

Privacy in the Open: How Attention Mediates Awareness and Privacy in Open-Plan Offices

Jeremy P. Birnholtz

Department of Communication
Cornell University
310 Kennedy Hall
Ithaca, NY 14853

jpb277@cornell.edu

Carl Gutwin

Department of Computer Science
University of Saskatchewan
110 Science Place
Saskatoon, Canada S7N 5C9

gutwin@cs.usask.ca

Kirstie Hawkey

Department of Electrical
Engineering and Computer Science
University of British Columbia
4044-2332 Main Mall
Vancouver, Canada V6T 1Z4

hawkey@ece.ubc.ca

ABSTRACT

The tension between privacy and awareness has been a persistent difficulty in distributed environments that support opportunistic and informal interaction. For example, many awareness systems that display ‘always-on’ video links or PC screen contents have been perceived as too invasive, even though functional real-world analogues, like open-plan offices, may provide even less privacy than their online counterparts. In this paper we explore the notion of privacy in open-plan real-world environments, in order to learn more about how it might be supported in distributed systems. From interviews and observations in four open-plan offices, we found that attention plays an important role in the management of both confidentiality and solitude. The public nature of paying attention allows people to build understandings of what objects in a space are legitimate targets for attention and allows people to advertise their interest in interaction. Our results add to what is known about how privacy works in real-world spaces, and suggest valuable design ideas that can help improve support for natural privacy control and interaction in distributed awareness systems.

Categories and Subject Descriptors

H.5.3 Group and Organization Interfaces

General Terms

Design, Human Factors

Keywords

Awareness, attention, privacy, group work

1. Introduction

Over the past twenty years, CSCW research has paid considerable attention to the design and development of awareness systems for distributed groups – including work on media spaces [e.g., 13, 42], instant messengers [e.g., 9, 18], and availability monitors [e.g., 3, 16]. The goal of these systems is to provide informal awareness and support casual interaction, which are hallmarks of effective work in face-to-face environments [29, 51].

One of the fundamental and persistently difficult design issues for

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GROUP 07, November 4-7, 2007, Sanibel Island, Florida, USA.

Copyright 2007 ACM 978-1-59593-845-9/07/0011...\$5.00.

these systems is how to balance the needs for awareness and communication with control over privacy [25]. By providing awareness information, systems necessarily reduce control over confidentiality; by allowing communication, they also allow interruption, thus reducing control over solitude. As a result, there has been some resistance to distributed awareness systems – in particular, those that have employed sharing of video, audio, or screen images, which have been perceived as invasive [8, 22].

In contrast, privacy seems not to be a major obstacle in some real-world settings, even when similar types of information are freely available. Open-plan office spaces, for example, are characterized by easy visibility of colleagues’ activity and frequent opportunistic interaction [4]. Confidentiality and solitude are reduced in these situations, as compared with conventional walled offices, but open-plan designs are both popular and effective [46]. This suggests that workers in these spaces are somehow able to mitigate concerns about reduced privacy.

How do they do it? How are open-plan office workers able to negotiate the balance between awareness and privacy while it remains such a stumbling block for distributed groupware? This question is vital for designers who wish to support awareness and collaboration in distributed work groups. If we can understand how privacy works in real-world open-plan offices, key lessons can be learned for the design of distributed awareness tools.

In this paper, we provide an initial set of answers to these questions. Our investigation consisted of visits to four commercial open-plan work environments, where we carried out observations and conducted detailed interviews. We gathered information about how people behaved when they interacted with one another, how they maintained awareness of what was going on in the office, how they felt about and maintained privacy in their workspace, and how they decided when to interrupt and talk to others.

Our findings show that attention plays an important role in balancing privacy and awareness in the open spaces we studied, and highlight two ways that attention is used in this process. First, the public nature of attention in an open space protects confidentiality by allowing for a shared understanding about what objects in the workplace are (and are not) legitimate targets for attention. Second, the public nature of paying attention to individuals can be a way to attract the attention of others, in service of initiating an interaction.

Our work raises several issues for designers to consider when building distributed awareness systems. Our study suggests that

making attention public in a distributed awareness system can help people to manage privacy more naturally, and that including traditionally illegitimate objects (such as computer screens) in a distributed awareness system may cause problems by artificially legitimizing that information as an object of attention.

2. Background and Literature Review

2.1 Awareness and Informal Interaction

While there are many definitions of awareness, the notion is rooted in the subtle and tacit forms of coordination that come from knowledge about the activities of others (see [43] for a detailed discussion). Schmidt's [43] review suggests that awareness information is gathered via "monitoring" for interesting behavior from others and "displaying" certain aspects of one's own behaviour that might be interesting to others, all based on prior experience or external knowledge. Such information can be used to spark informal interactions.

Informal workplace interactions—that is, those that are unplanned and occur opportunistically—have repeatedly surfaced as key elements of work practice, particularly where employees work in close proximity to one another [1, 29], and it is easy to access others or information about what they are doing [2, 4, 46, 51].

Many CSCW systems have sought to support similar behaviours in distributed groups – for example, through media spaces that use video to support mutual awareness and informal communication [13, 17, 44]. In these systems, users could see video images of remote colleagues in their offices, and used the visual information to assess availability and to initiate interaction.

Users, however, often reported that cameras intruded on their privacy [42]. Fish et al. [15] noted further that their system did not support the subtleties of starting and ending conversations. Building from this, Vertegaal [48] enhanced basic video by using eye-tracking to convey mutual gaze. Gaze, however, is only one of many cues used in initiating interaction.

Awareness information has also been provided via other media. Instant messaging and chat tools, for example, provide basic presence information, but availability is often indicated as a function either of recent keyboard activity or user-selected states such as 'busy,' or 'at lunch.' Keyboard activity, however, does not reliably indicate actual availability [16]. Moreover, users frequently do not update their status, and their actual availability may vary depending on the interrupter [11].

With this in mind, interruptions have garnered significant recent interest [e.g., 35]. When interrupting a colleague, there is utility in having some information about what that colleague is doing so that the interruption may be appropriately timed, but not so much that the information distracts the initiator [10] or invades the privacy of the target.

2.2 Privacy and Awareness

Privacy has been defined in several ways in CSCW, but two key common components emerge from prior work: confidentiality and solitude [6, 18]. Confidentiality refers to control over information moving outward from the self, and affects other people's access to personal information. Solitude refers to control over information moving toward the self (including interruptions), and determines how much of one's attention is consumed by that information. In terms of awareness, solitude refers to the extent to which one

receives cues about the state or presence of others, as well as the noticeability of those cues.

Where awareness in groups is concerned, there is a clear relationship between confidentiality and solitude. As Alice gets more information about what Bob is doing, his activities necessarily become less confidential. Moreover, the additional information that Alice receives about Bob increases the potential that she will be distracted from other tasks [3], thereby reducing her solitude. In this sense, the abundance of available awareness information in an open office has the potential to reduce both confidentiality and solitude. To some extent, this can be addressed by social norms [37]: for example, people concerned about the confidentiality of information on computer displays have been observed to re-position the displays so that it is difficult for others to view them.

In some cases, reduced privacy in open offices has been shown to have a negative impact. Industrial psychologists have studied employee satisfaction and privacy in open-plan spaces and found that, generally, moving workers from enclosed offices to an open-plan arrangement reduced satisfaction, increased distraction and reduced perceived privacy [12]. These findings must be interpreted cautiously, however, as participants were moved to open spaces to reduce costs, rather than to enable communication.

The larger point here is that there are clear tradeoffs between privacy and awareness in open offices [31]. Wulf and Hartmann [53] refer to a similar tradeoff in consideration of activity visibility in networks, noting that increased visibility of actions can improve coordination, but also increases the potential for surveillance. The literature suggests, however, that some of the people who work in these spaces value the interaction afforded by the open space, and have developed social norms and practices that facilitate periodic privacy when it is needed. Our question then becomes one of how these norms and practices work, and what lessons we might learn that could be applied in online environments.

In existing online awareness systems, one common approach to addressing tradeoffs involving solitude and confidentiality is reciprocity (i.e., the equal availability of persons and/or information). Fish et al. [15] note how visual reciprocity makes interaction feasible, by showing that others are present and available. Wiberg and Whittaker [52] observed 'conversational reciprocity;' their participants made themselves available for conversation in the hope that others would do the same.

In terms of confidentiality, some early media space developers were concerned about unequal amounts of information being shared (e.g., one user used a video camera and another elected not to). This led Clement [8] to discuss an "equality principle" in which "all parties to communication enter on a formally equal basis" (p. 78). Such a principle is implemented by Isaacs et al.'s [26] Piazza system, in which only communication media (e.g., video + audio, audio only) available to both interacting parties are used. One drawback to this approach is that, since each party has equal information about the other, it eliminates plausible deniability of receiving a request for interaction, which others have suggested is a positive attribute of some media [32, 38]

Even when there is equality of information sharing, however, there is ambiguity about the use of this information, which can heighten concerns about confidentiality [19, 24, 27]. Others using video for awareness have attempted to resolve this by modifying

the information – for example, by blurring a video image such that presence and identity can be detected, but not the details of activity [3]. This is effective in some cases, but may be ill-suited to sensitive environments such as people’s homes [39]. Fogarty et al. [16] took a different approach in MyVine, a system which integrates a range of information sources to provide estimates of availability. Only the estimate of availability is provided to observers in this case, and not the underlying details of activity.

From all of this, we see that initiation of interaction and control over confidentiality and solitude are closely related. On the one hand we have environments such as closed offices, where the explicit barriers in the environment mean that less awareness information is shared, and there tends to be less frequent interaction [2]. On the other hand, we have environments such as open-plan offices, where physical barriers are few, awareness information is readily available, and interaction is frequent. A range of situations, including media spaces and other online environments, exist between these extremes. We are particularly interested here in the open-plan category. In particular, we want to know why open-plan offices, which provide the least amount of privacy on this continuum, seem to raise fewer privacy concerns than the aforementioned media spaces. In other words, given that people are willing and able to work in open-plan offices spaces, how is it that they are able to exert adequate control over their solitude and confidentiality?

2.3 Attention As Mediator

Over a decade ago, Heath, et al. [22] reviewed the disappointing performance of media spaces in supporting collaboration, and raised many of these same issues, identifying attention as an important part of interaction. Schmidt [43] discusses the relationship between awareness and attention, noting that there is some disagreement on the extent to which awareness implies conscious or undivided attention. What seems clear, however, is that initiation of interaction with somebody involves a progressive increase in the amount of attention being paid to that individual.

As Vertegaal et al. [49] has pointed out, initiation of interaction is a process of mutual attention. If Alice wishes to interrupt Bob, she needs to know what he is doing (i.e., what he’s paying attention to), decide whether to interrupt or not, and then get his attention (i.e., attempt to become what he’s paying attention to, by paying attention to him). This can be difficult in media spaces for several reasons. First, cameras and monitors are not always aligned, so eye contact is not direct [22]. Second, using only a single camera for multiple remote viewers can erroneously convey attention, which has been addressed by using multiple cameras [44] or eye tracking [48]. Third, all of these systems require some conscious monitoring by the observed party to know when he/she is the focus of the observer’s attention.

All of this points to the different roles played by people in utilizing awareness information to initiate interaction and that these roles have unique information requirements. Benford et al. [5] and Rodden [41] show the influence of location and space on this information in their spatial model of awareness with its concepts of focus and nimbus. A range of systems have applied these ideas in using on-screen icon or avatar proximity or size as indicators of attention or engagement [14, 34]. These systems still rely heavily on monitoring by the person in the ‘focus’ role, however, which is not always a realistic expectation.

This prior work has shown that attention is an important mechanism. The remaining question is how attention is communicated in open plan spaces, such that confidentiality and solitude are not perceived to be under constant threat, and constant monitoring is not necessary. Despite extensive research on accommodating privacy concerns in distributed groups, few have studied open-plan offices with an eye toward better understanding the allocation of attention, and providing design suggestions for balancing privacy and awareness. These are the goals of the present study.

3. The Present Study: Attention and Privacy in the Open

The fact that people are willing and able to work in open-plan offices is testament to the fact that it is possible to adequately address issues related to confidentiality and solitude in these environments. We show below that, in the offices we studied, the control processes for both confidentiality and solitude are based on the combination of shared social structures and environmental constraints, along with the mechanism of attention. Attention is important for two main reasons. First, it is both a means for gathering information about the activities of others, as well as the basis for determining what can legitimately be attended to while gathering that information. Second, when attention is on public display, it becomes a way to attract the attention of others in the service of initiating (or postponing) an interaction.

The first reason raises the concept of attentional legitimacy. The offices we studied have conventions about which artifacts and activities are legitimate objects of attention. In the absence of physical constraints on attention (e.g., walls or doors), these understandings protect both solitude and confidentiality. For example, it is legitimate to notice objects in the periphery of a person’s workspace, but might be considered ‘nosy’ to stare at their computer screen.

The second reason involves the idea of public displays of attention. It has been noted that the initiation of interaction is a gradual process that usually involves more interaction [7, 30]. Heath and Luff [21] noted further that paying attention to someone is itself a communicative act – an implicit request for interaction that occurs through the common ground of attentional legitimacy. We will show in this paper that, in the open spaces we studied, there was a sort of ‘gradual engagement dance,’ that was, essentially, a progression of awareness-gathering ‘moves’ that involved certain types of responses from the other person. We argue that this dance plays a critical role in protecting solitude, and that the dance consists essentially of exchanges of attention.

In other words, interest in interaction on the part of the initiator is expressed by paying attention to his or her target in progressively more intrusive ways. This is rooted in a broader shared understanding of language-use conventions [7] and in spatial relationships such as proximity [20]. For example, if two people are not interacting, then it is not legitimate for one person to be close to and focus attention exclusively on the other – and doing so becomes a public and obvious act when the parties are physically in the same place.

3.1 Research Context

To explore attention in open-plan offices, we focused on environments where open spaces were used because of perceived creative, and not just economic, advantages. Four design and

architectural firms in a large North American city participated in our study. These firms, profiled in Table 1, range in size from 7 to 100 full-time staff. All were selected because they are well known in their fields for creative talent, and because they occupy well-designed open office spaces.

Table 1. Characterization of Firms and Interviews

Firm	Specialty	Employees	Number of Interviews	Subject IDs
A	Graphic Design	25	4	1-3, 8
B	Landscape Architecture	7	4	4-7
C	Interior Planning	100	10	9-18
D	Architecture	15	Observations only	

3.2 Methods

Our data gathering consisted of interviews and observations. Although it was not possible to conduct interviews and observations in the same firms, there was considerable consistency across all four firms. We therefore do not focus on differences between the firms, but instead treat them as a single data set. Where there were differences observed, however, these will be mentioned explicitly.

3.2.1 Interviews

We conducted eighteen semi-structured interviews that were 20-40 minutes in length. We spoke with designers, architects and project managers in Firms A, B and C. Interviews were conducted in late 2005 and early 2006, using a protocol that had been iteratively developed and piloted. All interviews were recorded and transcribed. The interviewer also toured each of the firms and took photographs.

Subjects were asked about their experience working in an open plan office, both with regard to perceived privacy and their interactions with others. We asked if and how they stay abreast of what colleagues are doing in the office; how often they interrupt colleagues and how they assess whether it is a good time to interrupt; and what they do when they require solitude or confidentiality.

Data analysis followed methods described by [36] and consisted of careful study of transcripts, followed by iterative coding first into categories based on the questions described above, and then into the two themes that we use to structure our presentation of results: attentional legitimacy and the public nature of attention.

3.2.2 Observations

Field observations were conducted at Firm D only. These were guided by our interview findings, and the data provide additional richness and fidelity. Observations were conducted near the centre of the open office, and recorded via detailed field notes [33]. There was no evidence that people modified their behaviour due to the presence of an observer (except when participants interacted with the observer, who would not otherwise be present).

Observations took place in 1-2 hour blocks at approximately regular intervals on 13 days over two months in early 2007, for a total of 25 hours. The intent in the observations was to focus on interpersonal interactions, defined here as any time a person

begins talking with one or more other persons with whom they were not already actively involved. Active involvement is indicated by temporally proximate utterances (within 1 minute of each other) and physical presence. Violation of either condition indicated a new episode. We focused particularly on how interaction episodes were triggered, where they took place, how long they lasted, and whom they involved.

In total, 618 interactions were observed involving 38 individuals. Some of these individuals were visitors or interacted infrequently, however, and the 14 most-frequent ‘interactors’ accounted for about 90% of the observed interactions.

Field data were analyzed by separating out individual interaction episodes and coding these according to how the episode began. After some iteration, a coding scheme was developed that included 7 types of interaction starts. The most common of these (55% of cases) was when one person walked over to another’s space. We will focus most heavily on this category in our results.

4. Results

4.1 Attentional Legitimacy in Information Seeking

People in the open offices we studied clearly rely on shared understandings about what are (and are not) legitimate targets of attention in open-plan spaces. While all of the offices we visited provided some explicitly separate spaces where people could go when they needed confidentiality or solitude, we do not focus on those areas because legitimacy plays a reduced role in governing behaviour in them, due to the physical constraints of the space.

We asked our interview participants what they pay attention to as they move about the office, and how they manage to maintain a sense of solitude and confidentiality in an environment where nearly everything is visible to others. Two clear themes emerged in participants’ responses. First, participants reported that they had low expectations regarding privacy in the open workspace. Second, they reported that, even though all is visible, certain spaces were more legitimate focuses of attention (i.e., were considered more public) than others.

4.1.1 Low Expectations of Privacy, High Expectations of Respect

In terms of expectations about privacy, our participants appeared to have a shared sense that, in the open space, there were virtually no guarantees of privacy, but there was a clear expectation that personal space would be respected. Moreover, as will be discussed below, it is typically clear when one’s personal space is being entered.

At the same time, many participants also shared the notion that they did not need privacy to do their work. When we asked how much privacy subjects felt they had in their workspaces, five explicitly responded with “zero” or “not a lot.” Only one participant mentioned that he did not have enough privacy, though two did make some suggestions (e.g., higher panels between desks) that might improve privacy for them. One participant also mentioned that he felt he had a lot of privacy, but that this was entirely because people did not spy on each other. Others indicated that minimal privacy was just part of life in their particular office and was not a positive or negative issue for them.

When asked if the lack of privacy was bothersome, most indicated that they “get used to it.” Fourteen participants indicated that when they did need to engage in confidential work such as employee reviews or personal financial matters indicated that they tended to conduct those activities away from their workspace. One participant noted, for example, that “for anything that is considered private from a privacy perspective, I do leave my desk.” (I17) One participant said he tended to work on these items in the evening when colleagues were not present, which is consistent with prior findings [e.g., 50].

Participants indicated that the lack of physical protections for privacy in their workspaces also moderated to some degree the personal activities in which they engaged. Three indicated that, if they needed to make a personal phone call they would go for a walk outside the office and use their cellular phone. This was also reflected in comments about what might be seen on their computer display:

I don't keep things on my screen that nobody should see. I don't see a reason why I should. I don't use this as a personal computer; I have one at home (I5).

I mean obviously you are not going to sit on like the internet [personal email, banking]... I think the open space kind of deters you from doing that because people can see your screens (I9).

Nonetheless, our observational data suggest that, at least in Firm D, some people seemed to modify their behaviour more than others. In one instance, our observer (and others in the office) was privy to a phone call about one person's personal financial details, even though that person knew the observer was there.

At the same time, however, all of our participants generally also seemed to expect respect for both solitude and confidentiality when they needed these things. This was evident in expectations regarding interruptions, as well as in norms about respecting the space and property of others in moving about the office.

The critical point from all of these examples is that there is a shared expectation of minimal privacy in the workspace that appears to mitigate people's concerns about confidentiality and solitude. This suggests that, despite shared understandings about legitimacy that we discuss below, low expectations for privacy (i.e., the known possibility that anybody can see what you're doing at any time, even if they're not likely to spy on you) constrain not just the attention of observers (as we describe below), but also the behaviour of the observed. In a sense, this reduces the cost of possible attentional error by an observer, where we can think of cost in social terms such as embarrassment.

4.1.2 *Understandings About Legitimacy*

In addition to low expectations that restrict behaviour, a second critical component of this story about legitimacy is a shared understanding of what actually are legitimate targets of attention. Generally, our participants noted that shared spaces and the peripheral areas of individual workspaces were more legitimate areas to focus on. Our observational data suggest that this is particularly true for people who don't know each other well, or who interact infrequently. Our participants were quite reflective in describing what they do and do not pay attention to as they move through the offices.

They described spaces that are explicitly public in nature, and are therefore legitimate targets of attention for anybody. In all of the

offices, for example, there are common spaces between at least some of the desks called “throw tables” or “collaboration spaces.” On the one hand these are likely an artifact of work in a field where large and numerous drawings are common, so space is needed to peruse and sort these. On the other, such spaces serve as semi-public venues for discussion. Descriptions and observations suggest that it is more socially acceptable to look at and discuss a drawing on a throw table than a drawing on somebody's desk:

So we use those tables as a way to collaborate and I can tell if they are obviously just conversing from a social perspective or if they are actually looking and doing work (I17).

Participants in the graphic design firm also indicated that everyone in their office has large, foam-core boards by their desks that are explicitly public and recognized as a way to advertise progress and to solicit feedback on work:

They are there for everybody to see and it feels different than something on your computer screen. Sort of an invitation to take a look (I3).

Moreover, such attention can foster opportunistic collaboration and discussion:

people as they are waiting...for something or getting their coat, they will look at it and make a comment and that is the way we can kind of use each other as a resource (I1).

Our observational data also showed that people frequently gathered around artifacts at shared tables, and others often saw these interactions and would join in. In this way, both the artifacts and the people in the shared space become legitimate targets of attention; and such spaces can provide useful cues as to what colleagues are doing and whether they might be free to interact.

As for individual workspaces, participants described the area in the centre of their space (usually the space adjacent to one's computer) as “more private” and peripheral spaces as “more public”. There appeared to be strong norms surrounding personal spaces in the open office, which most participants indicated to be important. One participant captured a notion expressed by many:

I feel that my space is respected. And I feel that if anybody needs to borrow, say, some of my reference material, they do come and ask. Nobody just comes and takes (I9).

Participants also indicated that they try to respect the personal space of others as they move about the office, and that there were certain areas that our interview subjects felt were not appropriate targets of attention. We asked, for example, if participants casually looked at colleagues' computer screens and physical desktops as they walked around the office. While three admitted to looking at screens occasionally, this was typically inadvertent:

it probably just happens because we have big computer screens and there are so many of them in the office. I wouldn't avoid doing it, but I don't really try to do it directly either (I7).

Another mentioned that doing so was “like looking in someone's window when you are outside on the street (I3).” This participant also indicated that, even if she saw something interesting or noteworthy on a colleague's screen, she would hesitate to talk to them about it for fear of seeming rude or nosy. This was confirmed by our observations; no interactions were observed that started as a result of something somebody saw on a screen in

walking by. There were, however, eight times when others were invited to look at a personal screen, as when providing assistance or discussing a particular detail. More often, however, participants printed out drawings and used these printouts for discussion.

It was considered acceptable, though, to look at people's workspaces from afar. Most people in the offices we visited can see one another, and although others always cannot be observed in detail, the first stage of information gathering involves attention to presence and availability from a distance:

Usually I can see the person that I need to talk to, like I'll just poke my head up (I1).

Because it's so open, you can generally look over and see if someone is on the phone or if they are free (I15).

If he's on the phone, I won't approach, but when I see that he is free, I'll get up and walk over to his station (I14).

Such examples were also prevalent in our observational data, and there were 16 instances where a person from one side of the office looked over at or walked to the other side and said something like "Oh, he's not there" and then walked away.

The long-distance cues discernable from such glances, such as being on the telephone or talking to another person, are usually implicit, but people can also give more explicit indications of their availability that can be seen from far away. Participants provided many examples, such as wearing headphones, which not only blocks out unwanted sound, but is also widely understood as an indicator that the wearer wishes not to be disturbed. These kinds of cues make it possible for both parties to carry out 'distant negotiation' that is lightweight for the observer and does not push the engagement to the point of interruption.

Audible activities were also reported to be an important and legitimate focus of attention. Participants expressed that they were able to maintain awareness of activities in the office through overhearing conversations and may at times interject when a conversation is pertinent to them. This was confirmed by our observational data, in which 28 (about 5% of the total) interactions were joined by a third party who was listening in.

4.2 Public Displays of Attention

Initiating an interaction involves a shift from paying attention to legitimate cues about a person's activities and availability from a distance, to a state of mutual attention or engagement. In other words, there is a gradual shift from one-sided attention (on the part of the initiator) to mutual attention. This raises the critical question of how this transition occurs. More specifically, this is a question of implicit signals of attention. How do I know what you are paying attention to? And, importantly, are you paying attention to me? We argue here that in the open offices we studied, this is accomplished by public displays of attention.

To learn about this, we asked our participants to describe in detail what it is like when they want to talk to or interrupt someone else in their office, and to describe situations where others interrupted them. We also, as noted above, paid close attention to the initiation of interactions in conducting our observations. The data suggest that attention is communicated in several ways.

4.2.1 Physical Proximity

In our observations, the majority (55%) of the interactions we saw occurred because one person walked up to another's workspace and interrupted that person. Proximity plays an important and

interesting role in this process, particularly in that it plays the dual role of allowing the initiator a more detailed view of the target's activity, while simultaneously making the target aware, by virtue of the initiator's presence in his/her periphery, that he or she is a target of the initiator's attention. In this way, proximity becomes a public display of the initiator's attention that is visible to the target and to everybody else in the office.

Interview participants reported that there are times when, even after walking over to somebody's workspace, they will walk away if the person in the target role does not notice them or seems particularly busy. We observed nine cases of this, in addition to two instances where an initiator started to walk away, probably assuming their target was too busy to notice them, but were then called back by the target, and subsequently began to interact.

Moreover, several interview participants distinguished between levels of activity, indicating they would not interrupt colleagues working "heads down" or "furiously in a corner." One participant did this at a very fine level of granularity:

If they were drawing or sketching something or researching something I could see with their facial expression, I wouldn't interrupt them. But if they were just, you know, casually sketching then I would think they were more approachable (I13).

It was also clear in our data that the people in the offices we studied are very conscious of proximity. Interview participants indicated, for example, that they approach people from particular directions, suggesting that the visibility of the approach is important:

So I always try to approach it from another angle, because I feel that the idea of having somebody come up behind you is uncomfortable (I3).

Generally people are pretty considerate, I guess, you know, no-one kind of creeps up behind you or anything like that (I6).

All participants talked, using various terms, about "hovering" just on the edge of the "private" part of a person's workspace. As one person stated,

People know that when they walk up to your workstation, you know, they don't come and sit on your lap. There is like that bubble – they kind of stand there, especially if you are on the phone (I12).

These public displays of attention are noticeable for reasons mentioned above – there are implicit social understandings about space and proximity. When violated, the violation itself can serve as an implicit request for further engagement:

Generally if you get within three feet of someone here then they'll know that you are wanting to talk to them. [...] You tend to notice so it is human instinct to look up. And so if you look up, and somebody makes eye contact with you, you kind of pick up that they want to talk to you (I15).

These claims were largely reinforced by our observations in Firm D, where people typically hovered within 2 or 3 feet of their target. There was typically a very brief moment when, if the target had not already heard the initiator approaching (see below), the initiator hovered until the target made eye contact (a public indicator of mutual attention) and the interaction began. When

targets were already engaged in conversation, however, hovering a few feet away was common.

Though we observed only two instances of aborted or postponed interactions (even when the target was on the phone or otherwise engaged) in Firm D, three of our interview participants did mention that sometimes they ask others to come back later:

If I am busy typing something, I mean, obviously I am not going to be rude. I'll acknowledge them and just say 'listen, I am just so busy now I can't talk. We'll talk later.' [...] You can also look up and make a verbal acknowledgement like 'give me two minutes. I'll finish this and come right over or, you know, just wait.' Or ... you can just engage in the conversation (I12).

Another subject mentioned that they can often take part in the negotiation even while on the phone with someone else:

And so if people walk by I'll either make a hand signal to tell them to wait a second, I'll be off in a minute, or I'll tell them that I'll come back to them – I'll just kind of cover the receiver up (I15).

The key point from all of this, however, is that proximity is important as a means for publicly displaying attention.

4.2.2 Audible Indicators of Attention

While visible proximity is an important way of conveying attention, sound can have a similar effect and was particularly prominent in our observations in Firm D. There were several ways in which sound conveyed attention.

First, the space had few sound-absorbing features, so it was easy to hear the footsteps of others as they approached. Given the ease of perceiving such sounds, it was not surprising that during many observed interactions in Firm D where, upon hearing the approach of another person, the target looked up and made eye contact with the initiator before he or she had arrived. There were, also cases, (including several involving the observer himself) where somebody would look up to see somebody approaching, but that person had another target in mind and did not return eye contact/visual attention. These cases were not formally recorded, but illustrate that it is through a combination of public cues that an initiator's focus of attention is discerned by a target.

Voices also provide public indication of attention, in three ways. First, some people, particularly those who sit near each other and didn't need to get up in order to talk, tended to look over to their target and begin talking. Often this talk was slow and tentative until the target of the interaction responded or made eye contact, and then the initiator would speak more quickly. In this way, a tentative use of the voice may substitute for visible proximity as an indicator of attention, when the parties are already proximate.

Second, some participants at Firm D tended to talk to their targets as they approached (24 episodes). This was particularly true when the pair had interacted recently, or were working together on a project and the initiator had some sense that the target was likely available. Here, voice combines with proximity to provide public information not just about the initiator's visual focus of attention, but also some notion of what the impending interaction is about.

Third and finally, voice is an indicator of the target's attention when that target is already involved in a conversation, and this conversation is overheard by the initiator. Sometimes existing involvement in a conversation is a cue to come back later. At

other times, however, as we explained above, it can be a cue that the subject of the conversation is relevant to the initiator as well. As we indicated, we observed several cases where initiators used these cues as indicators of attention.

5. Discussion

5.1 Theoretical Implications

While substantial additional fieldwork is necessary to determine the extent to which these results apply to other settings, there are some potentially useful theoretical implications from what we have learned. First, we have shown how careful consideration of interpersonal attention can, at least in the settings we studied, provide a critical bridge between theories of awareness and privacy, both of which play roles in the initiation of informal interaction. This shift in focus is consistent with recent interest in attention-aware interfaces at the individual level [e.g., 23, 40].

The relationship between awareness and attention has been studied before; however, our study raises two important points that have not been discussed. First, we present evidence suggesting that awareness requires at least some active attention (even if not necessarily exclusive). Second, and more significantly, our study shows that awareness is not always a reciprocal process, particularly in terms of interpersonal attention.

This latter point blurs Schmidt's above-mentioned distinction between "monitoring" and "displaying" behaviours in that, in the offices we studied, monitoring itself was on display. And the public nature of the display facilitated social behaviour that helped preserve privacy. In Rodden's [41] terms, this pushes the bounds of the focus/nimbus dichotomy in that there is a sense in which, when attention is displayed publicly, one's focus is actually a part of one's nimbus.

When these distinctions fade into the background, what becomes important is the conveyance of attention and facilitation/identification of mutual attention. In other words, how do we let people know, without interrupting them, that we want to know if they can be interrupted?

Several others have attempted to address this concern, often via the use of spatial metaphors that rely on the importance of proximity in the real world. In virtual reality environments (e.g., [5]), this may make sense. Attempts to apply a generalized spatial model [41] in environments where proximity is not physical, however, have not always resulted in systems that are intuitive or successful [14, 34]. One critical problem is that there is an implicit assumption that awareness information that is being provided is always being attended to. Having information available about the activities of others is not the same as knowing about those activities. Thus, our work extends the awareness literature in suggesting that mere awareness is not enough. There is a need to focus on mechanisms of attention [28] as they relate to interpersonal interaction and smooth movement from awareness to conversation.

This brings us to theories of privacy, which are concerned with the flow of information toward and away from the self. We believe that the predominant focus in the literature on control of this flow can be helpfully guided by the notion of interpersonal attention. Our results suggest that privacy is not a problem in open plan office spaces in part because low expectations cause some censoring of behaviour, but also because attention is on public display in ways that are easy and cheap (from a cognitive

standpoint) to process. The combination of a clear sense of attentional legitimacy and easy-to-process indicators of attention appeared sufficient to address most of the privacy concerns that our participants had. Thus, there may be utility in focusing less on control of the privacy of all information, and more on useful ways of indicating that attention is being paid. Indicating that attention is being paid is a question of awareness; doing this in ways that are subtle and noninvasive are, however, questions of privacy.

5.2 Practical Implications

Here we return to our initial motivation for this study and discuss issues for designers who are attempting to provide tools for awareness and interaction.

5.2.1 *Translating real-world proximity*

In the real-world situations we studied, physical proximity was an important cue in determining people's focus of attention. Although proximity does not specifically indicate attentional direction and focus, it does clearly indicate what in the environment can be attended to, and at what level of detail. In addition, when people are closer to you, it is much easier to determine their actual attentional focus through cues such as gaze direction and body orientation.

Proximity, however, has proven difficult to translate to distributed awareness systems. Some systems have done so with two-dimensional proximity of avatars [e.g., 14], but this approach does not provide the same degree of salience that occurs in the real world. Rather than attempt to directly map proximity onto screens, we suggest that designers separate the idea of proximity from the information that proximity conveys in the real world. We argue that it is this information that is critical in awareness, and that there are ways of representing this information that are better suited to online environments. In particular, this might be accomplished in two respects: making attention public, and supporting multiple levels of observation.

Make attention public

First, proximity in the real world allows for subtle (but increasingly salient, as physical proximity increases) and public indication of attentional focus. This suggests that there will be value in allowing people to see what others are paying attention to and, importantly, when others are paying attention to them, which requires that information about the direction and focus of a person's attention be both gathered (at the source) and displayed (at the remote sites).

The gathering of attention information can be divided into situations where attention is conveyed through deliberate user action (e.g., pointing and clicking on icons representing other users) and situations that require implicit tracking (e.g., glancing). Explicit attention at the interface is relatively easy to gather, and users can determine exactly what information will be obtained from the attentional focus (since this is controlled by the system). In contrast, tracking implicit attention is difficult. Until reliable gaze-tracking becomes commonplace, it is not easy to determine which part of a display a person is attending to. Some prior systems, such as the Hydra videoconferencing tool [44], simplify the problem by moving each person's representation to a separate display, although this requires additional hardware for each collaborator. In the short term, designers could move awareness information inside the application interface, such that it can only be obtained through explicit actions. While this has costs in terms

of user effort and application switching, it does fit well with the idea that people must expend additional effort to get additional awareness information, and would allow each "level" of a person's observing behaviour to become more noticeable.

In displaying this information to the users being monitored, designers may wish to re-think prior notions of reciprocity in awareness and communication systems. We argue that strict equality of information is not what matters, but rather that there should be an appropriate correlation between the amount of information that is being gathered by an observer, and the obviousness of that activity to others. In other words, the more information about Alice that Bob is getting by observing her, the more obvious Bob's observation should be to Alice.

Support multiple levels of observation

Closely related to the first point, the ability to provide observers with varying levels of information about others means that multiple levels of observation must be supported.

Our participants indicated that many interruptions in open-plan offices are postponed or halted because it is clear from a distance that a colleague is busy. In particular, it was easy for them to see from across the room that someone was talking to another person, or on the phone. Given the goals of reducing observer effort and protecting the solitude of the person being observed, distributed open-plan systems should indicate, at a basic level, whether a person is on the phone, in an existing conversation and whether another person in the system is already checking the availability of this person. On the one hand, providing this information may seem to violate confidentiality. At the same time however, our results indicate a clear tradeoff on this dimension. There is a very clear sense that people want others to be able to observe them at a distance – even without their knowledge – because this reduces the likelihood of poorly timed interruptions.

At the same time, it should also be possible to "get closer" to a person and learn more about what they are doing, perhaps getting their permission to do so in the process. For example, one could imagine displaying a number of potential indicators of activity, such as the number of active windows on the desktop, the title of the current window, or even a screen snapshot. For observers, this allows for the equivalent of increasing proximity in the real world, and allows them to better assess their colleagues' activity. For those being observed, this allows certain, trusted colleagues (who are likely to be the most frequent interrupters) to have more information about what they are doing, which can aid in coordination and communication.

5.2.2 *Considering attentional legitimacy*

Importantly, we are not advocating that all attention or all activities be made public. It was clear in our results that one of the reasons privacy was not a significant issue in the offices we observed was because people had low expectations and had developed workarounds (i.e., going to more private spaces). These behaviors should also be supported by designers of systems for distributed workgroups, by providing alternative private spaces for interaction and activities that are not observable by others.

At the same time, however, clearly legitimate public display spaces should also be provided. The idea of a public display space (such as the 'throw tables' or display boards described above) that is clearly legitimate could be a valuable addition to a distributed system. Placing objects in such a space requires explicit action,

but there is some evidence that people are willing to advertise their activities in this way (e.g., [45]), and this mechanism could be valuable as a way to invite opportunistic collaboration without compromising confidentiality.

Finally, designers must be aware of the legitimacy of the information they make available in a distributed awareness tool. Adding information that is not traditionally a legitimate target for attention into an awareness tool may implicitly promote that information to legitimacy. For example, our studies showed that computer screens are often not legitimate as objects of attention; therefore, this information must be treated very carefully in a distributed awareness system. Screen content is considered to be private (and known as such), making this information available in a distributed awareness tool might be seen to implicitly legitimize paying attention to this information. This is not to say that private information such as video, audio, and screen images cannot be part of an awareness system – only that this information should probably be given out only by explicit action. That is, there is no privacy violation if a person explicitly decides to share her screen image with another [47].

5.3 Limitations and Future Work

There are several limitations that must be considered in interpreting these results. First, our study has a limited scope in that we have focused on a small number of firms. While constrained focus allows for substantial detail and reflection at the level of the individual participant, it makes generalization problematic. More research is needed to determine how these results apply to other settings.

Moreover, our fieldwork focused more on observing and describing existing behavior, without exploring the origins and enforcement of these norms. Thus, more work is also needed to better understand how these behaviors came about, and how they are influenced by such factors as culture (organizational and national) and gender

Second, it is currently difficult to reliably and accurately discern an individual's focus of attention at any given moment, and to attribute motive to observed actions. While our use of two methods likely increases the validity of these results, our data rely heavily on observational data both retrospectively, by the interview parties, and our own observations as activities were taking place. Such data allow for high-level descriptions of what took place, which is important given the public nature of attention that we describe, but there is also a need for more micro-level analyses to supplement these descriptions.

There are several directions that this study opens up for further work. First, the observational work can be extended to other organizational settings, and we plan to see whether attention plays a similar role in open-plan settings such as research labs.

Second, we are currently implementing a distributed awareness and communication system called OpenMessenger that will allow us to explore and test our results and design principles. OpenMessenger is similar to existing IM applications, but adds several of the ideas discussed above: explicit representation of the focus and degree of each person's attention (both audible and visual); several levels of information about a person (with a corresponding increase in the noticeability of these actions); configuration of which information sources will be conveyed; and display space for each person that can show work artifacts such as

text, images, links, documents, and screen contents. We plan to carry out both focused experiments and longer-term usage studies with this prototype to explore the implementation of design ideas, and user comfort with and reaction to different representations.

6. Conclusions

We have considered the question of why people are reluctant to share certain kinds of personal information through distributed awareness systems, even though similar information is often readily available in face-to-face environments. We carried out interviews and observations in four open-plan offices to better understand how people manage privacy and awareness, and to explore this discrepancy between real and virtual environments. We found that attention, the public display of attention, and attentional legitimacy played major roles in mediating both awareness and privacy. These ideas contribute to existing discussions of informal interaction and casual awareness, and provide a number of issues for designers to consider in developing distributed awareness systems.

7. References

- [1] Allen, T.J. *Managing the Flow of Technology*. MIT Press, Cambridge, MA, 1977.
- [2] Becker, F. and Sims, W. *Offices That Work: Balancing Communication, Flexibility and Cost*, Cornell University, International Workplace Studies Program, iwsp.human.cornell.edu/pubs/pdf/IWS_0002.PDF, 2001, 1-77.
- [3] Begole, J., Matsakis, N. and Tang, J., Lilsys: Inferring Unavailability Using Sensors. In *ACM CSCW(2004)*, 511-514.
- [4] Bellotti, V. and Bly, S., Walking Away from the Desktop Computer: Distributed Collaboration and Mobility in a Product Design Team. In *ACM CSCW(1996)*, 209-218.
- [5] Benford, S., A spatial model of interaction in large virtual environments. In *ECSCW*, (1993), 109-124.
- [6] Boyle, M. and Greenberg, S. The Language of Privacy: Learning from Video Media Space Analysis and Design. *ACM TOCHI*, 12, 2 (2005), 328-370.
- [7] Clark, H.H. *Using language*. Cambridge University Press, New York, 1996.
- [8] Clement, A. Considering privacy in the development of multi-media communications. *CSCW*, 2, (1994), 67-88.
- [9] Cockburn, A. and Greenberg, S., Making contact: getting the group communicating with groupware. In *ACM GROUP*, (1993), 31-41.
- [10] Dabbish, L. and Kraut, R., Controlling Interruptions: Awareness displays and social motivation for coordination. In *ACM CSCW(2004)*, 182-191.
- [11] Davis, S. and Gutwin, C., Using Relationship to Control Disclosure in Awareness Servers. In *Graphics Interface* (2005), 75-84.
- [12] De Croon, E.M., Sluiter, J.K., Kuijer, P.P.F.M. and Frings-Dresen, M.H.W. The Effect of Office Concepts on Worker Health and Performance: A Systematic Review of the Literature. *Ergonomics*, 48, 2 (2005), 119-134.
- [13] Dourish, P. and Bly, S., Portholes: Supporting awareness in a distributed work group. In *ACM CHI(1992)*, 541-547.
- [14] Erickson, T., Smith, D., Kellogg, W., Laff, M., Richards, J. and Bradner, E., Socially translucent systems: Social proxies, persistent conversation and the design of "Babble". In *ACM CHI*, (1998), 72-79.

- [15] Fish, R., Kraut, R. and Root, R., Evaluating Video as a Technology for Informal Communication. In *ACM CHI*, (1992), 37-48.
- [16] Fogarty, J., Lai, J. and Christensen, J. Presence versus Availability: The Design and Evaluation of a Context-Aware Communication Client. *International Journal of Human-Computer Studies (IJHCS)*, 61, 3 (2004). 299-317.
- [17] Gaver, W., Moran, T., MacLean, A. and Lovstrand, L., Realizing a Video Environment: EuroPARC's Rave System. In *ACM CHI*, (1992), 27-35.
- [18] Greenberg, S., Peepholes: Low cost awareness of one's community. In *ACM CHI*(1996), 206-207.
- [19] Grudin, J. Desituating action: digital representation of context. *Human Computer Interaction*, 16, 2-4 (2001). 269-286.
- [20] Hall, E.T. Proxemics. *Current Anthropology*, 9, 2/3 (1968). 83-108.
- [21] Heath, C. and Luff, P. Collaboration and Control: Crisis Management and Multimedia Technology in London Underground Line Control Rooms. *CSCW*, 1, 1 (1992). 24-48.
- [22] Heath, C., Luff, P. and Sellen, A., Reconsidering the virtual workplace: flexible support for collaborative activity. In *ECSCW*, (1995), 83-99.
- [23] Horvitz, E., Kadie, C., Paek, T. and Hovel, D. Models of attention in computing and communication: from principles to applications. *Commun. ACM*, 46, 3 (2003). 52-59.
- [24] Huang, E.M. and Mynatt, E.D., Semi-Public Displays for Small, Co-located Groups. In *ACM CHI* (2003), 49-56.
- [25] Hudson, S.E. and Smith, I., Techniques for addressing fundamental privacy and disruption tradeoffs in awareness support systems. In *ACM CSCW*, (1996), 248-257.
- [26] Isaacs, E., Tang, J. and Morris, T., Piazza: A desktop environment supporting impromptu and planned interactions. In *ACM CSCW*(1996), 315-324.
- [27] Jancke, G., Venolia, G.D., Grudin, J., Cadiz, J. and Gupta, A., Linking public spaces: Technical and social issues. In *ACM CHI*(2001), 530-537.
- [28] Kastner, S. and Ungerleider, L.G. Mechanisms of Visual Attention in the Human Cortex. *Annual Review of Neuroscience*, 23, (2000). 315-341.
- [29] Kraut, R., Egidio, C. and Galegher, J., Patterns of Contact and Communication in Scientific Research Collaboration. In *ACM CSCW*(1988), 1-12.
- [30] Kristoffersen, S. and Ljungberg, F., An Empirical Study of How People Establish Interaction: Implications for CSCW Session Management Models. In *ACM CHI*(1999), 1-8.
- [31] Lai, J., Levas, A., Chou, P., Pinhanez, C. and Viveros, M. Bluespace: personalizing workspace through awareness and adaptability. *International Journal of Human-Computer Studies (IJHCS)*, 57, (2002). 415-428.
- [32] Lederer, S., Hong, I., Dey, K. and Landay, A. Personal privacy through understanding and action: five pitfalls for designers. *Personal and Ubiquitous Computing*, 8, 6 (2004). 440-454.
- [33] Lofland, J. and Lofland, L.H. *Analyzing Social Settings: A Guide to Qualitative Observation and Analysis*. Wadsworth Publishing Co, Detroit, 1995.
- [34] McEwan, G. and Greenberg, S., Supporting social worlds with the community bar. In *ACM GROUP*, (2005), 21-30.
- [35] McFarlane, D.C. and Latorella, K.A. The scope and importance of human interruption in human-computer interaction design. *HCI*, 17, 1 (2002). 1-61.
- [36] Miles, M.B. and Huberman, A.M. *Qualitative Data Analysis: An Expanded Sourcebook*. Sage Publications, Thousand Oaks, 1994.
- [37] Moor, J.H. Towards a theory of privacy in the information age. *ACM SIGCAS*, 27, 3 (1997). 27-32.
- [38] Nardi, B., Whittaker, S. and Bradner, E., Interaction and outercation: Instant Messaging in action. In *ACM CSCW*(2000), 79-88.
- [39] Neustaedter, C., Greenberg, S. and Boyle, M. Blur filtration fails to preserve privacy for home-based video conferencing. *ACM TOCHI*, 13, 1 (2006). 1-36.
- [40] Roda, C. and Thomas, J. Attention aware systems: Theories, applications and research agenda. *Computers in Human Behavior*, 22, (2006). 557-587.
- [41] Rodden, T., Populating the application: a model of awareness for cooperative applications. In *ACM CSCW*, (1996), 87-96.
- [42] Root, R., Design of a Multi-Media Vehicle for Social Browsing. In *ACM CSCW*(1988), 25-38.
- [43] Schmidt, K. The problem with 'awareness'. *CSCW*, 11, (2002). 285-286.
- [44] Sellen, A., Buxton, B. and Arnott, J., Using spatial cues to improve videoconferencing. In *ACM CHI*, (1992), 651-652.
- [45] Smale, S. and Greenberg, S., Broadcasting information via display names in instant messaging. In *ACM GROUP*(2005), 89-98.
- [46] Teasley, S., Covi, L.A., Krishnan, M.S. and Olson, J.S. Rapid software development through team collocation. *IEEE Transactions on Software Engineering*, 28, 6 (2002).
- [47] Tee, K., Greenberg, S. and Gutwin, C., Providing artifact awareness to a distributed group through screen sharing. In *ACM CSCW*(2006), 99-108.
- [48] Vertegaal, R., The GAZE groupware system: mediating joint attention in multiparty communication and collaboration. In *ACM CHI*, (1999), 294-301.
- [49] Vertegaal, R., Shell, J.S., Chen, D. and Mamuji, A. Designing for augmented attention: towards a framework for attentive user interfaces. *Computers in Human Behavior*, 22, (2006). 771-789.
- [50] Werner, C.M. and Haggard, L.M. Avoiding intrusions at the office: privacy regulation on typical and high solitude days. *Basic and Applied Social Psychology*, 13, 2 (1992). 181-193.
- [51] Whittaker, S., Frohlich, D. and Daly-Jones, O., Informal Workplace Communication: What is It Like and How Might We Support It? In *ACM CHI*(1994), 131-137.
- [52] Wiberg, M. and Whittaker, S. Managing availability: supporting lightweight negotiations to handle interruptions. *ACM TOCHI*, 12, 4 (2005). 356-387.
- [53] Wulf, V., Hartmann, A. The ambivalence of network visibility in an organizational context. in Clement, A.K., P. Wagner, I. ed. *NetWorking: Connecting workers in and between organizations*, North Holland, Amsterdam, 1994, 143-152.