Draft only. To appear in P. Markopoulos, B. de Ruyter and W. Mackay, eds.: Awareness Systems: Advances in Theory, Methodology and Design. Berlin, Heidelberg, New York: Springer Verlag.

SAMEER PATIL, ALFRED KOBSA

PRIVACY CONSIDERATIONS IN AWARENESS SYSTEMS

Designing with Privacy in Mind

1. INTRODUCTION

The earlier chapters of this book present a conceptual understanding of awareness [cite Section 2]. A historical account [cite Section 1] as well as descriptions of various implementations [cite Section 3] illustrate how various systems have attempted to foster greater awareness. A common challenge that all awareness systems face is the tension with an individual's desire for privacy [Hudson and Smith 1996].

Interaction between awareness and privacy is not limited to awareness systems but is a characteristic of everyday life. As Schwartz [1968] notes, "*We are led to relinquish our private information and activities by the expediencies and reciprocities routinely called for in daily life. We all know, for example, that in order to employ others as resources it is necessary to reveal to them something of ourselves*".

In the case of awareness systems, the risks of reduced privacy for every individual whose information is disseminated, and the benefits of such information for all recipients must be carefully evaluated. Moreover, these systems require users to extend their existing practices regarding awareness and privacy, from the familiar physical domain to the newer digital domain. Situations that lack of familiarity are known to be problematic for privacy management though, and may lead to privacy violations [Romero and Markopoulos 2008].

Privacy management in the digital domain poses precisely such difficulties. Certain characteristics of the digital domain differ substantially from the physical world, namely, high-speed transmission, potential persistence, and enhanced computation. The digital domain may also result in disembodiment [Heath and Luff 1991] (e.g., one may be represented only by a screen name). Disembodiment thwarts the ability to exploit the rich cues that are readily used in face-to-face interactions (e.g., posture, expressions, intonation). In addition, dissociation of interaction [Bellotti 1998] could occur when only the results of people's actions are shared while the actions themselves are not visible (e.g., a Wiki page with no version history available).

[Editor(s) here] (ed.), [Book Titlen here], 1—18.

Owing to these differences, the transformation of expectations and behaviors from the physical to the digital world is not always effective, or even possible.

As a result, simultaneously satisfying users' awareness as well as privacy needs poses a significant challenge for designers of awareness systems. Insufficient attention to either of these needs could potentially undermine the usage of a system. When users are unable to achieve appropriate levels of awareness and privacy effortlessly, they may not exploit a system's potential fully. For instance, Lee et al. [Lee, Girgensohn et al. 1997] found that when privacy was desired, users of their Portholes video system preferred to simply turn their cameras off since fiddling with other privacy options, such as blurring the video, was too cumbersome. Herbsleb et al. [Herbsleb, Atkins et al. 2002] found it difficult to attain substantial usage of their chat system due to the fact that its default settings were too private. The system imposed significant initial setup efforts on its users to be able to provide awareness benefits. Likewise, we found in our own research that users who were forced to use instant messaging due to organizational requirements often resorted to circumvention tactics. They set their status to "away" or "busy" even when they were not, or conversely, changed their preferences so as to always appear online even when they were away from their desks [Patil and Kobsa 2004; Patil and Kobsa 2005]. Such circumvention results in suboptimal use of awareness systems.

Focusing on awareness without paying sufficient attention to its privacy aspects may evoke strong user backlash. A recent example that involves the popular social networking site Facebook (http://www.facebook.com) is an excellent illustration. Facebook introduced a new awareness feature that automatically presented to each user an aggregation of every single activity of their friends. Tens of thousands of users were outraged. The revolt ranged from online petitions and protest groups to threats of a boycott [Calore 2006]. This episode underscores that user opposition due to privacy concerns can translate into minimal use or even abandonment of a system. In such cases, organizations stand to lose their investment in deployed awareness applications. Moreover, organizations that design and build these systems, as well as their customers, face the prospect of longer-term damage to their trust and credibility in the users' eyes [Adams 1999; Adams and Sasse 1999].

Thus, it is important for awareness systems to respect users' privacy concerns. This chapter analyzes theoretical and empirical work to aid designers in building privacy-sensitive awareness systems.

2. PRIVACY

Before we can discuss how awareness systems could be made more privacy-sensitive, we must first take a detailed look at privacy itself. The concept of privacy is so intricate that there is no universal definition of it. The difficulty of defining privacy stems from its highly *situated* [Suchman 1987] and context-dependent nature. Even in the *same* situation, different individuals may have differing opinions and expectations regarding what privacy means to them (for example, Westin [1991] classified individuals as privacy fundamentalists, pragmatists, or unconcerned, based upon their preferences). This context dependency and variability between

individuals make dealing with privacy a difficult task. To quote Lederer et al. [Lederer, Hong et al. 2004] "*One possible reason why designing privacy-sensitive systems is so difficult is that, by refusing to render its meaning plain and knowable, privacy simply lives up to its name. Rather than exposing an unambiguous public representation for all to see and comprehend, it cloaks itself behind an assortment of meanings, presenting different interpretations to different people*".

There are three main perspectives from which the notions of privacy are commonly described and analyzed (see *Table 1*):

Perspective	**Concept of Privacy**	**Enacted by**	**Consequences of privacy violation**
Normative	Right or freedom	Laws, Contracts, Policies	Civil and/or criminal penalties
Social	Socially constructed	Individual and collective everyday social action	Potential embarrassment or breakdown in relationship(s) etc.
Technical	Control over data and information	Automated and/or manual access control	Identity theft, unauthorized access, illegal use of information

Table 1. *Three perspectives regarding the concept of privacy*

Normative: Analyzed philosophically, privacy is an ethical concept [Negley 1966; Johnson 1985; Mason 1986]. Privacy is viewed as a "right" of individuals, and, thus, as a matter of "freedom". For example, Warren and Brandeis [1890] characterized privacy as "*the freedom to be left alone*". With this perspective, privacy is a civil liberty that needs to be protected through legal and political means. Traditionally, the focus of privacy protection has been on laws, contracts and policies aimed at protecting an individual from large entities such as corporations and the government [Lessig 1999]. Increasingly, however, legislation is being extended to privacy protecting one's privacy from other individuals (for instance, laws against hacking, stalking or voyeurism).

Social: From a social perspective, privacy has psychological roots [Westin 1967; Schwartz 1968]. Privacy is "socially constructed" based on the behavior and interaction of individuals as they conduct their day-to-day affairs. For instance, in Goffman's [1959] analysis "*the expressive component of social life has been treated as a source of impressions given to or taken by others*", where expression "*has been treated in terms of the communicative role it plays during social interaction*". This manifests itself in Rachels' [1975] claim that "*privacy is necessary if we are to maintain the variety of social relationships with other people that we want to have*". Thus, managing privacy allows us to manage social relationships. Altman [1975] has described the process of privacy management as a "*dialectic and dynamic*

boundary regulation process" – conditioned by the expectations and experiences of the parties involved and under continuous negotiation and refinement. Given differences in norms, expectations, experiences, behaviors, and laws across cultures, it is no surprise that privacy manifests itself differently in different cultures [Westin 1967; Milberg, Burke et al. 1995]. Viewed socially, the notion of privacy evolves as external changes bring about changes in expectations and behavior, or as technology introduces new forms or means of interaction.

Technical: The technical perspective views privacy from the functional characteristics of digital systems. Discussions from this perspective tend to investigate *how* ethical and social considerations could be operationalized. Privacy is thus treated as the desire for selective and adequate control over data and information – both incoming and outgoing. For example, Stone et al. [Stone, Gueutal et al. 1983] describe privacy as "*ability of the individual to personally control information about oneself*" whereas Samarajiva [1998] extends it to "*control of outflow of information that may be of strategic or aesthetic value to the person and control of inflow of information including initiation of contact*". The issues under consideration include the capture, storage, ownership, usage, and access of personal data. For example, the code of Fair Information Practices was developed from this perspective [U.S. Department of Health 1973].

To summarize, the social perspective focuses on what practices relate to privacy, while the normative discussions look at whether a particular behavior is ethically (or legally) justified. The technical discourse is concerned with how the ethical and social understandings can be represented formally, and implemented practically in an operational system. The three perspectives are not mutually exclusive but interdependent. Privacy laws may be enacted based on normative, technical or social considerations, while social interactions may be altered due to changing laws and technology.

Having laid out the foundational understandings of privacy, we now proceed to discussing how awareness and privacy interact with each other.

3. RELATIONSHIP BETWEEN AWARENESS AND PRIVACY

Given that the concepts of awareness and privacy are both related to disclosure, it should not be surprising that they interact with each other. This interaction between awareness and privacy is not new. Westin [1967] describes it as a balancing act:

> *"Privacy is neither a self-sufficient state nor an end in itself, even for the hermit and the recluse. It is basically an instrument for achieving individual goals of self-realization. As such, it is only part of the individual's complex and shifting system of social needs, part of the way he adjusts his emotional mechanisms to the barrage of personal and social stimuli that he encounters in daily life. Individuals have needs for disclosure and companionship every bit as important as their needs for privacy. As ancient and modern philosophers agree, man is a social animal, a gregarious being whose need for affiliation marks his*

> *conduct in every society. Thus, at one hour a person may want lively companionship and group affiliation; at another moment, the intimacy of family or close friends; at another the anonymity of the city street or the movie; at still other times, to be totally alone and unobserved. To be left in privacy when one wants companionship is as uncomfortable as the inability to have privacy when one craves it.*
>
> *[…] All individuals are constantly engaged in an attempt to find sufficient privacy to serve their general social roles as well as their individual needs of the moment. Either too much or too little privacy can create imbalances which seriously jeopardize the individual's well-being."*

In the context of awareness systems, this equilibrium corresponds to a reconciliation of the benefits of awareness for improving effectiveness and efficiency, and the potential risks of reduced privacy.

In the physical setting of everyday life, individuals utilize the spatial and architectural features of the environment [Schwartz 1968] (e.g., a door), biological and cognitive features of humans [Westin 1967] (e.g., limitations of human memory), and shared understanding of norms [Westin 1967] to meet their awareness and privacy needs. Thus, situations in which one's familiarity with aspects of day-to-day affairs breaks down (e.g., moving to a foreign country) have been observed to be problematic for privacy management.

Privacy is managed based on one's familiarity with these features and understanding of norms, acquired through daily life experiences. This, of course, does not imply that privacy violations could never occur in familiar everyday settings. In fact, privacy violations due to accidental disclosure are not uncommon [Schwartz 1968]. When a violation of privacy does occur, and is detected, individuals typically engage in *social negotiation* until a commonly agreed upon (or comfortable) state of privacy is reached for everyone involved. Westin [1967] describes practices such as covering one's face, averting others' eyes, or facing the wall. As Palen and Dourish [2003] point out, "*Privacy is understood to be under continuous negotiation and management, with the boundary that distinguishes privacy and publicity refined according to circumstance*".

Recent technological developments, e.g., in the fields of Computer Supported Collaborative Work (CSCW), have introduced the digital domain as an additional arena in which awareness and privacy need to be reconciled [Agre and Rotenberg 1998]. The next subsection describes how the digital domain, due to its relative novelty and its unique characteristics, poses new challenges in this regard.

3.1 Digital Domain

We noted earlier that situations in which familiarity breaks down are problematic for privacy management, and may lead to privacy violations. Awareness systems create exactly such problems since they require users to extend their privacy management practices from the familiar physical domain to the relatively new digital domain.

Additionally, certain characteristics that distinguish the digital domain from the physical world are important from a privacy standpoint. Salient among these are:

Transmission: The ease, speed, and low cost with which data is transmitted in the digital domain are major reasons why it is attractive for awareness systems. However, these advantages come at the expense of increased risk for unauthorized access through technical means such as hacking and network sniffing, and higher potential damages that may result from such attacks.

Persistence: With the advent of practically infinite storage capacity, the digital domain may increase the temporal dimension of data indefinitely. In contrast, information about the vast majority of routine activities in the non-digital world generally can be trusted to be merely ephemeral. The digital "trails" of one's activities undermine the "plausible deniability" [Nardi, Whittaker et al. 2000] of facts and actions that one does not want to admit to. It also separates information from the context in which it was generated [Dix 1990]. Moreover, the storage of personally identifiable information introduces legal issues of accountability, liability etc. For example, a Chinese journalist was convicted of leaking state secrets based on records of his Internet activities provided by Yahoo! Inc. [Kahn 2005]

Computing power: Data in digital form is amenable to kinds of analysis that are almost impossible in a non-digital world. Computing power also makes it possible to automate such analyses. For example, techniques like data mining, pattern detection, social network analysis, event notification, visualization etc. can be used for inference, prediction, profiling, surveillance, and much more.

Disembodiment and dissociation: As mentioned earlier, by being situated in between interacting individuals, the digital domain can cause disembodiment [Heath and Luff 1991] and dissociation [Bellotti 1998]. Disembodiment and dissociation hinder one's ability to present oneself as effectively to others as in a face-to-face setting, and result in a breakdown of social and behavioral norms and practices [Bellotti 1998]. For example, while Goffman [1959] describes how people present different appropriate "faces" in real life quite seamlessly, a direct operationalization of this metaphor in a digital system turned out to be unsuccessful [Lederer, Beckmann et al. 2003]. Moreover, disembodiment could result in individuals being forced to be explicit about certain information that is otherwise intuitive or implicit [Bellotti 1998].

As a result of these distinctions, the digital domain can inhibit behaviors that may be fluid and seamless in the social realm. Ackerman [2000] pointed out that in the digital domain privacy encounters the *social-technical gap*, i.e., "*the divide between what we know we must support socially and what we can support technically*". On the other hand, characteristics of the digital domain enable actions that may otherwise be impossible or prohibitively difficult to achieve socially. Lessig [1999] sums this up rather nicely: "*In the 1790s the technology was humans; now it is machines. Then the technology noticed only what was different; now it notices any transaction. Then the default was that searchable records were not collected; now the default is that all monitoring produces searchable records. These differences add up.*"

4. RELEVANT RESEARCH

Over the past few years, the importance of taking action on privacy issues engendered by awareness systems has gained increased attention. Research that tackles privacy in awareness systems falls along three major themes: users studies of specific awareness systems, design principles and guidelines derived from theoretical considerations, and privacy-enhancing technical solutions. We discuss each of these below.

4.1 User Studies

Initial findings related to privacy were primarily noted as side effects in studies aimed at evaluating experiences with the awareness aspects of systems. Dourish [1993] characterizes privacy controls along a "social-technical continuum". On the social side, social pressures and norms are relied upon to prevent misuse of the system. On the technical side, technology prevents attempted misuse. Social controls are likely to work well only within a small, relatively tight-knit community [Dourish 1993; Ackerman, Starr et al. 1997]. Even then, they may result in very strong protection behavior such as turning the system off, or altering one's work habits [Mantei, Baecker et al. 1991]. In contrast, technical privacy protections cause increased acceptance and adoption of a system due to increased user trust that the system will protect privacy [Dourish 1993]. Later studies confirmed that trust in the system is an important factor implicit in privacy assessments [Adams 1999; Adams and Sasse 1999; Patil and Lai 2005].

Palen [1999] found that socio-technical mechanisms controlled privacy even in highly open network calendaring environments. Users managed privacy partly via technical access control, partly via the norm of reciprocity[1], partly via practices such as cryptic entries, omissions, defensive scheduling, and partly via the social anonymity within the organizational context. Lee et al. [Lee, Girgensohn et al. 1997] suggest that in addition to the need for privacy control users also desire a lightweight mechanism to address it. As Herbsleb et al. [Herbsleb, Atkins et al. 2002] discovered, the lack of such mechanisms increases setup time. Grinter and Palen [2002] illustrate (albeit with teenagers) that users adapt system capabilities to their own ends. Teens in their study made enterprising use of access permissions, profiles, status messages and screen names to manage privacy. Nardi et al. [Nardi, Whittaker et al. 2000] show that plausible deniability of presence is used for managing privacy in instant messaging.

Recently, studies of awareness systems have started targeting privacy as the primary object of investigation. The user's relationship with the information recipient, the purpose or usage of the information, the context, and the sensitivity of the content have all been found to be important criteria used when users make privacy judgments [Adams 1999; Adams and Sasse 1999; Lederer, Mankoff et al. 2003; Patil and Kobsa 2004; Consolvo, Smith et al. 2005; Olson, Grudin et al. 2005]. Lederer et al. [Lederer, Mankoff et al. 2003] also showed that a-priori manual

[1] Palen [1999] found that individuals with unusually restrictive, or liberal, calendar access settings often had immediate colleagues with similar access configurations.

configuration of privacy preferences is better than automatic strategies – especially for information that users deem important.

Generic privacy attitudes and behaviors could also come into consideration in awareness systems. Therefore, it is instructive to look at a few privacy studies conducted in other contexts. For instance, as mentioned above, Westin [1991] classified individuals into three main clusters – privacy fundamentalists, pragmatists, and unconcerned. This distinction may also apply to privacy concerns in the context of awareness systems. Milberg et al. and Bellman et al. [Milberg, Burke et al. 1995; Bellman, Johnson et al. 2004] reported that privacy concern varies by country. At the same time, they mentioned that "secondary use" and "improper access" rank as the top two concerns across most nationalities. Cranor et al. [Cranor, Ackerman et al. 1999] listed anonymity and information sensitivity as important privacy-related factors for Internet users. Finally, Fox [2002] showed that users are often ignorant of the basic concepts underlying their digital domain activities, and typically do not utilize available tools for privacy protection.

4.2 Theories, Principles and Guidelines

Privacy is recognized to be a nuanced and situated concept that escapes universal definition. The rich body of literature on privacy in the social sciences is testimony to its intricate connections with the broader social context [Dourish and Anderson 2005]. Due to this complexity, technology designers have found it difficult to analyze and frame the privacy issues unveiled by user studies. Researchers have tried to address this problem by attempting to articulate theoretical insights into privacy in forms that are more accessible to system designers. For instance, Boyle and Greenberg [2005] describe a vocabulary of privacy that designers can employ for an unambiguous discussion of privacy issues. To suggest ways of thinking about privacy in socio-technical environments, Palen and Dourish [2003] outline a model of privacy that is based on the theory of social psychologist Irwin Altman. It views privacy as a process that regulates the boundaries of disclosure, identity and temporality. This process is both dynamic (i.e., shaped by personal and collective experiences and expectations) and dialectic (i.e., under continuous boundary negotiation).

Researchers also compiled various privacy-related findings from user studies into design principles and guidelines in order to enable better privacy management. Bellotti and Sellen [1993] propose a design framework based on feedback and control regarding information capture, construction, accessibility, and purpose. The purpose of feedback mechanisms is to provide users with information that helps them make judgments regarding privacy, while the purpose of control is to empower them to take appropriate actions to manage privacy. In addition, Bellotti and Sellen provide eleven criteria for evaluating design solutions – trustworthiness, appropriate timing, perceptibility, unobtrusiveness, minimal intrusiveness, fail-safety, flexibility, low effort, meaningfulness, learnability, and low cost. Langheinrich [2001] draws upon Fair Information Practices [U.S. Department of Health 1973] in proposing that privacy-sensitive systems ought to notify the user appropriately, seek consent, provide choice, allow for user anonymity or pseudonomity, limit scope with proximity as well as locality, ensure adequate security, and implement appropriate

information access. Iachello and Abowd [2005] provide an additional principle of proportionality (*"any application, system tool, or process should balance its utility with the rights to privacy of the involved individuals"*). In contrast, Lederer et al. [Lederer, Hong et al. 2004] outline five pitfalls to avoid: obscuring potential information flow, obscuring actual information flow, emphasizing configuration over action, lacking coarse-grained control, and inhibiting existing practice. Hong et al. [Hong, Ng et al. 2004] describe privacy risk models to analyze how well a system meets such principles or avoids pitfalls. These risk models are a set of information sharing questions pertaining to the social and organizational context in which the system is situated, and to the technology which is used to implement the system. To incorporate user perceptions, Adams and Sasse [1999] provide a privacy model based on information sensitivity, information receiver and information usage, in which each of the three factors interacts with the others.

4.3 Design Techniques

Incorporating principles and guidelines into working systems continues to pose challenges for designers. Improving privacy management requires addressing multiple conflicting concerns simultaneously [Hudson and Smith 1996], such as privacy vs. awareness, risks vs. benefits, control vs. overhead, and feedback vs. disruption. To complicate matters further, an acceptable solution to these tradeoffs is highly dependent on the user and the context.

Several techniques have been proposed and explored for the implementation of such principles. These include:

- encryption [Diffie and Hellman 1979];
- access control via preferences, policies, and roles [Edwards 1996; Wickramasuriya, Datt et al. 2004];
- mechanisms to reduce the burden of preference specification such as lightweight interfaces [Lau, Etzioni et al. 1999], or grouping and templates [Olson, Grudin et al. 2005; Patil and Lai 2005];
- automatic or manual control of the granularity of disclosed information [Dourish and Bly 1992; Lee, Girgensohn et al. 1997; Palen 1999; Consolvo, Smith et al. 2005];
- feedback via visualization [Gross, Wirsam et al. 2003], sound [Gaver, Moran et al. 1992], intelligent agents [Ackerman and Cranor 1999], and contextual disclosure [Kobsa and Teltzrow 2005];
- distortion of disclosed information [Boyle, Edwards et al. 2000];
- support for anonymity (or pseudonymity) [Appelt 1999];
- inference of appropriate awareness disclosure based on modeling [Begole, Tang et al. 2002].

Describing these techniques is beyond the scope of this chapter. The reader is referred to the cited works for details. In practice, no technique alone can satisfy all requirements and constraints. A typical awareness system would likely combine multiple privacy management approaches.

5. POSITIONING AWARENESS SYSTEMS

In order to choose the most relevant insights from prior work, we propose that awareness systems be positioned in a space of three independent dimensions (see *Figure 1*). We discuss these dimensions below.

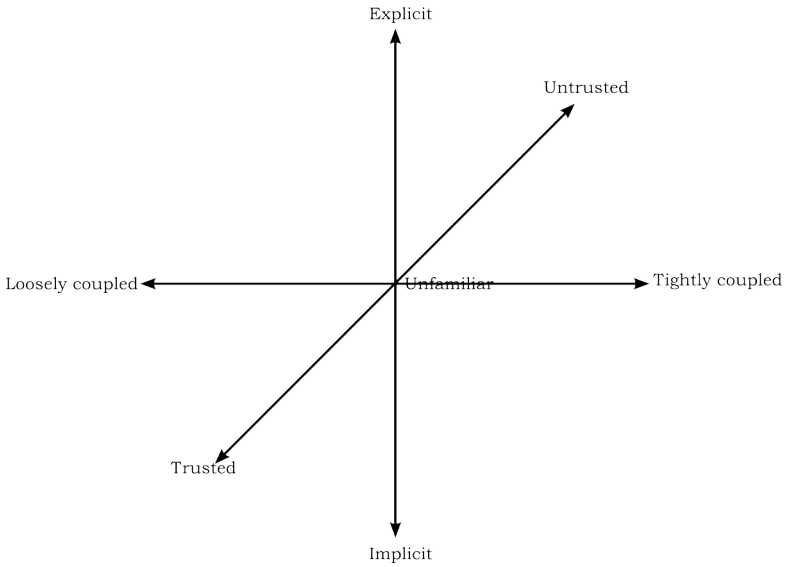

Figure 1. *Positioning Awareness Systems along Privacy-Relevant Dimensions*

5.1 Nature of Awareness Mechanisms

By their very nature, all awareness systems deal with capturing, storing, analyzing, disseminating, and/or displaying awareness information in some form. However, there is a distinction to be made between systems that are built specifically for awareness purposes (e.g., [Dourish and Bly 1992; Appelt 1999; Cadiz, Gupta et al. 2000]), and those that provide awareness implicitly by virtue of their use [Bellotti 1996]. For example, the primary purpose of email is to communicate the content of a message. Yet, by virtue of the timestamp, IP address, server names, and other header information, email reveals additional information implicitly. (It is also important to note that researchers have been exploring systems that could be built on top of other systems to make implicit aspects of awareness more explicit [Fisher and Dourish 2004; Froehlich and Dourish 2004].) Thus, awareness systems can be characterized to lie along a continuum ranging from explicit to implicit awareness functionalities (see *Figure 1*). For example, a system like Instant Messaging (IM) that provides communication mechanisms along with awareness could be positioned somewhere in the upper half [Nardi, Whittaker et al. 2000].

Systems that deal with awareness information explicitly, try to expose the benefits of awareness in a direct manner. As a result, they may also draw direct attention to the associated privacy issues. In contrast, when awareness is implicit or secondary to the function of a system, the primary attention of the user is on other aspects of the task carried out with the system (e.g., the user is much more likely to focus on the contents of an email message rather than on the IP address from which the email is being sent). Consequently, privacy aspects remain invisible in such cases [Bellotti 1996].

5.2 Activity Coupling:

User activities that awareness systems support lie along a continuum from loosely to tightly coupled [Olson and Teasley 1996; Olson and Olson 2000; Neale, Carroll et al. 2004]. For instance, the work of software developers working on two separate modules of the same program is less tightly coupled than that of a developer and a tester working on the same module.

As Olson and Olson [2000] explained, tightly coupled activities typically require "frequent, complex communication among the group members, with short feedback loops and multiple streams of information". Thus, when the work is tightly coupled, the awareness among collaborators of each other's activities is automatically improved as a side effect of more frequent and prolonged interactions. Given the shared (and often synchronous) focus on the same activity, awareness functionalities in these circumstances are mainly concerned with ensuring that the parties involved are aware of the focus and understanding of others [Dourish and Bellotti 1992]. On the other hand, when collaborative activities are loosely coupled, awareness is impoverished. In such cases, a variety of factors may affect awareness unfavorably. These include less frequent and asynchronous interaction between collaborators, less shared context, and simultaneous involvement in multiple tasks and projects [Olson and Teasley 1996; Pinelle and Gutwin 2003]. Thus, the looser the coupling, the greater is the need for external support by awareness systems.

Similarly, the privacy expectations in loosely coupled distributed activities can be expected to be greater than in the case of tightly coupled work. This may be caused by the same factors that engender impoverishment of awareness (i.e., less frequent and asynchronous interaction, less shared context, multi-tasking etc.) Additionally, if the work is geographically distributed across different countries, different privacy attitudes and laws of different nationalities need to be considered [Milberg, Burke et al. 1995; Bellman, Johnson et al. 2004]. In contrast, tightly coupled activities involve more focused (and often synchronous as well as collocated) interactions that allow one to monitor privacy closely.

5.3 Nature of Relationships

The nature of the relationships among various users of an awareness system forms the third dimension. These relationships can range from trusted and familiar (e.g., a colleague with whom one shares an office) to unfamiliar (but known, e.g. an employee in a different branch of the organization) to untrusted (e.g., a stranger who might read one's blog).

The degree of familiarity with the individual with whom one interacts is important in shaping attitudes and behaviors. For instance, greater familiarity reduces the importance of static awareness information [Danis 2000] because collaborators are likely to already know it, or can ask for it directly [Lederer, Mankoff et al. 2003]. In terms of privacy, Lederer et al. [Lederer, Mankoff et al. 2003] point out differences in privacy considerations when dealing with familiar as opposed to unfamiliar parties. While a great deal of research and legislation focuses on privacy protection from unknown organizations and people (e.g., governments, corporations, hackers), the other side of the continuum has received lesser attention. Yet, this side – ranging from the trusted to the unfamiliar – is important when dealing with awareness systems.

6. DESIGNING WITH PRIVACY IN MIND

Designers can utilize the above work of others to tackle privacy issues in their own awareness systems. Yet, we believe that in order to improve the privacy-sensitivity of awareness systems, a focus on privacy is needed right from the earliest conceptual phases of system development. As the term "awareness system" implies, the *purpose* of the system is to foster awareness. Hence, system designers have so far focused on providing awareness while privacy has only received secondary attention. We urge designers to treat privacy on an equal footing with awareness when building systems. The abovementioned principle of proportionality [Iachello and Abowd 2005] is a step in that direction. However, it deals mainly with a cost-benefit analysis of awareness and privacy to decide whether or not a system should be built at all. We take one step further and advocate that even after using this principle at the beginning of the design process to decide that an awareness system should be built, designers must still consider privacy at every subsequent step of the design cycle.

Two examples from our own research – one positive and one negative – illustrate why keeping privacy in mind at all stages of the design is essential.

6.1 Workplace Awareness Application

We designed an awareness application called mySpace to support the collaborative work of knowledge workers who were co-located in the same building [Patil and Lai 2005]. The mySpace application is a browser-based interactive visualization of a user's physical workplace that provides dynamically updated information about people, places and equipment. Recognizing that mySpace would lead to privacy concerns, we sought to empower users to manage their privacy according to their own preferences. Our initial intuition (based on experience with the organizational culture) was to allow users to specify one set of preference for their immediate team, and another set for all others. Additionally, we wanted to make the operation of the system transparent for users by disclosing all pieces of individual information to which mySpace had access. Yet, we feared that doing so would scare users into

selecting more privacy-protective preferences, thereby undermining the awareness benefits.

Instead of proceeding to build the system as envisioned, we conducted a user study of an early prototype. To our surprise, we found that our intuition was not aligned with users' desires. Users wanted to manage privacy at a finer grain by specifying preferences differently for multiple groups of contacts. Also, increased system transparency promoted trust in the system and seemed to reassure users that the system would honor their preferences. This resulted in increased awareness being provided to close, trusted groups of contacts.

Studying at the prototype stage how users deal with privacy aspects allowed mySpace to be sensitive to the practices of its target population. It also spared the costs and difficulties of correcting ill-chosen privacy management mechanisms retroactively after deployment.

6.2 Instant Messaging Privacy Plugin

Our experiences gained from improving privacy management in existing IM systems illustrate the weaknesses of the retrofitting approach. Based on several interviews and a survey of IM users [Patil and Kobsa, 2004; Patil and Kobsa, 2005], we had identified several avenues for improving IM privacy management. However, not having access to the servers of the commercial IM networks made the task of implementing our improvements challenging. We thus packaged our privacy management extensions as a plugin for the open source multiple IM client GAIM (now Pidgin). The plugin is called PRISM (for PRivacy-Sensitive Messaging). The architecture of the enhanced IM system is shown in *Figure 2*.

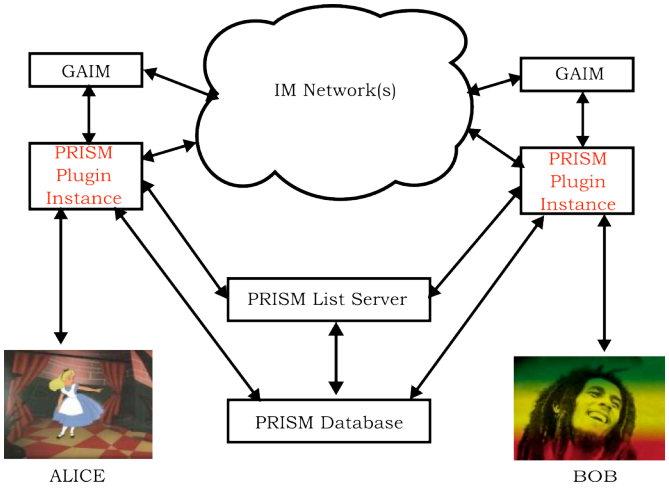

Figure 2. *System Architecture for PRISM*

As can be seen, PRISM maintains a separate server and a database in order to provide some of its privacy extensions. One enhancement that PRISM provides is to allow users to view the activities of others at a group level in order to facilitate a comparison with one's own activities. The PRISM database logs various user actions of interest, such as when users log in, log off, or change their availability status. Since PRISM does not have access to the servers of the IM networks, such logging is essential for generating visualizations of activities of a group of users. Ideally, the servers of the IM network would need to be extended to incorporate these functions.

More significantly, we often ran into limitations imposed by the specifics of the IM protocol(s). For instance, we aimed at empowering users to specify their privacy preferences differently for different groups of contacts. However, the IM protocol(s) lacked sufficient nuance to achieve this for all settings. For example, we were able to allow users to specify a different status for different groups but unable to provide a way to specify that only certain groups could view the length of time they were idle. Such deficiencies reflect inadequate attention to user privacy practices in the development of the IM protocol(s).

Finally, we aimed at generic privacy enhancements that did not rely on specifics of any single IM system. Because IM systems differ in the details of their protocols, and of their server implementations, ensuring a common cross-IM experience is a challenging task. For example, some IM systems allow one to broadcast the length of users' idle times, others don't; some IM systems allow multiple simultaneous logins, others don't. We found that catering to the lowest common denominator limits the extent to which the client can add, or improve upon, privacy management features. The only remaining option is to treat each protocol differently. The approach may confuse users since then the privacy management experience and expectations are no longer uniform.

Overall, PRISM serves as a cautionary example and illustrates the challenges and difficulties that designers are likely to face when attempting to retrofit privacy enhancements rather than designing systems with privacy in mind right from the outset.

7. CONCLUSION

Handling user privacy appropriately is a significant challenge for awareness systems. Inadequate attention to privacy issues can be a barrier to their success. To build awareness systems that are sensitive to the privacy needs of their users, designers ought to pay attention to privacy at every stage of system design. In order to be effective in this task, designers need to be aware of the various ways in which privacy can be understood. They should also pay attention to the special characteristics of the digital domain that may affect privacy management. Fortunately, designers can draw upon a substantial body of insights regarding privacy in the research literature. Appropriate techniques need to be chosen based on a careful evaluation of the context of work activities and social relationships within which the awareness system under consideration operates. Designing awareness

systems with privacy in mind has the potential to enhance privacy sensitivity significantly, and to empower users to satisfy their awareness as well as privacy needs.

8. ACKNOWLEDGEMENTS

Some of the research described was supported by NSF Grant No. 0205724. We wish to acknowledge our collaborator Jennifer Lai. We also wish to thank the subjects who participated in our study on mySpace. Finally, we are grateful to Mihir Mahajan for helping proofread various drafts of this chapter.

9. REFERENCES

Ackerman, M. S. (2000). The Intellectual Challenge of CSCW: The Gap between Social Requirements and Technical Feasibility. Human-Computer Interaction, 15, 179-203.

Ackerman, M. S., & Cranor, L. (1999). Privacy Critics: UI Components to Safeguard Users' Privacy. In CHI '99: CHI '99 Extended Abstracts on Human Factors in Computing Systems (p. 258-259). New York, NY, USA: ACM.

Ackerman, M. S., Starr, B., Hindus, D., & Mainwaring, S. D.(1997). Hanging on the 'Wire: A Field Study of an Audio-only Media Space. ACM Trans. Computer-Human Interaction, 4 (1), 39-66.

Adams, A. (1999). Users' Perception of Privacy in Multimedia Communication. In CHI '99: CHI '99 Extended Abstracts on Human Factors in Computing Systems (p. 53-54). New York, NY, USA: ACM.

Adams, A., & Sasse, M. A. (1999). Privacy Issues in Ubiquitous Multimedia Environments: Wake Sleeping Dogs, or Let Them Lie? In Seventh IFIP Conference on Human-Computer Interaction INTERACT'99 (p. 214-221). The Netherlands: IOS Press.

Agre, P. E., & Rotenberg, M. (Eds.). (1997). Technology and Privacy: The New Landscape. Cambridge, MA, USA: MIT Press.

Altman, I. (1975). The Environment and Social Behavior: Privacy, Personal Space, Territory, Crowding. Monterey, California: Brooks/Cole.

Appelt, W. (1999). WWW Based Collaboration with the BSCW System. In SOFSEM '99: Proceedings of the 26th Conference on Current Trends in Theory and Practice of Informatics on Theory and Practice of Informatics (p. 66-78). London, UK: Springer-Verlag.

Begole, J. B., Tang, J. C., Smith, R. B., & Yankelovich, N. (2002). Work Rhythms: Analyzing Visualizations of Awareness Histories of Distributed Groups. In CSCW '02: Proceedings of the 2002 ACM Conference on Computer Supported Cooperative Work (p. 334-343). New York, NY, USA: ACM.

Bellman, S., Johnson, E., Kobrin, S., & Lohse, G. (2004, November-December). International Differences in Information Privacy Concerns: A Global Survey of Consumers. The Information Society, 20 (5), 313-324.

Bellotti, V. (1996). What You Don't Know Can Hurt You: Privacy in Collaborative Computing. In HCI '96: Proceedings of HCI on People and Computers XI (p. 241-261). London, UK: Springer-Verlag.

Bellotti, V. (1997). Design for Privacy in Multimedia Computing and Communications Environments. In Technology and Privacy: The New Landscape (p. 63-98). Cambridge, MA, USA: MIT Press.

Bellotti, V., & Sellen, A. (1993). Design for Privacy in Ubiquitous Computing Environments. In ECSCW'93: Proceedings of the Third European Conference on Computer-Supported Cooperative Work (p. 77-92). Norwell, MA, USA: Kluwer Academic Publishers.

Boyle, M., Edwards, C., & Greenberg, S. (2000). The Effects of Filtered Video on Awareness and Privacy. In CSCW '00: Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work (p. 1-10). New York, NY, USA: ACM.

Boyle, M., & Greenberg, S. (2005). The language of privacy: Learning from Video Media Space Analysis and Design. ACM Trans. Computer-Human Interaction, 12 (2), 328-370.

Cadiz, J. J., Gupta, A., & Grudin, J.(2000). Using Web Annotations for Asynchronous Collaboration Around Documents. In CSCW '00: Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work (p. 309-318). New York, NY, USA: ACM.

Calore, M.(2006, September). Privacy Fears Shock Facebook. Wired News.

Consolvo, S., Smith, I. E., Matthews, T., LaMarca, A., Tabert, J., & Powledge, P.(2005). Location Disclosure to Social Relations: Why, When, & What People Want To Share. In CHI '05: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (p. 81-90). New York, NY, USA: ACM.

Cranor, L. F., Reagle, J., & Ackerman, M. S. (1999). Beyond Concern: Understanding Net Users' Attitudes About Online Privacy. AT&T Labs-Research Technical Report, TR 99.4.1.

Danis, C. M. (2000). Extending the Concept of Awareness to Include Static and Dynamic Person Information. SIGGROUP Bulletin, 21 (3), 59-62.

Diffie, W., & Hellman, M. E. (March 1979). Privacy and Authentication: An Introduction to Cryptography. Proceedings of the IEEE, 67 (3), 397-427.

Dix, A. J. (1990). Information Processing, Context And Privacy. In INTERACT '90: Proceedings of the IFIP TC13 Third International Conference on Human-Computer Interaction (p. 15-20). Amsterdam, The Netherlands : North-Holland Publishing Co.

Dourish, P. (1993). Culture And Control In A Media Space. In ECSCW'93: Proceedings of the Third European Conference on Computer-Supported Cooperative Work (p. 125-137). Norwell, MA, USA: Kluwer Academic Publishers.

Dourish, P., & Anderson, K. (2005). Privacy, Security... and Risk and Danger and Secrecy and Trust and Identity and Morality and Power: Understanding Collective Information Practices. Institute for Software Research (ISR) Technical Report, University of California, Irvine, UCI-ISR-05-1.

Dourish, P., & Bellotti, V. (1992). Awareness And Coordination In Shared Workspaces. In CSCW '92: Proceedings of the 1992 ACM Conference on Computer-Supported Cooperative Work (p. 107-114). New York, NY, USA: ACM.

Dourish, P., & Bly, S. (1992). Portholes: Supporting Awareness in a Distributed Work Group. In CHI '92: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (p. 541-547). New York, NY, USA: ACM.

Edwards, W. K. (1996). Policies and Roles in Collaborative Applications. In CSCW '96: Proceedings of the 1996 ACM Conference on Computer Supported Cooperative Work (p. 11-20). New York, NY, USA: ACM.

Fisher, D., & Dourish, P. (2004). Social and Temporal Structures in Everyday Collaboration. In CHI '04: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (p. 551-558). New York, NY, USA: ACM.

Fox, S. (2000). Trust and Privacy Online: Why Americans Want to Rewrite the Rules. Pew Internet & American Life Project.

Froehlich, J., & Dourish, P. (2004). Unifying Artifacts and Activities in a Visual Tool for Distributed Software Development Teams. In ICSE '04: Proceedings of the 26th International Conference on Software Engineering (p. 387-396). Washington, DC, USA: IEEE Computer Society.

Gaver, W., Moran, T., MacLean, A., Lövstrand, L., Dourish, P., Carter, K., & Buxton W. (1992). Realizing a Video Environment: EuroPARC's RAVE System. In CHI '92: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (p. 27-35). New York, NY, USA: ACM.

Goffman, E.(1959). The Presentation of Self in Everyday Life. Garden City, New York: Doubleday.

Grinter, R. E., & Palen, L. (2002). Instant Messaging In Teen Life. In CSCW '02: Proceedings of the 2002 ACM Conference on Computer Supported Cooperative Work (p. 21-30). New York, NY, USA: ACM.

Gross, T., Wirsam, W., & Graether, W. (2003). AwarenessMaps: Visualizing Awareness In Shared Workspaces. In CHI '03: CHI '03 Extended Abstracts on Human Factors in Computing Systems (p. 784-785). New York, NY, USA: ACM.

Heath, C., & Luff, P. (1991). Disembodied Conduct: Communication Through Video in a Multi-Media Office Environment. In CHI '91: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (p. 99-103). New York, NY, USA: ACM.

Herbsleb, J. D., Atkins, D. L., Boyer, D. G., Handel, M., & Finholt, T. A. (2002). Introducing Instant Messaging And Chat In The Workplace. In CHI '02: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (p. 171-178). New York, NY, USA: ACM.

Hudson, S. E., & Smith, I. (1996). Techniques for Addressing Fundamental Privacy and Disruption Tradeoffs in Awareness Support Systems. In CSCW '96: Proceedings of the 1996 ACM Conference on Computer Supported Cooperative Work (p. 248-257). New York, NY, USA: ACM.

Johnson, D. G.(1985). Computers and Privacy. Computer Ethics. Englewood Cliffs, NJ, USA: Prentice-Hall

Kahn, J. (2005, September). Yahoo helped Chinese to Prosecute Journalist. International Herald Tribune.

Kobsa, A., & Teltzrow, M. (2005). Contextualized Communication of Privacy Practices and Personalization Benefits: Impacts on Users' Data Sharing and Purchase Behavior. In D. Martin & A. Serjantov (Eds.), Privacy Enhancing Technologies: Fourth International Workshop, PET 2004 (p. 329-343). Springer.

Langheinrich, M. (2001). Privacy by Design - Principles of Privacy-Aware Ubiquitous Systems. In UbiComp '01: Proceedings of the 3rd International Conference on Ubiquitous Computing (p. 273-291). London, UK: Springer-Verlag.

Lau, T., Etzioni, O., & Weld, D. S. (1999). Privacy Interfaces For Information Management. Communications of the ACM, 42 (10), 88-94.

Lederer, S., Hong, J., Dey, A. K., & Landay, J. (2004). Personal Privacy through Understanding and Action: Five Pitfalls for Designers. Personal Ubiquitous Computing, 8 (6), 440-454.

Lederer, S., Mankoff, J., & Dey, A. K.(2003a). Towards a Deconstruction of the Privacy Space. In Ubicomp 2003 Workshop on Ubicomp Communities: Privacy as Boundary Negotiation. http://guir.berkeley.edu/privacyworkshop2003

Lederer, S., Mankoff, J., & Dey, A. K. (2003b). Who Wants to Know What When? Privacy Preference Determinants in Ubiquitous Computing. In CHI '03: CHI '03 Extended Abstracts on Human Factors in Computing Systems (p. 724-725). New York, NY, USA: ACM.

Lederer, S., Mankoff, J., Dey, A. K., & Beckmann, C. (2003). Managing Personal Information Disclosure In Ubiquitous Computing Environments. Technical Report, Computer Science Division, University of California, Berkeley, UCB-CSD-03-1257.

Lee, A., Girgensohn, A., & Schlueter, K. (1997). NYNEX Portholes: Initial User Reactions and Redesign Implications. In GROUP '97: Proceedings of the International ACM SIGGROUP Conference On Supporting Group Work (p. 385-394). New York, NY, USA: ACM.

Lessig, L.(1999). Code and Other Laws of Cyberspace. New York, NY, USA: Basic Books, Inc.

Mantei, M. M., Baecker, R. M., Sellen, A. J., Buxton, W. A. S., Milligan, T., & Wellman, B.(1991). Experiences in the Use of a Media Space. In CHI '91: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (p. 203-208). New York, NY, USA: ACM.

Mason, R. O. (1986, March). Four Ethical Issues of the Information Age. MIS Quarterly, 10 (1), 5-12.

Milberg, S. J., Burke, S. J., Smith, H. J., & Kallman, E. A. (1995). Values, Personal Information Privacy, and Regulatory Approaches. Communications of the ACM, 38 (12), 65-74.

Nardi, B. A., Whittaker, S., & Bradner, E. (2000). Interaction and Outeraction: Instant Messaging in Action. In CSCW '00: Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work (p. 79-88). New York, NY, USA: ACM.

Neale, D. C., Carroll, J. M., & Rosson, M. B. (2004). Evaluating Computer-Supported Cooperative Work: Models and Frameworks. In CSCW '04: Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work (p. 112-121). New York, NY, USA: ACM.

Negley, G. (1966). Philosophical Views on the Value of Privacy. Law and Contemporary Problems, 31 (2), 319-325.

Olson, G. M., & Olson, J. S. (2000). Distance Matters. Human-Computer Interaction, 15 (2/3), 139-178.

Olson, J. S., Grudin, J., & Horvitz, E. (2005). A Study of Preferences for Sharing and Privacy. In CHI '05: CHI '05 Extended Abstracts on Human Factors in Computing Systems (p. 1985-1988). New York, NY, USA: ACM.

Olson, J. S., & Teasley, S. (1996). Groupware in the Wild: Lessons Learned from a Year of Virtual Collocation. In CSCW '96: Proceedings of the 1996 ACM Conference on Computer Supported Cooperative Work (p. 419-427). New York, NY, USA: ACM.

Palen, L. (1999). Social, Individual and Technological Issues for Groupware Calendar Systems. In CHI '99: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (p. 17-24). New York, NY, USA: ACM.

Palen, L., & Dourish, P. (2003). Unpacking "Privacy" for a Networked World. In CHI '03: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (p. 129-136). New York, NY, USA: ACM.

Patil, S., & Kobsa, A. (2005a). Privacy in Collaboration: Managing Impression. In The First International Conference on Online Communities and Social Computing.

Patil, S., & Kobsa, A. (2005b). Uncovering Privacy Attitudes and Practices in Instant Messaging. In GROUP '05: Proceedings of the 2005 International ACM SIGGROUP Conference On Supporting Group Work (p. 109-112). New York, NY, USA: ACM.

Patil, S., & Lai, J. (2005). Who Gets to Know What When: Configuring Privacy Permissions in an Awareness Application. In CHI '05: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (p. 101-110). New York, NY, USA: ACM.

Pinelle, D., & Gutwin, C. (2003). Designing for Loose Coupling in Mobile Groups. In GROUP '03: Proceedings of the 2003 International ACM SIGGROUP Conference On Supporting Group Work (p. 75-84). New York, NY, USA: ACM.

Rachels, J. (1975). Why Privacy Is Important. Philosophy and Public Affairs, 4 (4), 323-333.

Romero, N., & Markopoulos, P. (2008). Grounding Privacy with Awareness Systems.

Samarajiva, R.(1997). Interactivity as though Privacy Mattered. In Technology and Privacy: The New Landscape (p. 277-309). Cambridge, MA, USA: MIT Press.

Stone, E. F., Gueutal, H. G., Gardner, D. G., & McClure, S. (1983). A Field Experiment Comparing Information-Privacy Values, Beliefs, and Attitudes Across Several Types of Organizations. Journal of Applied Psychology, 68 (3), 459-468.

Suchman, L. A. (1987). Plans and Situated Actions: The Problem of Human-Machine Communication. New York, NY, USA: Cambridge University Press.

U.S. Department of Health Education and Welfare. (1973). Records, Computers and the Rights of Citizens. Report of the Secretary's Advisory Committee on Automated Personal Data Systems, Publication No. 1700{00116.

Warren, S. D., & Brandeis, L. D. (1890, December). The Right to Privacy. Harvard Law Review, 4 (5), 193-220.

Westin, A. F. (1967). Privacy and Freedom. New York: Atheneum.

Westin, A. F. (1991). Harris-Equifax Consumer Privacy Survey 1991.

Wickramasuriya, J., Datt, M., Mehrotra, S., & Venkatasubramanian, N.(2004). Privacy Protecting Data Collection in Media Spaces. In MULTIMEDIA '04: Proceedings of the 12th Annual ACM International Conference On Multimedia (p. 48-55). New York, NY, USA: ACM.