# Using Social Geometry to Manage Interruptions
and Co-Worker Attention in Office Environments

Maria Danninger
Interactive Systems Labs
Universität Karlsruhe
Karlsruhe, Germany
maria.danninger@ira.uka.de

Roel Vertegaal
Human Media Lab
Queen's University
Kingston, ON, Canada
roel@cs.queensu.ca

Daniel P. Siewiorek
HCI Institute
Carnegie-Mellon University
Pittsburgh, PA, USA
dps@cs.cmu.edu

Aadil Mamuji
Human Media Lab
Queen's University
Kingston, ON, Canada
mamuji@cs.queensu.ca

*Abstract*

Social geometry is a novel technique for reasoning about the engagement of participants during group meetings on the basis of head orientation data provided by computer vision. This form of group context can be used by ubiquitous environments to route communications between users, or sense availability of users for interruption. We explored problems of distraction by co-workers in office cubicle farms, applying our method to the design of a cubicle that automatically regulates visual and auditory communications between users.

*Keywords: Social Geometry, Attentive User Interface*

## 1 Introduction

The large increase in the number of computing appliances that surround today's user comes at the cost of an ever-increasing number of potentially disruptive notifications [8,20]. With the increased volume of wireless communications, users are now frequently interrupted by auditory or visual alerts from cell phone calls, email and instant messaging notifications. This is because each of the user's computing appliances is designed to independently vie for the attention of the user with each message delivered, regardless of the user's current engagement with other devices or people. Meier showed early on that interruptions that distract a user from a focus task are an important source of work-related stress [17]. By designing computing devices such that they sense and respond to a worker's availability for communications we may be able to develop notification strategies that are less disruptive, and better coordinated between devices and users.

### 1.1 Motivation

Interruptions generated by office workers' computing devices are especially troublesome in situations where co-workers share the same space, such as in shared offices or cubicle farms, and in public locales. This is because auditory notifications generated for one particular user may lead to disruption of other co-workers' focus tasks. However, the act of interfering with a co-worker's focus task is not limited to interruptions generated by computer systems. Active ad-hoc and co-located group communications among co-workers may be equally distracting and problematic. This paper explores how regulating communications by ubiquitous sensing and reasoning about the social orientation of co-workers may alleviate such problems in future office scenarios. To address the problem of disruption in office environments, researchers are developing context-aware systems that sense when users might be available for communications [1,3,6,11,10,21]. Allowing such systems to function reliably, however, may require sensing and reasoning about large numbers of complex contextual variables. In particular, knowledge about participation in group activities provides an important contextual indicator for a person's interruptibility. This is because users engaged in group interactions, such as when speaking to an audience, are less likely to be available for outside communications. Conversely, social orientation of users is also useful for determining when users *are* interested in communications. However, determining when a user is engaged in conversation with another person is not a straightforward problem. According to Edward Hall's theory of proximity in communicative space [9], interaction with other persons takes place within a certain distance range, which varies with culture. Proximity levels include intimate, personal, social and public. Intimate space is used for touching one another, personal space is used for conversations, social distances involve groups of people, and public spaces go beyond this. Moreover, Argyle showed that gaze and mutual gaze are an important part of face-to-face communication and influence social intimacy [2]. For example, during face-to-face interactions, users are likely to orient themselves towards each other before initiating a conversation [7]. However, augmenting an environment with eye trackers for sensing such behavior is expensive. In this paper, we therefore explored the use of more course-grained indicators of attention, including user proximity and head orientation towards objects or people of interest.

### 1.2 Social Geometry and Paper Outline

One of the defining features of social groups is that they typically present geometric clusters of people interacting with each other. Figure 1 shows how the structure of these clusters may allow classification into group types based on relatively simple geometric properties of relationships between individual bodies. Examples of group geometries are the 1-to-many or lecture geometry (Figure 1a), where most members are oriented towards a single speaker, or the many-to-many geometry of round table meetings (Figure 1b). In both geometries, the orientation and proximity of participants leads to a distinctive clustering, distinct from an arbitrary grouping, as shown in Figure 1c.

In this paper, we will show how such Social Geometries can be detected in arbitrary arrangements of individuals using simple computer vision techniques. We present a new method for reasoning about group membership based on straightforward geometrical and temporal properties of mutually shared attention between people. The Social Geometry provided by clusters of individual users provides a powerful yet computationally inexpensive method of analysis, one that allows use of coarse grained measures such as body orientation in reasoning about attention for and membership of a social group.

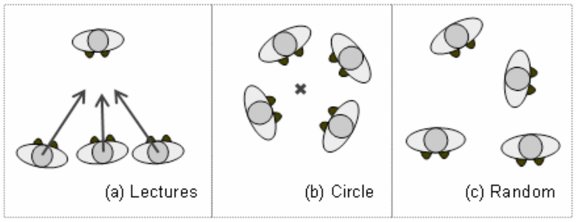

**Figure 1**. Examples for social groups of 4 persons in different group geometries.

After our review of the literature, we present our prototype computer vision engine, which allows simple capturing of user head location and orientation data. We then discuss how our social geometry engine deduces social group membership on the basis of simple spatial-temporal properties of mutually shared attention between users. Finally, we discuss how we applied our engine in the conceptual design and implementation of a ubiquitous computing environment for a future office. We implemented an office cubicle that automatically mediates interruptions between co-workers on the basis of social geometry data provided by overhead computer vision cameras.

## 2 Background

The management of interruptions generated by communication technologies in office environments has recently become an important topic of study in HCI [11,12]. The deployment of algorithms that reason about the importance of messages generated by computers [10], and the implementation of notification interfaces that determine when and through what channel to notify users have been crucial first steps towards the development of an integral and orchestrated management of user attention in their interactions with computers [20].

Notification Interfaces, developed by Horvitz [11], are user interfaces that use Bayesian reasoning to identify how and when to notify a user of an incoming message. While Horvitz experimented with the use of measures of head orientation of users towards a computer system, most notification interfaces are characterized by their ability to *reason* about user attention as a resource, rather than actively *sense* attention. We believe the development of systems that directly measure overt characteristics of user attention is crucial for the success of notification systems. The Attentive User Interface (AUI) paradigm [20] has tried to address this problem by managing user attention through a combination of explicit sensing and reasoning techniques.

### 2.1 Attention Management for Remote Group Communications

One category of applications in which user attention and interruptibility has been deployed successfully is the management of distributed group conversations.

Many desktop videoconferencing systems are ineffective due to deficiencies in gaze awareness and a lack of social orientation information. A number of systems have been designed to allow gaze awareness in teleconferences. CAVECAT [14] and Hydra [19] video systems demonstrated very early the importance of understanding how humans interact with and communicate through technology. The GAZE [27] video conferencing system uses eye trackers to measure where participants look in order to communicate who is talking to whom during mediated group collaborations.

The Attentive cellphone [25] is an example of a mobile attentive user interface that uses eye tracking and speech activity measures to determine when a user is engaged in a conversation. To preempt interruption, the phone communicates user engagement estimates to people in its contact list. SenSay [21] is a mobile phone that adapts to changing user states by manipulating

**Figure 2.** Retroreflective markers on people's headsets ease the motion capture process.

ringer volume, vibration, and phone alerts. The SenSay system uses a variety of sensors to determine when the user is busy and not to be interrupted.

## 2.2 Attention Management for Co-located Interactions

There are, to date, few examples of systems that measure participant attention in order to manage interruption during co-located meetings. Stiefelhagen [22] developed a system that tracked head orientation of participants using computer vision and a neural network during four-person meetings. Although computationally expensive, Stiefelhagen's system has been successfully applied to automated editing of meeting recordings [23]. In the auditory domain, Basu and Pentland developed a pair of Smart Headphones that detected and relayed sounds in the environment through to the user's headset, but only if they were classified as human speech [4]. Mueller and Karau improved upon this concept with Transparent Headphones, headphones augmented with an amplifier that picked up and modified sound in real time, before user auditioning [18]. One of the applications of this system was to help a user listen to mp3s while still being accessible to surrounding individuals. By mounting proximity sensors on the headphones, the system detected when a person approaches a user, presumably to engage in conversation. Upon co-location of individuals, the headphones would pause the mp3 player.

## 3 Tracking Head Orientation in Large Groups Using Real-Time Computer Vision

A key problem with many of the sensing techniques deployed in the above systems is that they typically do not scale well across wide areas or large numbers of users. In particular, facial recognition systems and eye contact sensors can be expensive and cumbersome to install. Associated hardware cost and computational complexity make these technologies, for now, unsuitable for widespread deployment in office environments. To address this issue, we developed a computer vision system that uses head orientation to determine joint attention between individuals across wide areas, through simple overhead computer vision. While eye contact between individuals provides one of the most direct and reliable measures of engagement between two individuals [26], head tracking provides a more tractable problem when dealing with large groups [23]. The use of an overhead camera provides the additional advantage of making head location data readily available. Studies show that head orientation provides a reasonable estimate of a person's direction of regard, accurate to within 15 degrees of visual angle [22].

To allow for scalability, real-time performance and reliability at low cost, we based our prototype tracking system on the detection of retro-reflective markers. To ease the motion capture process, we use a webcam augmented with infra-red LEDs, to track retro-reflective markers throughout the environment. Figure 2 shows the image from a camera located in the ceiling of an office environment. In our system, users wear headsets with retro-reflective markers. This facilitates identification of head location and orientation in two dimensions and at very low camera resolutions. We aim to develop future versions that track head orientation without the need for retro-reflective markers.

### 3.1 Tracking Performance

We evaluated various placements of retroreflective markers for use in our head tracking system to optimize robustness. The optimal position is one that is tilt independent, asymmetric and uniquely identifiable by the system. The use of unique identifiers allows support for identification of individuals. It also allows robustness to movements in and out of the camera's field of view. The markers are curved to ensure tilt independence. This allows the camera to track two-dimensional orientation robustly as people move and tilt their heads. Asymmetry in the markers reduces the ambiguity in directional information obtained by the system. Our motion tracking algorithm uses the central marker for coarse detection of a user's location. Four markers are placed to the right of the central marker. The first marker provides directional information of the user. By choosing to remove or keep any of the last three markers, our algorithm detects a binary pattern which it uses to uniquely identify eight users.

Next, we discuss how the coordinates and orientation data of individual users' head movements is combined by our social geometry engine for real-time classification of groupings.

## 4 Social Geometry Analysis

To establish social group membership from individual head movement measurements, we implemented an engine that performs analysis of the social geometries formed by the virtual connections between a group of co-oriented heads. By our definition, individuals may share mutual attention if their heads are oriented towards each other for a certain minimum percentage of time, and located within a certain maximum social distance from one another [2]. Rather than determining group geometries on a frame-by-frame basis, our system uses a dynamical approach that relies on statistical definitions of co-orientation over time.

### 4.1 Estimating Social Interaction

Figure 3 shows how we defined shared attention as a network link in a graph that connects two people with each other.

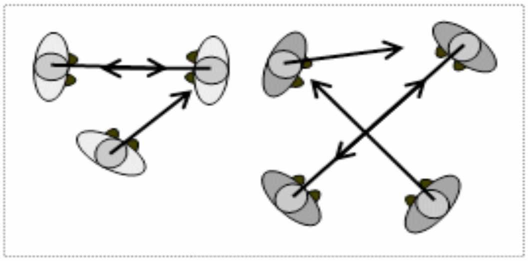

**Figure 3.** Social interactions are represented by attentional links from one person to another.

In more mathematical terms, shared attention can be represented using a binary relationship between individuals. Our system determines the relationship to either be true or false, i.e., present or not present on the basis of a fuzzy statistical model. We modeled shared attention as ordered pairs of a Cartesian product *Person x Person*, yielding a function *(Person1, Person2)* that indicates whether *Person1* is likely to be paying attention (i.e., be oriented towards) *Person2*. This function is directional in that *(Person1, Person2)* may indicate an attentional relationship even when *(Person2, Person1)* may not. Situations are thus modeled as digraphs *G(P, A)* with people *(P)* as nodes and attentional links *(A)* as directed edges. This allows reasoning about social geometries with straightforward graph theoretical constructs. To determine the attentional weight of the connection between each node in the graph, the system first calculates each person's field of view as a wedge with the person's current angle of orientation as its center ( Figure 4).

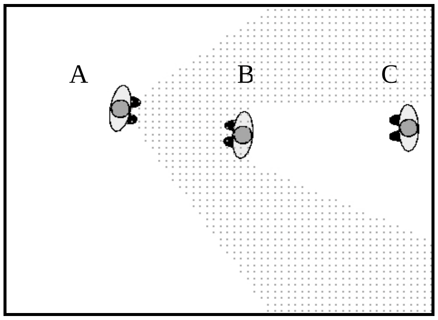

**Figure 4.** Obstacles such as people are subtracted from the calculated field of view (grey) by determining their shadow effect (white).

### 4.2 Calculating Pairwise Attention

Next, the system determines the amount of shared attention for each pair of nodes in the graph. It does so by finding the optimal potential communication partner for each node in the graph (see Fig 4). To find the best match for node A's communication partner, the system only considers nodes within A's field of view that are oriented towards A. The distance and angle to node A is then calculated for each candidate, with the best matching candidate selected as A's communication partner.

We used the following mathematical transfer functions to determine the overall weight of a connection between two nodes. For our proximity metric, we weighed connections between nodes along an inverse logistic growth function or sigmoid. This provides nodes within a threshold social distance (T) with high connection weights, and nodes outside social distance with low connection weights. The social distance threshold is a variable in our engine, and can be adjusted according to circumstance. Hall [9] estimated that social distance is usually less then 10 feet. For field of view calculations we deployed a simple Gaussian window with the current orientation of the node as its centroid. The Gaussian provides high connection weights for nodes that are close to the center of another node's field of view, and low weights towards the extreme angles. The weights for field of view and proximity are summated for each candidate node. The node in the graph with the highest resulting connection weight is subsequently selected as the best matching conversational partner.

### 4.3 Social Groups and Group Geometries

As said, a social group is defined in our system as a cluster of people with mutually shared attention for each other. Social groups can be classified by their structure or group geometry.

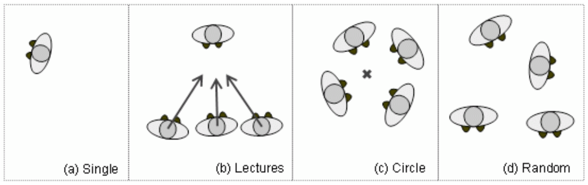
(a) Single (b) Lectures (c) Circle (d) Random

**Figure 5.** Examples for social groups of 4 persons.

Our Social Geometry Engine currently identifies groups of nodes as *Single* when there are no apparent communication partners; *Lecture*, when nodes are paying attention to a single other node; and *Circle* when nodes are facing each other, typical for round table meetings. All other candidates are defined as *Random* (see Figure 5).

In a static grouping approach, groups are modeled as a Boolean graph with people as nodes and attentional connections as directed edges. Social groups are represented by the equivalence class of nodes that are connected by undirected paths in our graph. This is equivalent to the definition for weakly connected components. This static approach works well for small, clearly defined groups. However, a node may interact with multiple other nodes over time. Due to divided attention, this node may only be connected to a single node at any given moment in time. This will cause any static grouping approach to inherently fail.

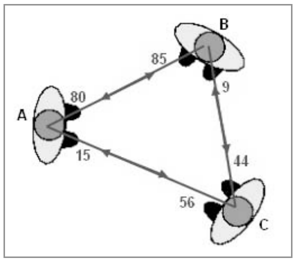

**Figure 6.** Graph with statistically weighted edges.

We therefore explored the use of a probabilistic dynamical model, which assumes that a person is at least periodically oriented towards every person he or she interacts with (see Figure 6) [26]. In this probabilistic graph, weighted edges between nodes in the graph represent the relative amount of time during which the two connecting nodes were selected as each other's communication partner (see Figure 6). Groups are identified after thresholding these probabilistic connection weights.

### 4.4 Modeling Focus Of Attention Toward Objects

Our method, up to this point, only considers social interaction between individuals. However, people may equally well interact with objects, such as white boards, notebooks or computers, which may cause the above algorithm to fail.

Figure **7** shows a situation where three people are looking at a single display, such as a projection screen. In this case, there is a sharing of attention that allows individuals to be classified as members of the same social group. In order to address this concern, our algorithm models the overall focus of attention independently from the group classification method.

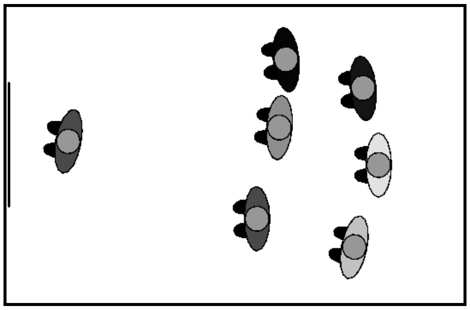

**Figure 7**. Shared focus of attention towards a display.

Joint focus of attention is determined by calculating the weighted field of view for every person. The weights are calculated according to the same functions defined for social interactions. By summating the weighted fields of view of all persons and determining local maxima, the system estimates whether the individuals may be paying attention to the same object. Thus, the persons facing a display can be recognized as one social group with a joint focal point. This allows our system to act robustly in cases where dynamical grouping breaks down, reporting at any time, for any node, the type of social grouping it is in.

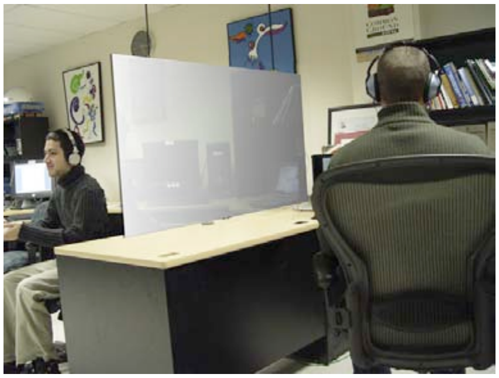

**Figure 8.** Attentive cubicle wall in opaque mode

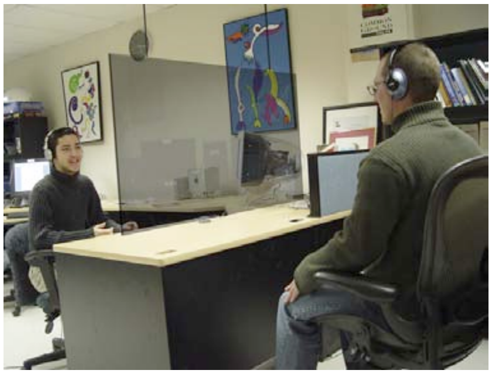

**Figure 9.** Attentive cubicle wall in transparent mode

## 5 Management of Interruption in Co-located Office Scenarios

We are currently working on the application of our social geometry engine in various Attentive User Interface technologies. In particular, video and avatar-based teleconferencing applications may benefit from being able to transmit and reason about social geometry information [27]. Similarly, knowledge about group membership can be used by ubiquitous environments to route communications between users, or sense availability of users for interruption by computing devices or other users.

With the trend towards increasingly flexible office environments, we aimed to explore potential problems of co-worker distraction in the cubicle farm of the future. We applied our social geometry method to the design of an office cubicle that automatically regulates communications between co-workers by sensing whether users are candidate members of the same social group. The goal of this project was to provide technological support for the co-worker's ability to move fluently and seamlessly from focused attention for their tasks to distributed attention for colleagues, following the metaphor and use of the "open office door" in traditional offices [5].

### 5.1 Problem

Problems of managing attention between co-workers can be particularly acute in office cubicle farms, where many users share the same workspace. In our pilot studies, cubicle workers often resorted to wearing noise-canceling headsets in order to avoid distraction by co-workers. Such headsets lower the level of ambient noise, allowing workers to focus better on their tasks. However, the use of noise-canceling headsets also places serious constraints on the effectiveness of collaborations in cubicle farms. This is because noise-canceling headsets reduce co-worker awareness of his or her environment, effectively inhibiting social interactions between office workers. To address this problem, we designed a context-aware office cubicle system that automatically mediates both auditory and visual communications between co-workers on the basis of information about their social-geometrical relationships.

### 5.2 Cubicle Walls as Office Doors

The walls of our cubicle prototype are made of a special translucent material called Privacy Glass™ (see Figure 8) [24]. Privacy glass is essentially a glass pane with an embedded layer of liquid crystals. When powered off, the crystals are aligned randomly, making the glass appear frosted and opaque. When a voltage is applied, the liquid crystals in the glass align, allowing light to go through the pane, thus rendering the glass transparent. When the privacy glass is opaque, cubicle workers cannot be seen by others, and are not distracted by visual stimuli from outside their cubicle. When the privacy glass is transparent, a cubicle worker can interact visually with workers on the other side of his cubicle wall. This allows users of our cubicle to establish an open office door policy regulated through our social geometry engine. To allow management of the auditory attention as well as tracking of individuals by our social geometry engine, cubicle workers wear a noise-canceling headset augmented with retroreflective markers and a microphone. Our system manages the cubicle worker's auditory attention by turning the noise-cancellation circuit in the headset on or off according to the user's measured focus of attention. When noise-cancellation is turned on, all ambient sound is attenuated by –20 Db.

The sound of the microphone can be routed directly to the headset, or to any other co-worker's headset. Depending on preferences, requests for attention of a person inside an opaque cubicle can be made by co-workers approaching the cubicle wall, or by co-workers knocking on its privacy glass pane. Knocks are detected through a contact microphone mounted on the glass. Upon reception of a request for attention, the system notifies the worker inside by temporarily rendering her headset transparent to ambient sound. This allows the cubicle worker to hear normally and respond to a co-worker's request for attention. When a cubicle worker responds by orienting towards the wall, the engine detects co-orientation and automatically renders the wall as well as the headphones transparent (see Figure 9).

### 5.3 Scenario of Use

The following scenario illustrates the use of our attentive office cubicle. User Alex is busy finishing a report. Alex has a tight deadline, and needs to file the report by the end of the day. While Alex is trying to focus on his writing, his colleague Jeff is discussing a design strategy with Laurie, a co-worker, in the next cubicle. All three individuals are wearing a noise-canceling headset that is tracked by our engine. The social geometry engine recognizes Laurie and Jeff are co-located and oriented towards each other, without any physical barriers between them. It reports each as a potential communication partner to the other person's headset, causing the system to pipe the audio from Jeff's microphone to Laurie's headset, and vice versa. The social geometry engine detects that co-worker Alex is not a part of any group, as he is oriented towards the wall of his cubicle. This causes Alex's headphones to engage noise cancellation and render his cubicle's privacy glass opaque. When Jeff and Laurie require Alex's assistance, Jeff makes a request for Alex's attention by knocking on the cubicle's privacy glass. The request is forwarded to Alex's headset, and informs the geometry engine to consider the wall between the two individuals removed. It causes Alex's noise cancellation to be turned off temporarily, allowing him to hear the request. As Alex responds to the request, he orients himself to the source of the sound. The social geometry engine detects the co-orientation of Jeff and Alex and sends a signal over an X10 interface to the privacy glass between them. This causes the glass to be rendered translucent, allowing Jeff and Alex to interact normally. After the conversation is completed, Jeff moves away from the cubicle wall, continuing his discussion with Laurie. Alex turns his attention back towards his computer system, causing the geometry engine to conclude Alex and Jeff are no longer candidate members of the same social group. The system responds by turning on noise-cancellation in Alex's headset, and by rendering the privacy glass of his cubicle opaque again. The above scenario illustrates how office cubicles can be designed to balance social as well as privacy needs of co-workers in a dynamical fashion. The above scenario can be applied to any location in the office by allowing sound from the headset microphones to be distributed to all people within an apparent social group. This allows co-workers that are part of an apparent social group to converse with each other without interfering with conversations of other groups.

### 5.4 Initial Experiences

Initial experiences with our systems are encouraging. We verified the accuracy of our group classification algorithm by running the system with four participants in our laboratory. Participants were asked to perform conversations in a number of group settings. Initial results suggest our engine correctly identified all types of social groups depicted in Figure 11, as well as transitions between them. Initial studies also suggest that the active management of co-worker attention provides a promising approach, with 3 out of four respondents suggesting our attentive cubicle reduced distraction levels. In the near future, we plan to evaluate our system more rigorously with larger numbers of participants and in real cubicle farms.

## 6 Conclusions

In this paper, we presented a method for determining the social engagement of participants of group meetings, through reasoning about the Social Geometries formed by their bodies during meetings. A vision-based system tracks user location and orientation through head-mounted retroreflective markers. Our social geometry engine infers group membership by examining co-alignment of individuals. We applied the engine in the design of an office cubicle that automatically mediates communications between co-workers by sensing whether they are candidate members of the same social group. The cubicle regulates visual interactions through the use of privacy glass, which can be rendered opaque or transparent upon detection of joint orientation. Auditory interactions are regulated through noise-canceling headphones that automatically become transparent to ambient sound upon co-orientation.

## 7 References

1. Abowd, G.D. and Mynatt, E.D.: Charting Past, Present and Future Research in Ubiquitous Computing. ACM TOCHI 7(1), pp. 29-58, March 2000.
2. Argyle, M. and Cook, M.: *Gaze and Mutual Gaze*. Cambridge University Press, 1976.
3. Barkhuus, L., Dey, A.K.: Location-based Services for Mobile Telephony: a Study of Users' Privacy Concerns. INTERACT 2003. September 1-5, 2003.
4. Basu S. and Pentland A.: Smart Headphones. In Extended Abstracts of CHI'01. Seattle: ACM Press, 2001.
5. Buxton, W.: Ubiquitous Video, Nikkei Electronics, 3.27 (no. 632), pp. 187-195, 1995.
6. Crabtree, A., Rodden, T., Hemmings, T., and Benford, S.: Finding a Place for UbiComp in the Home. In Proceedings of Ubicomp'03, pp. 256-263, 2003.
7. Frolich, D., et al. Informal Workplace Communication: What is It Like and How Might We Support It? HP Tech. Report, 1994.
8. Gibbs, W. Considerate Computing. *Scientific American*, January 2005.
9. Hall, E.: The Silent Language. Garden City, N.Y.:Doubleday, 1959.
10. Horvitz, E., Jacobs, A., and Hovel, D.: Attention-sensitive Alerting. In Proceedings of UAI'99. Stockholm: Morgan Kaufmann, pp. 305-313, 1999.
11. Horvitz, E., Kadie, C. M., Paek, T., and Hovel, D.: Models of Attention in Computing and Communications: From Principles to Applications, Communications of the ACM 46(3), pp. 52-59, March 2003.
12. Hudson, S., et al. Predicting Human Interruptibility with Sensors: a Wizard of Oz Feasibility Study. In Proceedings of CHI'03, pp. 257-264, 2003.
13. Mamuji, A. et al. AuraLamp: Contextual Speech Recognition In An Eye Contact Sensing Light Appliance. In Extended Abstracts of Ubicomp'03, 2003.
14. Mamuji A., Vertegaal R., Sohn C., Cheng, D. Attentive Headphones: Augmenting Conversational Attention with a Real World TiVo. In Extended Abstracts of CHI'05, Portland, OR, 2005.
15. Mantei , M., Baecker , R., Sellen, A., Buxton, W., Milligan, T., Wellman, B.: Experiences in the use of a media space, In Proceedings of CHI'91, p.203-208, 1991.
16. Matlock, T., et al.: Designing Feedback for an Attentive Office. In Proceedings of INTERACT'01, 2001.
17. Meier, R. L.: A Communications Theory of Urban Growth. Boston: MIT Press, 1962.
18. Mueller F. and Karau M.: Transparent Hearing. In Extended Abstracts of CHI'02, ACM Press, 2002.
19. Sellen, A. J.: Speech patterns in video-mediated conversations, in Proceedings of CHI'92, pp.49-59, 1992
20. Shell, J., Selker, T., and Vertegaal, R.: Interacting with Groups of Computers. In Special Issue on Attentive User Interfaces, Communications of ACM 46(3), March 2003.
21. Siewiorek, D., et al. SenSay: A Context-Aware Mobile Phone. In Proceedings of ISWC'03, pp. 248-249, 2003
22. Stiefelhagen, R.: Tracking Focus of Attention in Meetings. In Proc. of ICMI'02, 2002.
23. Stiefelhagen, R., Zhu, J.: Head Orientation and Gaze Direction in Meetings. In Proceedings of CHI'02, 2002.
24. SwitchLight. Privacy Glass™ Specification. www.switchlite.com, 2003.
25. Vertegaal, R., et al.: Designing Attentive Cell Phones Using Wearable EyeContact Sensors. In Extended Abstracts of ACM CHI'02, 2002.
26. Vertegaal, R., et al.: Eyegaze Patterns in Conversations: There is More to Conversational Agents than Meets the Eyes. In Proceedings of CHI'01, 2001.
27. Vertegaal, R., et al. GAZE-2: Conveying Eye Contact in Group Video Conferencing Using Eye-Controlled Camera Direction. In Proceedings of CHI'03, 2003.
28. Vertegaal, R., Van der Veer, G.C. and Vons, H.: Effects of Gaze on Multiparty Mediated Communication. In Proceedings of GI 2000. Montreal, pp. 95-102, 2000.
29. Weiser, M.: The Computer for the 21st Century. Scientific American, pp. 94-110, 1991.