

Improving the match between callers and receivers: A study on the effect of contextual information on cell phone interruptions

D. AVRAHAM^{*†}, D. GERGLE[‡], S. E. HUDSON[§] and S. KIESLER[¶]

Human-Computer Interaction Institute, Carnegie Mellon University, Pittsburgh, PA, USA

A problem with the location-free nature of cell phones is that callers have difficulty predicting receivers' states, leading to inappropriate calls. One promising solution involves helping callers decide when to interrupt by providing them contextual information about receivers. We tested the effectiveness of different kinds of contextual information by measuring the degree of agreement between receivers' desires and callers' decisions. In a simulation, five groups of participants played the role of 'Callers', choosing between making calls or leaving messages, and a sixth group played the role of 'Receivers', choosing between receiving calls or receiving messages. Callers were provided different contextual information about Receivers' locations, their cell phones' ringer state, the presence of others, or no information at all. Callers provided with contextual information made significantly more accurate decisions than those without it. Our results suggest that different contextual information generates different kinds of improvements: more appropriate interruptions or better avoidance of inappropriate interruptions. We discuss the results and implications for practice in the light of other important considerations, such as privacy and technological simplicity.

Keywords: Context-aware computing; Awareness systems; Mobile communication; Interruptions

1. Introduction

The world has witnessed a tremendous growth in people's use of mobile technology. This development, combined with the rise of other communication techniques such as email and instant messaging, has been described as a shift from a 'place-to-place' communication to a 'person-to-person' communication (Wellman 2001). Despite their many benefits, cell phones create problems that arise from a mismatch between their use and the situations in which they are being used. The aspect of this mismatch that we address here involves callers making phone calls at inconvenient, disruptive, or even dangerous moments for the receivers. Inappropriate calls include those made when the call is socially inappropriate, for example when the receiver is in a meeting or at the theatre, or when the receiver is engaged in an activity requiring attention, such as driving (Green 2000). Even when such interruptions are not dangerous, they can increase receivers' stress and errors

(see Eyrolle and Cellier 2000) or cause annoyance to bystanders (Wei and Leung 1999, Monk *et al.* 2004a, Monk *et al.* 2004b).

We report a study examining a possible solution to this problem. The solution involves presenting information to callers about receivers' context, to help them decide whether or not to call.

1.1 The call process

To understand the roots of the problem of inappropriate incoming calls, let us consider the way in which callers decide whether or not to make a telephone call.

First, we assume the call decision typically follows a process called 'naturalistic decision-making' described by Zsombok and Klein (1996). That is, rather than carefully weigh rational costs and benefits of a call, a prospective caller makes a quick judgement about the identity of the intended receiver and their relationship with them,

*Corresponding author. Email: {†nxd, ‡dgergle, §scott.hudson, ¶kiesler}@cs.cmu.edu

the urgency of the call, and the likely state of the receiver – their identity, location and activities, based in part on the phone number, day and time (Lacohee and Anderson 2001).

Cell phones have introduced a change to the information available to callers. With landline phones, the caller may not be able to predict if the intended person will be the one to answer as opposed to, for example, another member of the household. However, once the intended is reached, their location is known. By contrast, calling a cell phone typically assures that the receiver will be the person intended. However, cell phones are usually location-free, and do not provide callers with knowledge about receivers' locations. The absence of location information can make it difficult for a caller to predict the possible activities the receiver might be engaged in and, as a result, whether a call at that time is appropriate.

Providing callers with information about receivers might alleviate the problem (see for instance Pedersen 2001, Tang *et al.* 2001), but the effectiveness of such a solution has not yet been established. Many questions have also been left open: What kinds of information would provide the best results? What information would be simplest to provide? What information would entail the least privacy intrusion?

To answer these questions, we need to consider the way receivers choose whether or not to accept a call. Receivers base their decisions on a number of factors (Hudson *et al.* 2002). These include the identity of the caller, their relationship with the caller, their current activities, and the social situation (such as the presence of others, the need for privacy and the possibility that the call will intrude on others). Receivers might base their decision also on the subject of the conversation, but this information is typically unknown to them in advance. Each of these factors may be considered, but one can imagine a few cases where one factor overpowers the rest. For instance, the receiver may ask the caller to let him/her know immediately about a change in a schedule; however, the receiver would not accept the call while s/he is in a meeting.

1.2 Review of existing solutions

1.2.1 Single rule solutions. The most common type of solution allows receivers to apply a single rule, typically by changing a single setting on the cell phone. This rule is then applied to all calls without exception. For example, the receiver can turn the cell phone off, thus refusing all calls, or decide to leave the phone on, allowing all calls to go through. If the phone is left on, the volume of the phone ringer can be set to a specific level, or the ringer can be turned off having the phone vibrate and/or flash when an incoming call arrives.

The main drawbacks to solutions of this type are that they rely on memory and appropriate planning by receivers

(for example, remembering to turn the phone off at the beginning of a meeting and then remembering to turn it back on again). Furthermore, by applying a single rule for all incoming calls, the rule ignores other important factors, such as the identity of the caller and urgency of the call.

1.2.2 Manual filtering solutions. These solutions allow the receiver to make a decision at the time of the call, based on the identity of the caller. This is implemented by the receiver through the use of Caller ID or through associating specific ringer tones with specific people or groups of people (using caller ID as its underlying mechanism). One proposed solution also provides the receiver with a subtle mechanism for refusing calls and providing the caller with some information (Pering 2002). Another allows the receiver to accept the call and converse without talking (Nelson *et al.* 2001). The main benefit of these solutions is that the receivers can apply temporary rules or exceptions to rules based on their situation.

These solutions, however, have two main drawbacks. The first is that they ignore the importance of a specific call, treating urgent and nonurgent calls equally. The second – and more important drawback – is that the receiver has to make a decision for every incoming call. Since even a brief interruption may be disruptive (Gillie and Broadbent 1989), demanding the receiver's attention for the purpose of filtering can be costly, ultimately transforming the filtering process into an interruption in itself.

1.2.3 Multiple rules solutions. Some solutions involve the receiver setting up 'profiles' in advance for different activities, locations and people. The profiles created using these solutions could be sophisticated, but would require that receivers anticipate and categorise situations in advance. These solutions ignore the dynamic nature of many contexts. Furthermore, they require receivers who use them to plan appropriately, and remember to reset the correct profile (for a discussion of similar issues see Grudin 1988).

1.2.4 Automation solutions. Automating the decisions made by the receivers seems to be a solution with great potential. These include solutions that gather information about the receiver and infer an appropriate behaviour from it. Using sensing and other data sources (such as location and schedule), some of these solutions try to model the activity of the receiver and infer their activity level (Horvitz *et al.* 1999). The actions these solutions offer include turning the phone on or off, setting the phone to a particular profile (Schmidt *et al.* 1999) and redirecting to a different medium (Schmandt *et al.* 2000). There is an ongoing conceptual debate as to whether this type of solution can work effectively (see, for example, Bellotti and Edwards 2001).

1.2.5 Caller-based solutions. All the solutions cited above focus on the receiver's decisions and actions, ignoring the caller as a potential decision-maker. We propose instead to include the caller in the decision process and ask whether providing callers with contextual information about the receiver could improve both parties' satisfaction with calls. Some of these solutions provide callers with continuous awareness information about the receivers in a manner similar to the awareness provided by Instant Messaging systems (Isaacs *et al.* 2002). These solutions require that such information be sensed, inferred, or manually set by the receiver. For example, see (Milewski and Smith 2000, Schmidt *et al.* 2000, Bellotti and Edwards 2001, Pedersen 2001, Tang *et al.* 2001).

This type of solution assumes that communication is a cooperative act. If that assumption is valid, the main benefit of this type of solution is that it redistributes the call decision, removing some of the cognitive and social burden from receivers and placing it in the hands of callers. For the caller, a promising aspect of this type of solution is that it could leverage human judgement in determining whether the subject of conversation and current social environment of the receiver yields an appropriate time for delivering the message. However, it is not obvious that this type of solution would work in practice. On the one hand, the information provided to callers may be insufficient. On the other hand, more detailed information, even if available, could compromise the privacy of receivers. Solutions that require receivers to monitor and set the information provided to callers also have the same drawbacks as those of multiple rule solutions (see section 1.2.3 above).

1.3 Requirements of caller-based solutions

From this brief review, it seems that a good solution should at least have the following properties: take into account the factors that receivers use when deciding whether to accept a call, take into account the changing nature of these factors, and require minimal work and attention from receivers. The solution of providing callers with sensed contextual information about receivers seems to fit these requirements.

2. Experiment 1

In the first experiment we attempted to assess the likely value of a caller-based solution. We examined whether providing callers with information about receivers' context could improve the match between callers' behaviour and receivers' desires, and whether more detailed contextual information generated improved caller behaviour. We also investigated the role of inherent differences in caller and receiver motivations (O'Conaill and Frohlich 1995, Palen *et al.* 2000).

3. Method

3.1 Design

3.1.1 Participants as Receivers or Callers. One group of participants played the role of Receiver and four groups of participants played the role of Caller. Receivers were asked to indicate whether a particular message in a given situation should be delivered as a phone call (immediate interruption), or by voice mail (delayed interruption). Each of three groups of Callers received some information about a Receiver, while a fourth group, the control group, received no information. For every message, each Caller had to choose between making the call and leaving a voice-mail message. All participants were then asked to rate the urgency of each message.

In addition to the specific content of the message given to each group we varied the situational and environmental context. For example, a participant playing the role of Receiver might be presented with the following statement: 'You are in your office reading your email. A colleague is about to call you to tell you that John will be late for your meeting.' The Receiver was asked to decide whether the colleague should call or leave a voice mail.

The relationship between the caller and receiver was described as colleagues. That allowed us to investigate messages of both personal and work-related nature. The choices made by Receivers were used to establish baseline ratings for the particular messages in given contexts. These ratings were later used to evaluate how accurate the Callers were in assessing whether or not to call with a given message.

3.1.2 Messages. The messages presented to the participants were pruned from an initial set of 24. The initial set was pretested and a subset of messages was kept. The rated urgency of these messages tended to be of low variance across the raters and did not cause confusion. In an attempt to create a balanced sample of call types (in order to avoid introducing bias to the responses), this subset was further pruned to a final set of 16 messages. These were balanced for work-related and personal calls, and good, neutral or bad news. The messages used in the experiment can be found in figure 1.

3.1.3 Contextual awareness conditions. To assess the value of particular types of contextual information, we varied the amount and type of contextual information that we gave Callers along with the messages. In this experiment, we focused on providing the following four levels of contextual awareness:

1. None: In this condition, participants had no information regarding the current situation of the receiver. A participant in this condition would see

- **Your favourite football team has lost**
- **You have a movie overdue**
- **Your paper was accepted by the journal**
- **Your paper was rejected by the journal**
- **Happy Birthday!**
- **Your project got funding**
- **You have a package in the mail room**
- **Your son is ill and is being sent home**
- **Your parents' flight will be late**
- **Your next meeting is cancelled**
- **John will be late for your meeting**
- **We are at the coffee shop. You are welcome to join us**
- **Your car is ready for pickup**
- **Would you like us to buy you a ticket for the concert?**
- **Your boss wants to talk to you**
- **Your assistant wants to schedule a meeting**

Figure 1. Messages presented to participants.

a statement such as, 'You are planning to contact a colleague on his cell phone to wish him a happy birthday.'

2. Presence of People (People): This condition provided information about whether or not there were people located nearby the receiver. The levels of this condition included: Alone and Not Alone. We did not explicitly state whether the receiver was socially engaged with the other people, but rather whether there were other people around. An example statement would be, 'You are planning to contact a colleague on her cell phone to tell her that she has a package in the mailroom. There are no people around your colleague.'
3. Location: Here we provided the caller with information about the location of the receiver. The levels of this condition included: Home and Office. Here is an example of a statement in this condition: 'You are planning to contact a colleague on his cell phone to tell him that his project got funding. Your colleague is at home.'
4. Combined Information (Combo): This condition provided the caller with a combination of the information provided in the People and Location conditions (Alone vs. Not Alone and Home vs. Office). An example statement might be, 'You are planning to contact a colleague on his cell phone to tell him that his son is ill and is being sent home. Your colleague is at

the office. There is at least one person around your colleague.'

3.1.4 Experimental design tradeoffs. In our attempt to explore the mismatch between the Caller's willingness to transmit a call and the Receiver's desire to receive it, we had to choose a careful subset of environments and situations to present. For example, we chose not to vary the time of day or relationship of Caller and Receiver (we assume that additional shared information between caller and receiver would likely only improve the match between receivers and callers). We chose the aforementioned conditions because they represented achievable solutions using existing technologies and were previously suggested in the literature as viable solutions. We made the tradeoffs that allowed us to gain experimental control and allowed a detailed examination of how a caller-based solution might work.

3.2 Apparatus and materials

We collected data in an online environment. Java Server Pages (JSP) were automatically generated for every participant. The participants received a web address to go to, where they received instructions that varied depending on whether they played the part of Caller or Receiver.

3.3 Participants

Participants were 78 people (26 Receivers and 52 Callers) drawn from a relatively diverse sample that included both students and professionals. The average age of participants was 30.1 (SD = 10.7, Min = 18, Max = 68). Participants were randomly assigned to role and contextual awareness conditions.

3.4 Procedure

Upon arriving at the website, the participants were given an online consent form. Afterwards, a brief survey collected demographic information. The participants were then given instructions that pertained to their condition (Receiver or Caller). The task took approximately 10 minutes.

3.4.1 Receivers. Participants who played the role of Receiver were asked to choose between receiving a call or a voice mail for all of the messages in each of four combinations of location and presence of people. The order of the four situations was counterbalanced across Receivers, and the order of the messages was randomised for each situation (with the viewing order of message recorded to serve as a control and test for order effects). A within-subjects design was used for the ratings of the messages. When they were finished, Receivers were presented with each of the 16 messages in a random order and asked to ascribe a level of urgency to each message on a 5-point scale (1 = low urgency, 5 = high urgency).

3.4.2 Callers. Participants playing the role of Caller were assigned to one of the four different contextual awareness information conditions randomly (None, Location, People, Combo). They were asked to choose between making a call and leaving a voice mail once for each of the 16 messages – one message at a time, in a random order (again, with the viewing order of message recorded for control). Each message was presented with one of the levels of the appropriate contextual information (e.g. in eight of the messages presented to a Caller in the Location condition, the receiver was at home, and in eight the receiver was in the office). The combination of message and level of contextual information was counterbalanced across Callers in the same condition. A between-subjects design was used for the rating of the messages. As with the Receivers, the Callers were then presented with the messages in a random order and asked to rate their urgency on the same scale.

4. Results

We present the results in three stages. The first explores the Receivers' willingness to be interrupted in different

contextual conditions. Next, we compare the overall differences between the Callers and Receivers in their frequency of calls and their perceived levels of urgency. Finally, we examine the degree of agreement between the Callers and Receivers and explore the ways in which these differences interacted with the range of contextual information available. Four participants in the Caller role were excluded from the analysis as they chose to call 100 per cent of the time, regardless of the message or situation. We suspect that these participants represent a segment of the population for which this solution may be ineffective (see the discussion section for further comments).

4.1 Receivers' preferences

To examine Receivers' preferences for calls or voice mails, we used a repeated measures analysis of variance in which Location (Home vs. Office), People (Alone vs. Not Alone), the interaction term (Location x People), Urgency (1–5), Message (1–16), and Message Viewing Order (1–16) were repeated. As each participant rated several questions, observations were not independent of one another. Participants were modelled as a random effect.

We found that Receivers were much more willing to receive calls when they were at home (M = 52.8%) than when they were at the office¹ (M = 44.6%) ($t = -4.53$, $p < .001$). Similarly, we found that Receivers were much more likely to take calls when they were alone (M = 61.5%) than when others were present (M = 35.9%) ($t = -13.92$, $p < .001$). A more detailed exploration of the interaction between the Receivers' location and the presence of people demonstrates that the presence of people affected Receivers' choices more when they were in the office (making them substantially less willing to accept calls) than when they were at home (for the interaction, $F(1,1601) = 57$, $p < .001$; see figure 2). Put another way, Receivers did not want to receive calls when in the office, and in particular when they were in the office and not alone.

Not surprisingly, the relationship between Receivers' ratings of message urgency and their desire to accept that message as a call, instead of a voice mail, was strong ($F[4,1601] = 54.5$, $p < .001$; see figure 3).

4.2 Differences between Callers and Receivers

We examined preferences for calls or voice mails and urgency in all conditions. These analyses were used to create a baseline score for accuracy, that is, to examine how similar were choices made by Callers and Receivers overall.

Prior literature suggests there may be a discrepancy in the preference for a call between receivers and callers (O'Connell and Frohlich 1995). To test this discrepancy, we began by looking at the overall proportion of calls the Receivers wished to accept in comparison to the Callers'

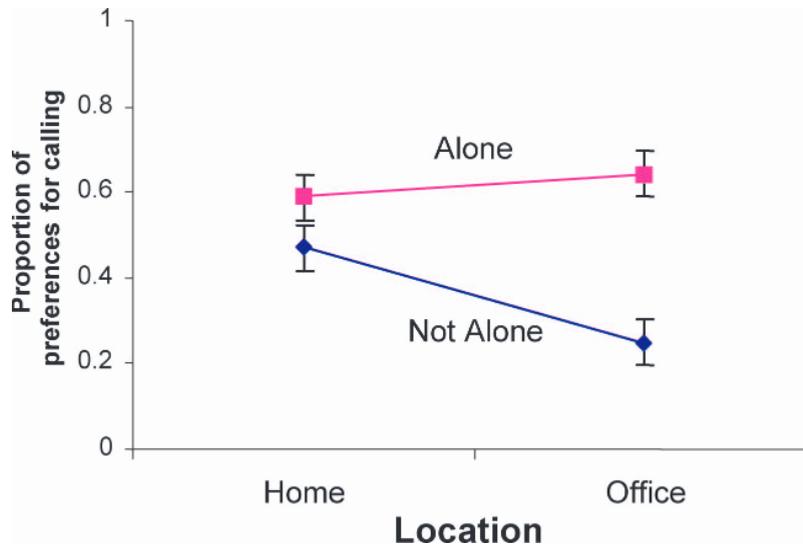


Figure 2. Proportion of preferences for calling, by Location and People (Receivers only data).

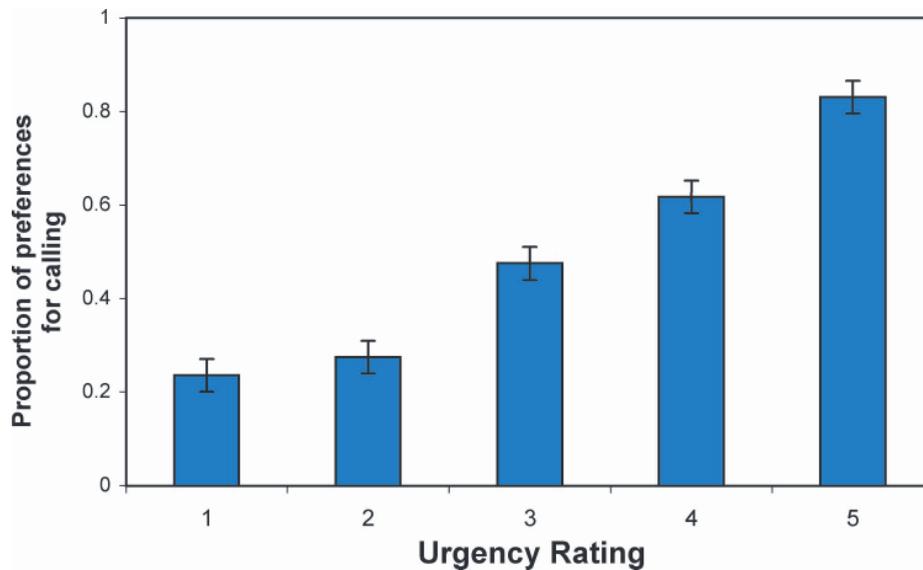


Figure 3. Proportion of preferences for calling, by rating of Urgency (Receivers only data).

intentions (see figure 4). For this analysis we used the complete data set and modelled it in a similar fashion to the previously described model with the inclusion of Role (Caller vs. Receiver) as a factor in the model. We found that Role did indeed have a significant effect on the proportion of calls for which participants chose a call over a voice mail ($F [1,2340]=9.26, p=.002$). Receivers preferred a call approximately 50 per cent of the time whereas the Callers chose to call approximately 60 per cent of the time.

One alternative explanation for this result might be that the participants who received no contextual information were the source of the difference. However, the result remained significant when the Callers in the None condition were removed ($F [1,2250]=8.68, p=.003$). This demonstrates that the overall difference between Callers' and Receivers' preferences for a call was not caused solely by Callers having no contextual information.

Another alternative explanation of why Receivers' desires and Callers' behaviour differed might be that they

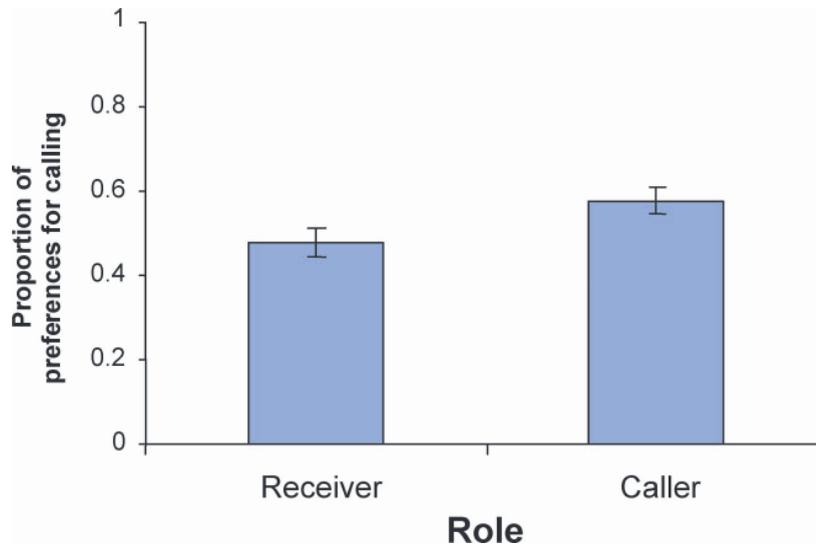


Figure 4. Proportion of preferences for calling, by Role (Receivers vs. Callers).

perceived the urgency of the messages differently. To explore this possibility, we first investigated the relationship between Callers' choice to place a call instead of a voice mail and their ratings of message urgency. Similar to Receivers, this relationship was strong ($F [1,697]=38.49$, $p < .001$). We then used the same model as described above to compare Callers and Receivers, using urgency as the dependent variable. Overall levels of urgency did not vary across Role ($M = 3.14$ vs. $M = 3.11$, *n.s.*), nor did any of the 2- or 3-way interactions.

4.3 Effects of contextual information

In order to determine the match between Receivers' desires and Callers' judgements, given different contextual information, we calculated a difference score between the average proportion of call versus voice mail choices given by the Receivers and each group of Callers. These differences tell us about the degree to which the Callers accurately anticipated Receivers' preferences for a call in the various contextual information conditions. The difference scores were calculated by subtracting from every Caller's choice (scored 1=Call, 0=Voice Mail) the corresponding average choice given by Receivers for the same message in a matching situation (except for the None condition where the average by message was used). For example, the choice of a Caller in the People condition for message two, when the information provided was that the receiver was alone, was compared to the average choice of Receivers for the same message when they were alone. The difference score for every contextual information condition was then computed as the absolute value of the average score by message. It was important to use absolute values

since both negative and positive deviation (over-calling and under-calling) should be taken into account and not allowed to cancel one another out. This procedure yielded scores ranging from 0 to 0.9, which were approximately normally distributed.

4.3.1 Differences between conditions. The scores were first subjected to a one-way ANOVA to investigate the differences between the conditions. The following results present the data at the summary level (i.e. one value for each participant). The overall model was significant ($F [3,700]=5.50$, $p = .001$; see figure 5). We used planned contrasts to examine where these differences occurred. We found that the Callers in the None condition were significantly worse (higher difference score) than in any of the other conditions ($p < .05$). However, we did not find significant differences in overall accuracy across the different contextual information conditions.

4.3.2 Differences within conditions. Next, we examined the effect of different levels of contextual information on the accuracy of Callers' judgements in the various conditions (Location, People and Combo). We subjected the scores of each group to a repeated-measures ANOVA in which the difference-score (between Caller and Receiver) was the dependent variable and Urgency (1–5), Message (1–16), Message Viewing Order (1–16) and level of contextual information were repeated. Participants were modelled as a random effect.

Callers in the Location condition made significantly more accurate judgements (lower difference score) when the receiver was in the office than at home ($M = 16.1\%$ vs.

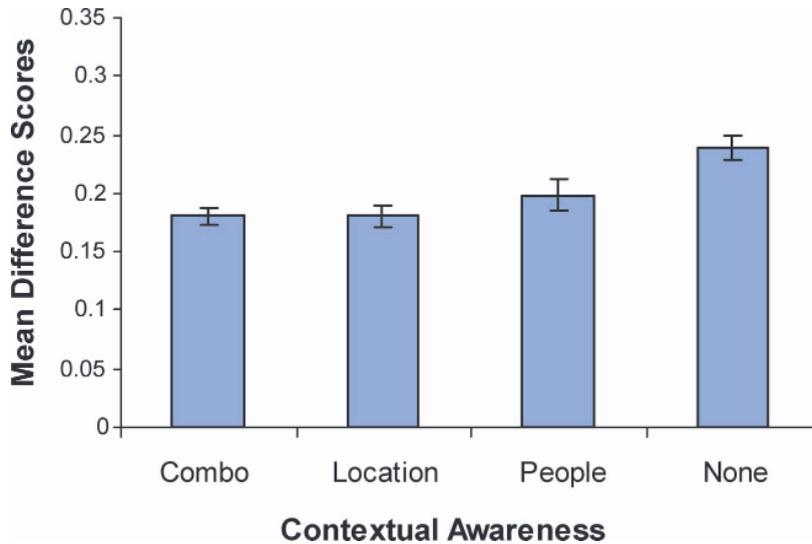


Figure 5. Mean absolute differences between the Receivers and Callers by contextual awareness condition (lower = better).

$M = 20.2\%$; $F [1,162] = 9.52$, $p = .002$; see figure 6a). One simple explanation could be that the office (or work in general for that matter) is typically perceived as a less appropriate context for interruptions. The availability of someone at home, on the other hand, does not have this characteristic. In other words, Callers in the Location condition performed significantly better when they received a less ambiguous indication. This indication (receiver in the office) helped them to *avoid interrupting* the Receiver inappropriately (from the Receiver's perspective).

Callers in the People condition made significantly more accurate judgements when the receiver was alone than when other people were around ($M = 15.7\%$ vs. $M = 23.7\%$; $F [1,132] = 20.65$, $p < .001$; see figure 6b). A similar explanation would be that being alone is a more acceptable context for interruptions. Higher ambiguity exists when someone is not alone and there is no indication of social engagement. In other words, Callers in the People condition performed significantly better when they received a less ambiguous indication, as in the Location condition. In this case, this indication (receiver alone) helped them to *interrupt* at an opportune time (from the Receiver's perspective).

As in the People condition, Callers in the Combo condition were significantly more accurate when the receiver was alone ($F [1,280] = 75.2$, $p < .001$), but, unlike the Location condition, they also were significantly more accurate when the receiver was at home ($F [1,280] = 21.83$, $p < .001$). An exploration of the interaction between the location and the presence of people showed that the Receiver's being alone caused improved accuracy (lower differences) with little effect of location. Accuracy when the Caller learned the Receiver was at the office and not

alone was drastically inferior (28.9%) (for the interaction, $F [1,280] = 8.98$, $p = .003$; see figure 6c). One explanation could be that Callers in the Combo condition did not reduce their call rate enough to match the significantly lower call rate desired by Receivers when they were not alone in their office (see figure 2).

These findings are interesting in that they suggest that providing different kinds of contextual information can result in different types of improvement in Callers' accuracy.

5. Experiment 2

In this experiment, a group of participants played the role of Caller with a new type of contextual information – a receiver's cell-phone ringer status. We felt that using the cell-phone ringer status as an indicator for contextual information had potential. Different ringer settings are a feature that already exists in cell phones, is familiar to cell phone users, and providing knowledge of it to callers is technologically simple. As in experiment 1, participants received some information about a receiver and for every message in a set, had to choose between making a call and leaving a voice-mail message (for the list of messages, see figure 1). Participants were then asked to rate the urgency of each message. Participants' responses were compared and contrasted with responses of Receivers and Callers in the None condition from experiment 1.

6. Method

The apparatus, materials and procedure of this experiment were identical to those of experiment 1.

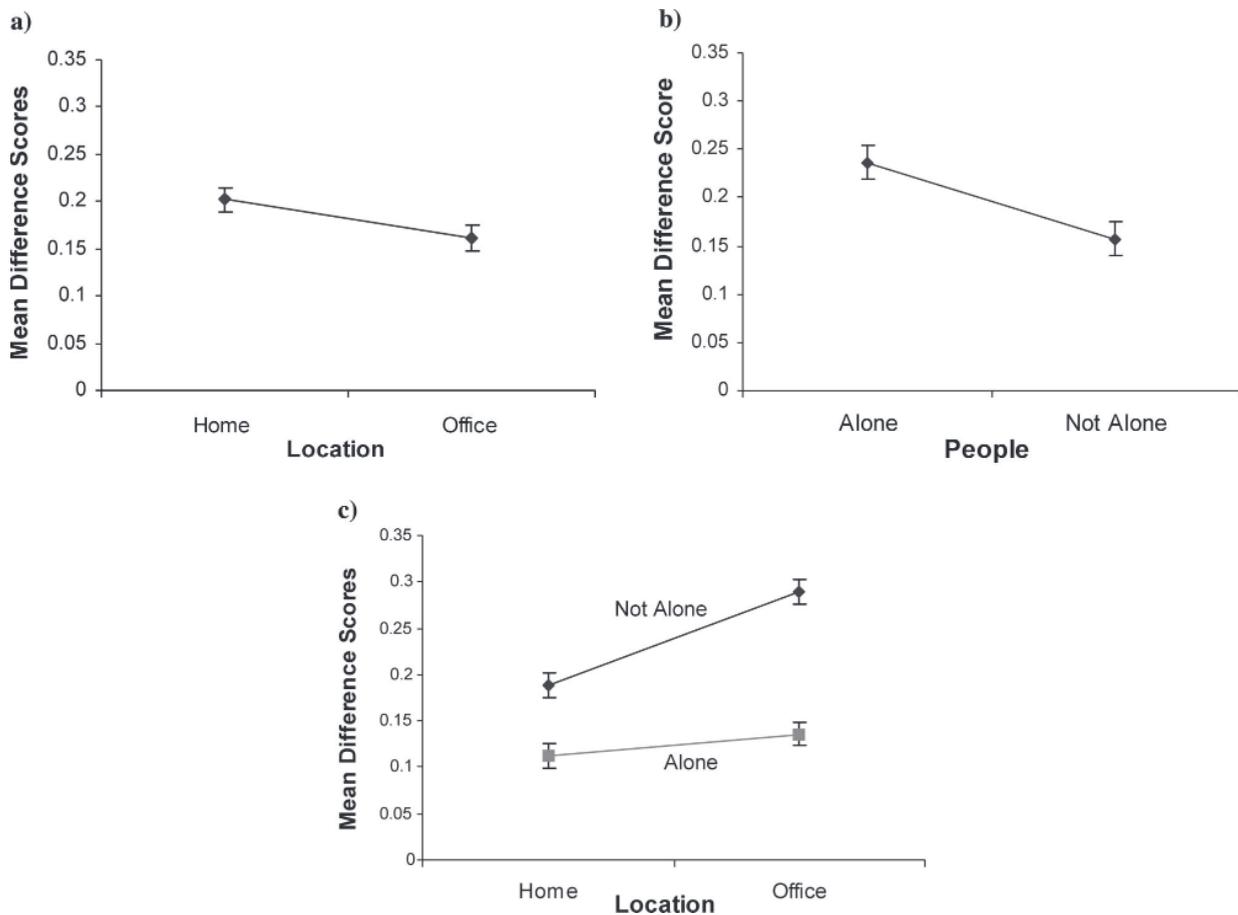


Figure 6. Mean absolute differences (a) between home and office in the Location condition (lower = better); (b) between alone and not alone in the People condition (lower = better); (c) by location and presence of people in the Combo condition (lower = better).

6.1 Design

6.1.1 Contextual awareness information. Cell-Phone Ringer Status (Ringer): Participants were provided information about the setting on the Receiver's cell phone. The levels of the condition were Loud and Silent (flash and vibrate). This information was presented in the following way, 'You are planning to contact a colleague on her cell phone to tell her that you are at the coffee shop and she is welcome to join you. The ringer on your colleague's cell phone is turned off (the phone will flash and vibrate).'

6.2 Participants

Participants were 12 people drawn from the same pool at the same time as experiment 1, which included a relatively diverse group of students and professionals. The average age was 27.7 (SD = 6.08, Min = 19, Max = 43).

7. Results

We present the results of this experiment in two stages. The first compares the overall differences between the Callers of this experiment and the Receivers (the same ones employed in experiment 1) in their frequency of calls and the perceived levels of urgency. Next, we examine the degree of agreement between the same Callers and Receivers and compare the degree of agreement in this experiment with that of Callers in the None condition from experiment 1.

7.1 Differences between Callers and Receivers

We examined the call versus voice-mail choices and urgency ratings made by Callers and compared them to the choices and ratings made by Receivers. We began by looking at the overall number of calls the Receivers wished to accept in comparison to the Callers' choices. For this analysis we

used the data for Callers gathered in this experiment and the complete data of the Receivers collected in experiment 1. We used a repeated-measures ANOVA in which Urgency (1–5), Message (1–16), Message Viewing Order (1–16), and Role (Caller vs. Receiver) were repeated. Participants were modelled as a random effect.

We found that while the difference was in the same direction as in experiment 1, role did not have a significant effect on the frequency with which participants in this Ringer condition chose a call over a voice mail ($M = 50\%$ vs. $M = 54\%$, n.s.). Urgency played a significant part in Callers' decision to place a call ($F [1,162] = 48.07$, $p < .001$). As in experiment 1, urgency did not significantly vary across role.

7.2 Effects of contextual information

To test the effect of the contextual information we compared the Callers in the Ringer condition to Callers in the None condition from experiment 1, using absolute difference scores. The difference scores were calculated by subtracting from every Caller's choice (1 = Call, 0 = Voice Mail), the corresponding average choice given by Receivers for the same message. Since the cell-phone ringer status was never presented to Receivers, the overall average by message was used (similar to the difference scores calculation for the None condition). As before, absolute values were used, to prevent negative and positive deviation (over-calling and under-calling) from cancelling one another out. The scores were subjected to a one-way ANOVA to investigate the differences between the conditions using difference score as the dependent variable.

We found that the Callers in the Ringer condition were significantly better (lower difference score) at choosing appropriate interruptions than Callers in the None condition ($F [1,252] = 107.48$, $p < .001$; see figure 7). This suggests that merely drawing attention to context may reduce unwanted interruptions. There was no significant performance difference between the two levels of contextual information (Loud vs. Silent).

8. Discussion

Cellular phones are an important part of the communication repertoire available to people today. Cell-phone users, however, are subject to the problem of incoming calls at inappropriate times.

A primary goal of this study was to evaluate the potential of providing callers with contextual information from sensors as a solution to this problem. We wanted to see whether callers who are provided with contextual information about receivers would be less likely to call at inappropriate times. The relatively poor performance of Callers in the None condition confirms the existence of this problem. The significant improvement shown by the Callers who received contextual information suggests that providing contextual information (location, presence of people and the cell-phone ringer status) may indeed reduce the number of inappropriate interruptions. Both experiments also confirmed our belief that Callers could adjust their behaviour to try and match the desires of the Receivers by combining their knowledge of the conversation subject with the contextual information provided. This result fits well with other studies discussing the value of providing

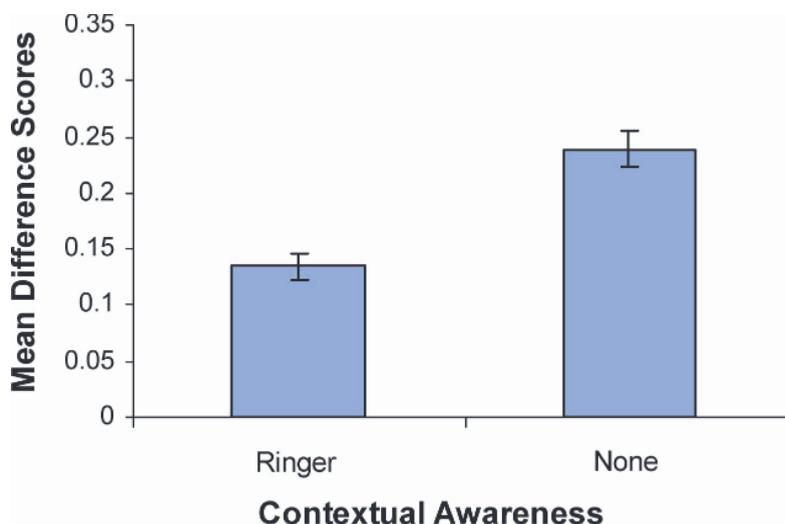


Figure 7. Mean absolute differences between Callers in the Ringer and None conditions (lower = better).

contextual awareness to distributed colleagues (see for example Perry *et al.* 2001, Dabbish and Kraut 2003).

Another goal of this study was to examine different types of contextual awareness information that could be provided and to compare their effects on callers' behaviour. Our results show that Callers did indeed use the contextual information provided to them when deciding whether to make a call or leave a voice mail. Callers, for instance, chose to call significantly less when the Receiver was in the office than when the Receiver was at home. Our results did not show significant dissimilarities between the mean differences of the various contextual information conditions. However, further investigation revealed that different types of contextual information caused different kinds of improvements in Callers' behaviour; knowledge that the receiver is in the office helped Callers avoid interrupting inappropriately (Location condition), while knowledge that the receiver was alone helped Callers interrupt when it was appropriate (People condition). We believe that this is due to the lower ambiguity associated with availability in these two conditions (office and alone). This finding is important when trying to predict whether certain sensed contextual information would create improvement in interruptions at appropriate times, or in avoiding interruptions at inappropriate times. A combination of the two types of sensed information (in the Combo condition) did not show superiority and may suggest that information about the presence or absence of people around the receiver took precedence over information about the receiver's location.

It is important to note, at this point, the four participants in Caller conditions who preferred an immediate interruption (call) over a delayed interruption (voice mail) 100 per cent of the time, regardless of the message or the contextual information provided to them. It is very likely that these participants represent a segment of the population for which this solution will not work (most readers are probably familiar with the experience of being interrupted by a telemarketer calling during dinner time). These four participants remind us that a caller-based solution cannot solve the problem if callers choose to ignore the information provided to them. Fortunately the vast majority of our callers (92%) chose otherwise.

A secondary goal of this study was to explore if merely the role of Caller or Receiver affects people's perception of urgency and their behaviour. Our findings suggest that role does affect behaviour, with Callers electing to call significantly more frequently than Receivers would like. This finding suggests that solutions that depend entirely on contextual information provided to Callers may not completely eliminate the problem of inappropriate calls. We demonstrated, however, that providing contextual information could significantly reduce the size of the discrepancy between callers' behaviour and receivers' desires. Additionally, our results suggest that perceived

urgency plays a significant role in the choices of both callers and receivers and that their perception of urgency is not significantly different across the range of contextual environments explored in this study. While this finding might seem intuitive, it should be a significant factor when deciding between an automated solution approach (which would not have knowledge of the urgency) and a caller-based solution as proposed in this paper.

8.1 Considerations for practice

The results presented above measure the benefit of different types of contextual information based on the match that they create between callers' behaviour and receivers' desires. While this is an important base measure, other factors need to be considered when judging a solution. These include the direction of errors (over-calling vs. under-calling), technological simplicity of the solution and privacy tradeoffs.

8.1.1 Over-calling vs. under-calling. When callers make an error, it can be of two sorts – calling when they should not and not calling when they should (or using voice mail when they shouldn't and not using it when they should). In our analysis, we regarded over-calling and under-calling to be of equal severity. Some, however, might regard these two types of errors as having very different values. For instance, one might consider being interrupted at an inappropriate time to be much worse than not being interrupted at a convenient time. These estimates might also differ between callers and receivers. In determining the value of a particular caller-based solution, such factors should be taken into consideration.

8.1.2 Technological issues. The choice of sensing technology can be critical for the deployment and acceptance of a caller-based solution. The options proposed in this study (location, people and ringer) each have different technological properties. Sensing the location of a person is indeed possible. However, sensing the exact location of a person indoors can be difficult and requires both knowledge of the structure (to determine the room the person is in) and a semantic knowledge of the structure (to know that the room is a 'meeting room'). Sensing whether there are people in the vicinity of the receiver is also possible (through vision, heat, or audio). Knowing that there are people around, however, does not tell us whether the receiver is socially engaged with these people. This might cause callers to under-call when there are people around the receiver even though the receiver is actually available for interruption (for a model of receiver's interruptibility see, for example Hudson *et al.* 2003). Finally, the state of the cell-phone ringer is a sensing solution with promising technological properties. The feature permitting different

ringer settings already exists in cell phones, is familiar to cell-phone users, and providing knowledge of it to callers is technologically simple. It does, however, have the drawback of depending on the receiver remembering to reset it.

8.1.3 Preserving privacy. A critically important issue surrounding the adoption of a specific strategy is that of privacy. Which information would people be willing to provide about themselves to reduce the problem of inappropriate incoming calls, and at what cost to their privacy? We did not measure or judge the privacy concerns associated with different information types, since it was of primary importance to understand whether or not a particular solution is feasible. However, providing information about a person's location can threaten privacy and expose sensitive information. Similarly, providing the knowledge that a person is not alone, regardless of information about social engagement, can also threaten the person's privacy. These possibilities raise a question as to the practicality of these two solutions. We do feel, however, that providing one's cell-phone ringer status does not raise many privacy concerns. Our results suggest that merely drawing attention to context even with very little contextual information may reduce unwanted interruptions. Thus an attempt should be made to provide information of minimal sensitivity.

9. Conclusions and future work

We have presented a study examining a solution to the problem of inappropriate incoming cell-phone calls. This solution involves providing callers with contextual information about the receivers, helping them to decide whether to call or not. We conducted two experiments to examine the effectiveness of different types of contextual information (location, presence of people and ringer status). In a simulation, participants playing either the roll of callers or receivers were asked to choose between making (or receiving) a call and leaving (or receiving) voice mail. Our results suggest that callers are able to make use of contextual awareness information provided to them to significantly improve their choices. Further, our results suggest that different kinds of contextual information generate different kinds of improvements: better interruptions and better avoidance of interruptions. Our results also suggest the existence of inherent differences between the roles of callers and receivers and that a good solution must allow the urgency of the conversation to be taken into account.

This research is part of a larger effort aimed at reducing the attention demands and interruptions that people are exposed to in their daily lives. We plan to extend this study to examine other caller-receiver relationships (such as a family member or boss), other locations (such as a coffee

shop or car), and other receiver activities (such as waiting at the airport). While providing the ringer status of the receiver's cell phone seems to be very promising, we need to study whether people use this feature consistently. As the need for control in our study dictated presenting a small number of contextual information levels, additional research is needed to examine finer levels of contextual information. In addition, future work should examine the extent to which real world callers modify their behaviours according to the findings we have presented here. We believe that the work reported here provides important direction and guidance for field researchers and system designers interested in designing caller-based solutions and studying their effects in the field. We also believe that this study makes a contribution in an implementation of a quick and comparatively inexpensive method for measuring the benefit and potential value of proposed sensing technologies early in the design cycle.

Acknowledgements

The authors would like to thank Susan Fussell, Robert Kraut and Yaakov Kareev for discussions and comments on the paper and analyses. This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. NBCHD030010, and by the National Science Foundation under grants IIS 0121560 and IIS 0325351.

Note

1. Because the independent variables were not completely orthogonal, we used Least Squared Means (LS Means) to compare experimental conditions. When calculating the means for an experimental condition, LS Means control for the value of the other independent variables. The means reported throughout this article are LS Means.

References

- BELLOTTI, V. and EDWARDS, K., 2001, Intelligibility and Accountability: Human Considerations in Context-Aware Systems. *Journal of Human-Computer Interaction*, **16**, pp. 193–212.
- DABBISH, L. and KRAUT, R., 2003, Coordinating Communication: Awareness Displays and Interruption. *Extended Proceedings of CHI*, pp. 786–787 (New York: ACM Press).
- EYROLLE, H. and CELLIER, J., 2000, The effects of interruptions in work activity: Field and laboratory results. *Applied Ergonomics*, **31**, pp. 537–543.
- GILLIE, T. and BROADBENT, D., 1989, What makes interruptions disruptive? A study of length, similarity and complexity. *Psychological Research*, **50**, pp. 243–250.
- GREEN, P., 2000, Crashes Induced by Driver Information Systems and What Can Be Done to Reduce Them. In *Proceedings of Convergence*, pp. 26–36 (SAE publication P-360) (Warrendale, PA: Society of Automotive Engineers).

- GRUDIN, J., 1988, Why CSCW applications fail: problems in the design and evaluation of organization of organizational interfaces. In *Proceedings of CSCW*, pp. 85–93 (NY: ACM Press).
- HORVITZ, E., JACOBS, A. and HOVEL, D., 1999, Attention-Sensitive Alerting. In *Proceedings of UAI*, pp. 305–313 (San Francisco: Morgan Kaufman).
- HUDSON, J.M., CHRISTENSEN, J., KELLOGG, W.A. and ERICKSON, T., 2002, Structure and Flow: 'I'd be overwhelmed, but it's just one more thing to do': availability and interruption in research management. In *Proceedings of CHI*, pp. 97–104 (NY: ACM Press).
- HUDSON, S.E., FOGARTY, J., ATKENSON, C.G., AVRAHAMI, D., FORLIZZI, J., KIESLER, S., LEE, J.C. and YANG, J., 2003, Predicting human interruptibility with sensors: a Wizard of Oz feasibility study. *Proceedings of CHI*, pp. 257–264.
- ISAACS, E., WALENDOWSKI, A. and RANGANATHAN, D., 2002, Hubbub: A sound-enhanced mobile instant messenger that supports awareness and opportunistic interactions. In *Proceedings of CHI*, pp. 179–186 (NY: ACM Press).
- LACOEHE, H. and ANDERSON, B., 2001, Interacting With The Telephone. Special Issue of the *International Journal of Human-Computer Studies* on 'Home Use of Information and Communications Technology', **54**, pp. 665–699.
- MILEWSKI, A.E. and SMITH, T.M., 2000, Providing presence cues to telephone users. In *Proceedings of CSCW*, pp. 89–96 (NY: ACM Press).
- MONK, A., CARROLL, J., PARKER, S. and BLYTHE, M., 2004a, Why are mobile phones annoying? *Behaviour & Information Technology*, **23**, pp. 33–41.
- MONK, A., FELLAS, E. and LEY, E., 2004b, Hearing only one side of normal and mobile phone conversations. *Behaviour & Information Technology*, **23**, pp. 301–305.
- NELSON, L., BLY, S. and SOKOLER, T., 2001, Interactive Quiet Calls: Talking Silently on Mobile Phones. In *Proceedings of CHI*, pp. 171–181 (NY: ACM Press).
- O'CONNAILL, B. and FROHICH, D., 1995, Timespace in the workplace: dealing with interruptions. In *Conference companion on Human factors in computing systems*, pp. 262–263 (NY: ACM Press).
- PALEN, L., SALZMAN, M. and YOUNGS, E., 2000, Going wireless: behavior and practice of new mobile phone users. In *Proceedings of CSCW*, pp. 201–210 (NY: ACM Press).
- PEDERSEN, E.R., 2001, Calls.calm: Enabling Caller and Callee to Collaborate. In *Extended Proceedings of CHI*, pp. 235–236 (NY: ACM Press).
- PERING, C., 2002, Taming of the ring: context specific social mediation for communication devices. In *Extended Proceedings of CHI*, pp. 712–713 (NY: ACM Press).
- PERRY, M., O'HARA, K., SELLEN, A., BROWN, B. and HARPER, R., 2001, Dealing with mobility: understanding access anytime, anywhere. *Transactions on Computer-Human Interaction (TOCHI)*, **8**, pp. 323–347.
- SCHMANDT, C., MARMASSE, N., MARTI, S., SAWHNEY, N. and WHEELER, S., 2000, Everywhere messaging. *IBM Systems Journal*, **39**, pp. 660–677.
- SCHMIDT, A., AIDOO, K., TAKALUOMA, A., TUOMELA, U., VAN LAERHOVE, K. and VAN DE VELDE, W., 1999, Advanced Interaction in Context. In *Proceedings of the First International Symposium on Handheld and Ubiquitous Computing*, pp. 89–101 (New York: Springer-Verlag).
- SCHMIDT, A., TAKALUOMA, A. and MANTYJARVI, J., 2000, Context-Aware Telephony Over WAP. *Personal and Ubiquitous Computing*, **4**, pp. 225–229.
- TANG, J.C., YANKELOVICH, N., BEGOLE, J., VAN KLEEK, M., LI, F. and BHALODIA, J., 2001, ConNexus to awarenex: extending awareness to mobile users. In *Proceedings of CHI*, pp. 221–228 (NY: ACM Press).
- WELLMAN, B., 2001, Physical Place and Cyber Place: The Rise of Personalized Networking. *International Journal of Urban and Regional Research*, **25**, pp. 227–252.
- WEI, R. and LEUNG, L., 1999, Blurring public and private behaviors in public space: policy challenges in the use and improper use of the cell phone. *Telematics and Informatics*, **16**, pp. 11–26.
- ZSAMBOK, C.B. and KLEIN, G. (Eds), 1996, *Naturalistic decision making* (Mahwah, N.J.: Lawrence Erlbaum Associates, Inc.).