy implications of their action. Similarly, Dourish and Redmiles [13] propose an architecture for enhancing user understandings of the security implications of their actions on networked computer systems.

Notwithstanding these investigations, the general state of understanding privacy concerns is limited in our fields of research and design. We feel that our goal—to better understand and describe the role of information technology in privacy management—is best met by returning to privacy theory that predates digital technologies.

Debates over privacy are not new, and did not arrive with information technology. The history of privacy is long and intricate, involving a wide range of concerns including social and legislative practice, cultural adaptation, and even urban and domestic architecture. A fuller treatment of privacy and technology merits a deeper examination of this background. However, we rely here on Altman's contemporary model of privacy and privacy regulation, which emerged from this long history.

## ALTMAN'S PRIVACY THEORY

Altman's fundamental observation is that privacy regulation is neither static nor rule-based. We know that setting explicit parameters and then requiring people to live by them simply does not work, and yet this is often what information technology requires, from filesystems to email filters to databases to cell phones. Instead, a fine and shifting line between privacy and publicity exists, and is dependent on social context, intention, and the fine-grained coordination between action and the disclosure of that action [6, 7].

### Privacy as Process

Altman conceptualizes privacy as the "selective control of access to the self" regulated as dialectic and dynamic processes that include multimechanistic optimizing behaviors [7: 67].

Altman describes privacy's dialectic and dynamic nature by departing from the traditional notion of privacy as a withdrawal process where people are to be avoided. Instead, Altman conceptualizes privacy as a *boundary regulation process* where people optimize their accessibility along a spectrum of "openness" and "closedness" depending on context. Privacy is not monotonic, that is, more privacy is

not necessarily better. Indeed, both "crowding" and "isolation" are the result of privacy regulation gone wrong. Privacy states are relative to what is desired and what is achieved; one can be in the presence of others but feel isolated or crowded depending on the degree of sociability sought. The goal of privacy regulation is to modify and optimize behaviors for the situation to achieve the desired state along the spectrum of openness and closedness. To that end, people employ "a network of behavioral mechanisms," which include

> *…verbal and paraverbal behaviors such as personal space and territoriality, and culturally defined styles of responding. Thus privacy regulation includes much more than just the physical environment in the management of social interaction. Furthermore, these behavioral mechanisms operate as a system. As such, they include properties of interdependence and of compensatory and substitutable action. That is, a person may use different mixes of behaviors to achieve a desired level of privacy, depending upon circumstances. Or different people and cultures may have unique blends of mechanisms to regulate privacy. [7: 67-68]*

### Caveats and Elaborations

Altman's theory is foundational, but has limitations for our purposes. He is concerned with the management of personal access in public spaces and other forms of interpersonal interaction; his attention is devoted primarily to situations where access is mediated by the everyday spatial environment. Information technology and the everyday environment mediate action in different ways [16].

Additionally, while Altman analyzes cultural differences, we attempt only to address conditions of *circumstance*, which we define to be a function of local physical environment, audience, social status, task or objective, motivation and intention, and finally, information technologies in use. Technologies and the forms of their use set conditions, constraints, and expectations for the information disclosures that they enable or limit. We view information technology not simply as an instrument by which privacy concerns are reflected, achieved, or disrupted; rather, it is part of the circumstance within which those concerns are formulated and interpreted.

## PRIVACY IN A NETWORKED WORLD

Privacy management is not about setting rules and enforcing them; rather, it is the continual management of boundaries between different spheres of action and degrees of disclosure within those spheres. Boundaries move dynamically as the context changes. These boundaries reflect tensions between conflicting goals; boundaries occur at points of balance and resolution.

The significance of information technology in this view lies in its ability to disrupt or destabilize the regulation of boundaries. Information technology plays multiple roles. It can form part of the context in which the process of boundary maintenance is conducted; transform the boundaries; be a means of managing boundaries; mediate representations of action across boundaries; and so forth. However, to better understand the role of technology, additional precision about these boundaries is needed.

We begin by describing three boundaries that we believe are central to the characterization of privacy management. One is the *Disclosure* boundary, where privacy and publicity are in tension. At this boundary, determinations are made about what information might be disclosed under what circumstances, albeit with varying degrees of direct control. The display and maintenance of *Identity* of parties on both sides of the information exchange occurs at another boundary. Features of identity, including institutional affiliation, are managed in tension with audience. *Temporality* describes the boundaries associated with time, that is, where past, present and future interpretations of and actions upon disclosed information are in tension.

Furthermore, the objectives that determine where each of these boundaries lies are in tension with each other. Actions around information disclosure are tempered by possibilities about what might happen in the future, or how the information might be judged as an artifact of the past. Actions around disclosure are balanced with how one wants to present oneself, or by knowledge of who might consume and re-use disclosed information, and so forth.

### The Disclosure Boundary: Privacy and Publicity

As Altman theorizes, privacy regulation in practice is not simply a matter of avoiding information disclosure. Participation in the social world also requires selective disclosure of personal information. Not only do we take pains to retain certain information as private, we also choose to explicitly disclose or publicize information about ourselves, our opinions and our activities, as means of declaring allegiance or even of differentiating ourselves from others (another kind of privacy regulation). Bumper stickers, designer clothing, and "letters to the editor" deliberately disclose information about who we are. We sit in sidewalk cafes to "see and be seen." We seek to maintain not just a personal life, but also a public face. Managing privacy means paying attention to both of these desires.

Furthermore, maintaining a degree of privacy, or "closedness" [6], will often *require* disclosure of personal information or whereabouts. The choice to walk down public streets rather than darkened back alleys is a means of protecting personal safety by living publicly, of finding safety in numbers. We all have thoughts or facts we would like to keep secret, but most of us also need to ensure that others know something about ourselves, for personal or professional reasons. For some, this management of personal and public realms is analogous to the job of a public relations agent who needs to make their client available and known in the world, while at the same time protecting them from the consequences of existing in this very public sphere. Celebrities operate in this space, but so do many lesser-known people: academics, for example, often feel compelled to maintain web pages, not only to advertise their expertise and experience, but also to keep requests for papers and other inquiries at bay. Therefore, one of the roles of disclosure can ironically be to *limit*, rather than *increase*, accessibility. Views of privacy that equate disclosure with accessibility fail to appreciate this necessary balance between privacy and publicity.

Active participation in the networked world requires disclosure of information simply to be a part of it. To purchase goods, we make ourselves visible in public space; in exchange for the convenience of shopping on-line, we choose to disclose personal identity information for transactional purposes. In so doing, we assume some risk of identity theft, although we might mitigate risk by shopping only at well-known web sites.

However, problems emerge when participation in the networked world is not deliberate, or when the bounds of identity definition are not within one's total control. A Google search, for example, can reveal a good deal of information about a person, including the artifacts and traces of past action, which may have been concordant with self-perception at a particular time—such as postings to Usenet—but not in later years. Information can also come from other, third-party sources, including public record data that was never as easily accessible as the web makes it today, such as the price paid for homes. Even friends might benevolently post photographs from a recent party that one would not post on one's own (for any number of reasons, including revealing behavioral as well as location-in-time information). When one's name is unique and therefore easily searchable, these concerns about public presentation of the self are magnified. One might even take deliberate action to formulate a public persona under these conditions by way of a personal web page, if only to mitigate or put in balance the perceptions one might gather from other sources. As these examples show, the tension around privacy and publicity is influenced by identity and temporal concerns, which we now address in turn.

### The Identity Boundary: Self and Other

The second tension central to privacy management is that which occurs around the Identity boundary; that is, the boundary between self and other. On first reflection, such a boundary might seem counter-intuitive or even nonsensical. Conventional formulations of privacy problems focus on the individual, and the boundaries of the individual would seem to be stably defined by the spatial extent of the body. However, when we look at privacy as a social phenomenon, this simple formulation becomes inadequate.

Affiliation and allegiance are complicating factors. The individualistic perspective assumes that people act primarily *as* individuals, which fails to take into account when people act as representatives or members of broader social groups, as they often do. Social or professional affiliations set expectations that must be incorporated into individual behavior, which is why disclaimers about corporate liability in email signatures exist, or even why employees are discouraged or barred from using corporate email address to post to public forums. Furthermore, adopting a particular set of attitudes towards appropriate information disclosure can even serve as a means of marking status or affiliation. ("Client confidentiality" is a marker of professional status for physicians and lawyers but not plumbers.) In other words, the inclusiveness or exclusiveness implied by self and other is continually enacted in and through one's actions in the world.

The tension between self and other is also problematized by the phenomenon of *recipient design*—the way that one's actions and utterances are designed with respect to specific others. That is, not only is "self" constructed with respect to a set of social arrangements, but "other" is not entirely undifferentiated—at different times, different "others" (professional colleagues, students, fellow bus riders, or whatever) can be distinguished from each other and will be treated differently. So, for example, when technology advocates argue that security cameras mounted in civic spaces offer no threat to individual privacy because one's actions are "already public," they fail to take into account that "public" is a broadly faceted concept, and that denying the ability to discern who might be able to see one's action can, in itself, constitute a violation of personal privacy.

Our *reflexive interpretability of action*—one's own ability to understand and anticipate how one's actions (and information, demeanor, etc.) appear to others—is sometimes compromised in information technology supported environment and has repercussions for privacy management. Assessing the efficacy of strategies for withholding or disclosing information is inescapably based on this reflexive interpretation. To withhold information, one needs to know from whom it is to be withheld and how that can be done, which requires an understanding of how our actions will be available or interpretable to others.

The fundamental problem of technology in interaction, then, is *mediation*. In the everyday world, we experience relatively unfettered access to each other, while in technological settings our mutual access is mediated by some technology that interposes itself, be that telephone, email or other computer system. Rather than interact directly with another person, we interact with and make assessments from a representation that acts in proxy. However, in technologically-mediated environments, these representations are often impoverished, and indictors of the boundary between privacy and publicity are unclear. We implicitly and constantly seek to understand how others want to be perceived along many dimensions, including their degree of availability and accessibility, but interactions can go awry when what is conveyed through the technological mediation is not what is intended. Privacy violations, then, can occur when regulatory forces are out of balance because intent is not adequately communicated nor understood.

Information persistence as a result of technological mediation further complicates regulation of the self/non-self boundary, and to what degree a person feels that information can act as proxy to the self. Explicit products of work or activity—such as manuscripts posted to a web page, or Usenet postings that we now know are archived—can be used to construct and control how we want to be perceived. In comparison, we have very little control over representations of ourselves that are artifacts of simply having been somewhere or done something at a particular time—such as visiting a cookie-enabled web page, or as being listed as a member of an email distribution list. How this kind information is interpreted is largely in the control of recipients. Those interpretations

subsequently vary in time, yielding even less direct control by the person for whom the information represents.

**Temporal Boundaries: Past, Present and Future**

Altman's view of the dialectic nature of privacy management is perhaps most obviously seen in the tension between past and future. The critical observation here is that specific instances of information disclosure are not isolated from each other; they occur as the outcome of a sequence of historical actions, and as the first of many expected actions stretching out into the future. The exercise of control, or the active management of privacy as the outcome of some decision-making process, needs to be seen in the context of this temporal sequence.

Past actions are a backdrop against which current actions are played. Our response to situations of potential information disclosure in the present are likely to draw upon or react to similar responses in the past, both those of our own and those of others. This is not to say, of course, that we blindly act in the same way every time a situation recurs; if this were true, than privacy regulation would not be the dynamic process we have attempted to illustrate here. Patterns, conventions, and genres of disclosure (see below) are made to be broken; conventions can be invoked by breaching as well as by following them. The relevance of future context is part of this same, continuous process; we orient not only to immediate circumstances but also to potential future situations. Current actions may be a means to affect future situations (as in our point above about academic web pages, where *current* disclosure is used to limit *future* accessibility.)

So, while past and future interpretations of information disclosure are out of our control, the way in which current privacy management is oriented towards events in the past or future is a matter of active control and management. Our response to situations of disclosure, or our interpretation of information we encounter, is framed and interpreted according to these other events and expectations. We negotiate boundary locations as we act in the world.

Technology's ability to easily distribute information and make ephemeral information persistent affects the temporal nature of disclosure. In our view, such ability is seen not as a fundamental blow to regulatory control, but rather as part of the ongoing management of tensions in which information permanence may play as much of a role as impermanence. The relevance of permanence and impermanence likes in the ways they constrain, undermine, or modify regulatory behavior. Because future uses of information cannot always be controlled, the nature or format of information might instead be governed. For example, distributing manuscripts as PDF rather than Microsoft Word format reduces the ease with which one's ideas can be modified; excerpts might be still extracted and changed, but the integrity of the whole remains.

The Disclosure, Identity and Temporality boundaries, and the tensions that occur with their negotiation, are the primary features of our framework. They demonstrate that privacy regulation is a dynamic, dialectic, negotiated affair. Technology itself does not directly support or interfere with personal privacy; rather it destabilizes the delicate and complex web of regulatory practices.

## GENRES OF DISCLOSURE

When considering these three tensions, it is important to bear in mind that they are not resolved independently; all three are part of the same, ongoing process of privacy management and regulation. At any given moment, the balance between self and other, privacy and publicity, and past and future must have a single coherent and coordinated resolution. As a unifying principle, we use the term *genres of disclosure* to highlight these socially-constructed patterns of privacy management. Using this term draws attention to the ways in which privacy management in everyday life involves combinations of social and technical arrangements that reflect, reproduce and engender social expectations, guide the interpretability of action, and evolve as both technologies and social practices change. Evolution occurs as the possibilities and consequences of particular technological arrangements should gradually be incorporated into practice. Integrating and resolving the various tensions, these regularly reproduced arrangements of people, technology and practice that yield identifiable and socially meaningful styles of interaction, information, etc., are what we refer to as genres of disclosure.

Our use of the term "genre" is deliberately suggestive of the work of researchers such as Yates and Orlikowski [27] and Erickson [15]. The important feature of their work is that they adopt a socially-constructed notion of genre, defined not simply by the structural properties of particular forms of communication, but by the social patterns of expectation and response that genres embody. Genres are encounters between representational forms and social practice. Privacy and technology is not least a matter of representation (of people and their actions), and the relevance of genre is precisely in how it sets expectations around these representations, integrating them into recurrent social practices. For example, Erickson [14] cites the example of a graduate student's reflections on a personal web page as a tool of self-expression and as a professional badge of entry into the job market; the issues of interpretation and the identification of the information as fitting into a commonly-understood pattern of communication was a central issue. Similarly, investigations of personal web pages have pointed to the importance of particular styles and interpretations of the information [10]. Personal web pages are clearly bound up with issues of disclosure; the idea of genres of disclosure extends this to other forms of interaction and potential information disclosure, from our expectations over the monitoring of our movements in public spaces to concerns over personal information requested on web-based forms.

An important feature of the notion of genre of disclosure is that, since genres are loosely defined, it can account for violations, or situations in which one feels that the "promise" of a genre was broken; that, for instance, personal information was misappropriated and used in ways that had not been anticipated. The very idea of misappropriation implies that information is disclosed with an expectation of appropriate use; the relationship between

forms of disclosure and expectations of use is precisely what the idea of genre is intended to capture.

## CASE STUDIES

The central motivation for this paper is to understand the complexity and multi-faceted nature of privacy in settings where information technology is present, which has necessitated the extensive conceptual exploration above. However, in an effort to make this perspective more tangible, we explore various cases drawn from our own observations as well as those of our colleagues; some of these cases have been discussed in research publications, while others are informal experiences reflected upon here for the first time. Analyses are necessarily abbreviated; our goal is only to illustrate how the tensions and considerations we have presented here might be used to express more nuanced understanding of privacy concerns.

### The Family Intercom

In an influential project, Georgia Tech researchers are building a residential research laboratory—the "Aware Home," a three-story house, intended for occupation—which is a testbed for the use embedded computing technologies in domestic settings. One of the program's projects is the Family Intercom [22], which allows family members to communicate seamlessly with one another when distributed throughout their home using a built-in sensing and tracking infrastructure.

This project can be interpreted as a reproduction, for domestic environments, of ubiquitous communication technologies that have been explored in research workplace settings (e.g. Active Badges, below, and media spaces [9]). This comparison illustrates a mismatch between the institutional arrangements of home and work life. Such concepts as "accessibility," "awareness," and "availability for interaction," which have been goals supported by research workplace studies, do not map conveniently onto equivalents in the home environment. A six-year old is not in a position to control her "availability for interaction" to her parents; she may display attentiveness or disinterest, but the concept of voluntary availability does not apply in this different institutional setting. Similarly, a sixteen-year-old may not appreciate his sibling's "passive awareness" of his actions. By the same token, the notion of the substitutability of media—that a conversation over an intercom is a replacement for a face-to-face conversation—does not apply in settings where what is being conducted is not simple communication but an exercise in power relations; when a parent calls to a child, what is often demanded is presence, not communication.

In other words, the genres of disclosure and accessibility in these two settings are quite different. The interesting questions, then, are which institutional arrangements are implicit in the design of the technology, and to what extent it is possible to blur these genres in practice.

### Shared Calendars

Enterprise calendar applications that allow users to share calendar data in the hopes of improving coordination, most typically in corporate settings, illuminate a range of privacy issues, from maintaining company security to protecting one's time. The benefits to temporal coordination have been demonstrated time and time again, and so we find that people are more willing than in the past to share information that was once considered private [24].

However, publicly available calendar data makes explicit the patterning and sequencing of information, and the interpretations that can be made from such patterns might inadvertently compromise privacy. The implications of this are particularly apparent in this case of conference room bookings: A technology company was rife with rumors about an impending layoff, but details were confidential. An employee using an online calendar system to book the local conference room for a meeting the following week found that it was already scheduled, and, in searching for alternatives, gradually discovered that every room had been booked, all day, by Human Resources. Although room bookings are not expected to be a channel by which large-scale corporate plans might be discovered, the employee was able to easily infer that layoffs were imminent.

This illustrates an interesting tension between publicity and privacy. One cannot book a room without disclosing that it has been booked (although other systems might disguise identities); publicity is relevant here, since advertising the rooms' unavailability is the point of the exercise. However, the ability to see aggregate temporal patterns of information ("all rooms booked by HR"), rather than individual data points, constitutes the disclosure problem. It is not simply that HR was forced into disclosing information they would have preferred to keep secret. Rather, it was that they desired *publicity* at the level of individual data points, but *privacy* at the level of the whole, with no system-supported means for making this distinction.

### Active Badges

In a classic case, Harper [19, 20] discusses experiences in the deployment of personal tracking systems based on Active Badges in two research laboratories. Of note is the variety of responses, differing both between the laboratories, and between different groups in each.

For example, the scientific staff in each lab placed different values on the technology, in part due to different styles of work. In one lab, staff often worked in a central lab, and so the ability to route phone calls to them was highly valued; in the lab where people worked at their desks, the active badge technology seemed less useful, and even intrusive. This speaks in part to the tension between publicity and privacy, and people's ability to control the balance.

On the other hand, Harper suggests a different orientation between different staff groups. The administrative staff (who often need to track down research staff's whereabouts) were more positively disposed to the technology than were the scientific staff. Harper proposes that this may be in part due to the professional self-image of the scientific staff, who were more likely to adopt a position of individual responsibility for their own actions, and perhaps then resent the introduction of a technology that potentially limits individual freedom and imposes greater organizational accountability. The administrative staff, perhaps, were less likely to feel this organizational accountability as a threat;

it was already a feature of their working lives. In our terms, this reflects different ways to resolve the tension between self and other; the scientific staff, understanding "self" in personal terms, saw the system as providing a new form of information disclosure, while the administrative staff saw less of an impact since their notion of "professional self" was already more strongly associated with the organization.

### Mobile Telephones

The public use of mobile phones has been the topic of much discussion in the popular and research literatures. Its rapid, widespread deployment has called attention to the norms and conventions that guide behavior in public places, as many people report that what constitutes appropriate use of mobile phones is violated there [25]. When such violations occur, it is not upon the user of the technology—as we usually see in other technology use scenarios—but rather upon the unassuming person who feels that a conversation has been thrust upon them, with violations made to their acoustical privacy. The boundary between privacy and publicity is challenged here: the telephone user may feel comfortable about the degree of openness they display, whereas the recipient who occupies the same physical space has little control of the degree of closedness they desire (short of asking the phone user to move, or to move him- or herself). The boundary between self and other is destabilized when phone users assume that they are without an audience, or that they somehow do not affect anyone else, or, worse yet, that their behaviors are of interest to those who surround them. The furor and upset that surrounds mobile telephone use emerges from the overlap between two domains of activity—personal conversation and public presence. Phone conversations are not subject to the same self-monitoring and responsiveness to setting that characterizes face-to-face interactions. What we are witnessing, then, is the gradual emergence of new communication forms that represent alternative resolutions of the tensions, organized around the new technological forms and their consequence for social opportunities [21].

### Instant Messaging

Instant messaging (IM) raises a range of privacy concerns, including tensions at the temporal boundary, created by the possibility of recording what is assumed to be ephemeral information for future use. Among teenagers, IM conversations tend to be informal and in-the-moment, much like face-to-face interaction [17]. Their IM communications are crafted for the present with the foreground assumption that friends can be trusted; however, the potential that their statements might instead be recycled for unauthorized purposes, by copying to a permanent record, keeps in check what is revealed to whom.

Teenagers' use of IM in their homes illustrates other kinds of privacy regulation behaviors [17]. Teens report preferring IM to the family phone because IM does not advertise to parents that they are engaged in conversation with others, perhaps at times when such communications would be discouraged or even prohibited. In the virtual meeting space that IM creates, teens want to advertise their publicity and availability to their friends; in contrast, in the physical space of the home, they want to minimize attention placed on their IM participation. In our privacy regulation terms, this tension occurs at the disclosure boundary, but also at the identity boundary, where teens pay attention to who they are expected (and want) to be in each of the spaces. Finally, we can say that the genres of disclosure for the two spaces are distinctly constructed and maintained.

### Summary

The main point we have tried to emphasize in our conceptual development is the dynamic and multidimensional nature of privacy. Privacy management involves satisfying a number of needs, and balancing a number of tensions. Taken individually, the examples presented here illustrate how our approach can illuminate specific issues in the interaction of privacy and information technology. Taken together, they demonstrate the diversity of privacy issues at work in everyday settings. In contrast to the traditional model of privacy as social withdrawal, we can see many different tensions at work. This, again, points to the need for interpretive frameworks, to help unpack and elucidate these different questions. Any encounter between privacy and technology will involve many or all of these different tensions. To understand the impacts of technology, we need to be able to see how these different tensions operate, separately and together.

## CONCLUSIONS

Our initial goal was to "unpack" the idea of privacy and propose a conceptual framework that would allow more specific and detailed statements about privacy and technology to be made in HCI analyses. Our central arguments have been that privacy management is a dynamic response to circumstance rather than a static enforcement of rules; that it is defined by a set of tensions between competing needs; and that technology can have many impacts, by way of disrupting boundaries, spanning them, establishing new ones, etc. Using case studies, we have attempted to show how such a perspective might illuminate our understanding of privacy regulation. What are the consequences of this perspective for technologists and designers? We submit four possibilities here.

First, our view emphasizes that, when considering privacy concerns raised by the development of new technologies, the whole of the social and institutional setting in which technologies are deployed should be considered. What is important is not what the technology does, but rather how it fits into cultural practice. As we suggested earlier, in Orwell's *1984*, it is the culture of pervasive mutual monitoring that constitutes the threat to individuals.

Second, our perspective on privacy requires attention to the historical continuity of practice. Privacy regulation is oriented both to the past and the future. Adequate analyses of the affects of technology on privacy practices would interpret those practices not as arbitrary decontextualized provisions, but as part of a trajectory of action.

Third, this perspective shows that privacy management is something of a balancing act, a resolution of tensions not just between people but also between their internal conflicting requirements. The significance here is that small changes may have disproportionately large effects.

Finally, our perspective demonstrates that the active process of privacy management takes place in the context of the possibilities that are offered by one or another technology. So, what technology enables is as important as how it is actually used; it is the possibilities, rather than the actual practice, around which privacy regulation is performed. Therefore, we need to be as responsible for what we make possible as for what we make real.

In offering both a framework and a vocabulary for talking about privacy and technology, our goal is to foster discussion between technology users, designers and analysts, and to encourage a more nuanced understanding of the impacts of technology on practice. Privacy will continue to be a significant factor in the design of information technology; our understanding of what privacy is and how it operates will need be as sophisticated as the technologies involved. We hope this is a first step.

## ACKNOWLEDGMENTS

This work was supported in part by a grant from the National Science Foundation (#IIS-9977952). We are grateful to Mark Ackerman, Victoria Bellotti, Tom Erickson, Jonathan Grudin, John King and anonymous reviewers for valuable comments.

## REFERENCES

1. Ackerman, M. & Cranor, L. 1999. Privacy Critics: UI Components to Safeguard Users' Privacy. *Proc. ACM Conf. Human Factors in Computing Systems CHI'99*, 2, 258-259.
2. Agre, P. 1995. Conceptions of the User in Computer Systems Design. In Thomas (ed), *The Social and Interactional Dimensions of Human-Computer Interfaces*, 67- 106. Cambridge: CUP.
3. Agre, P. 1997. Beyond the Mirror Worlds: Privacy and the Representational Practices of Computer Science. In Agre & Rotenberg (eds), *Technology and Privacy: The New Landscape*. Cambridge, MA: MIT Press.
4. Agre, P. 1999. The Architecture of Identity: Embedding Privacy in Market Institutions. *Information, Communication, and Society*, 2(1), 1-25.
5. Agre, P. 2001. Changing Places: Contexts of Awareness in Computing. *Human-Computer Interaction*, 16(2-4), 177-192.
6. Altman, I. 1975. *The Environment and Social Behavior: Privacy, Personal Space, Territory and Crowding*. Monterey, CA: Brooks/Cole Pub. Co., Inc.
7. Altman, I. 1977. Privacy Regulation: Culturally Universal or Culturally Specific? *Journal of Social Issues,* 33 (3), 66-84.
8. Bellotti, V. & Sellen, A. 1993. Design for Privacy in Ubiquitous Computing Environments. *Proc. Third European Conf. Computer-Supported Cooperative Work ECSCW'93* (Milano, Italy), 77-92. Dordrecht: Kluwer.
9. Bly, S., Harrison, S., & Irwin, S. 1993. Media Spaces: Bring People Together in Video, Audio and Computing Environments. *Communications of the ACM*, 36(1), 28-47.
10. Bly, S., Cook, L., Bickmore, T., Churchill, E., & Sullivan, J. 1998. The Rise of Personal Web Pages at Work. *Proc. ACM Conf. Human Factors in Computing Systems CHI'98,* 2, 313-314.
11. Clement, A. 1994. Considering Privacy in the Development of Multimedia Communications. *Computer-Supported Cooperative Work*, 2, 67-88.
12. Dourish, P. 1993. "Culture and Control in a Media Space," *Proc. Third European Conf. Computer-Supported Cooperative Work ECSCW'93* (Milano, Italy), 125-138. Dordrecht: Kluwer.
13. Dourish, P. & Redmiles, D. 2002. An Architecture for Usable Security based on Event Monitoring and Information Visualization. *Proc. New Security Paradigms Workshop NSPW'02* (Virginia Beach, VA).
14. Erickson, T. 1996. The World-Wide Web as Social Hypertext. *Comm. of the ACM*, 37(1), 15-17.
15. Erickson, T. 1997. Social Interaction on the Net: Virtual Community as Participatory Genre. *Proc. Hawaii Int'l Conference on Systems Science HICSS.* Los Alamitos, CA: IEEE Computer Society Press.
16. Gaver, W. 1992. The Affordances of Media Spaces for Collaboration. *Proc. ACM Conf. Computer-Supported Cooperative Work CSCW'92* (Toronto, Ontario), 17-24. New York: ACM.
17. Grinter, R. & Palen, L. 2002. Instant Messaging in Teen Life. *Proc. ACM Conf. Computer-Supported Cooperative Work CSCW'02* (New Orleans, LA), 21-30. New York: ACM.
18. Grudin, J. 2001. Desituating Action: Digital Representation of Context. *Human-Computer Interaction, 16,* 269-286.
19. Harper, R. 1992. Looking at Ourselves: An Examination of the Social Organization of Two Research Laboratories. *Proc. ACM Conf. Computer-Supported Cooperative Work CSCW'92* (Toronto, Ontario), 330-337. New York: ACM.
20. Harper, R., Lamming, M. & Newman, W. 1992. Locating Systems at Work: Implications for the Development of Active Badge Applications. *Interacting with Computers*, 4 (3), 343-363.
21. Laurier, E. 2001. Why People Say Where They Are During Mobile Phone Calls. *Environment and Planning D: Society & Space*, 19(4), 485-504.
22. Nagel, K., Kidd, C., O'Connell, T., Dey, A., & Abowd, G. 2001. The Family Intercom: Developing a Context-Aware Audio Communication System. In Abowd, Brumitt and Shafer (eds), *Ubicomp 2001, LNCS 2201*, 176-183. Berlin: Springer-Verlag.
23. Orwell, G. 1949. *Nineteen Eighty-Four*. London: Martin Secker & Warburg.
24. Palen, L. 1999. Social, Individual and Technological Issues for Groupware Calendar Systems. *Proc. ACM Conf. Human Factors in Computing Systems CHI'99* (Pittsburgh, PA), 17-24. New York: ACM Press.
25. Palen, L., Salzman, M. & Youngs, E. 2000. Going Wireless: Behavior and Practice of New Mobile Phone Users. *Proc. ACM Conf Computer Supported Cooperative Work CSCW 2000* (Philadelphia, PA), 201-210. New York: ACM Press.
26. Reagle, J. and Cranor, L. 1999. The Platform for Privacy Preferences. *Comm. ACM*, 42(2), 48-55.
27. Yates, J. & Orlikowski, W. 1992. Genres of Organizational Communication: A Structurational Approach to Studying Communication & Media. *Academy of Management Science Review,* 17(2), 299-326.