

# Managing Availability: Supporting Lightweight Negotiations to Handle Interruptions

MIKAEL WIBERG

Umeå University, Sweden

and

STEVE WHITTAKER

Sheffield University, UK

---

Interruptions are a central aspect of working life. The prevalence of remote co-workers and the use of mobile technology mean that interruptions are more prevalent, and workers have to learn to manage availability. To understand general issues in availability management, we carried out a naturalistic study of how interruptions are handled in face-to-face situations. We found that availability management requires negotiation, that it is also highly dependent on awareness about the availability of others, and that it demands cognitive effort to shift attention to the interruption. On the basis of these observations, we developed a technology, named The Negotiator, that embodies three main design requirements: (a) support for negotiation, (b) contextual information about when a recipient is available for a call, (c) lightweightness to reduce attention overhead. We carried out an experimental study of interruption management using this technology. The interface satisfied the original design requirements, that is, people, were able to use it effectively to negotiate times to talk, while successfully carrying out an intellectually demanding activity. Contrary to our expectations, however, people preferred to take responsibility for returning calls rather than delegating them, and they preferred to schedule calls as soon as possible rather than deferring them. We suggest that there are social reasons why people do this. They feel a social obligation to return calls as soon as possible so as not to inconvenience others and also to be responsible for making these calls themselves. They also take calls sooner to avoid having to remember future conversational commitments. We discuss the theoretical and technical implications of these findings.

Categories and Subject Descriptors: H.5.3 [**Information Interfaces and Presentation**]: Group and Organization Interfaces—*Asynchronous interaction, synchronous interaction, computer-supported cooperative work*; H.4.3 [**Information Systems Applications**]: Communications Application; H.1.2 [**Models and Principles**]: User/Machine Systems; H.5.2 [**Information Interfaces and Presentation**]: User Interfaces—*Prototyping, user-centered design, theory and methods*; H.4.1 [**Information Systems Application**]: Office Automation—*Time management, groupware*; K.4.3 [**Computers and Society**]: Organizational impacts—*Computer supported Collaboratory work*

General Terms: Design, Experimentation, Human Factors

---

Authors' addresses: M. Wiberg, Department of Informatics, Umeå University, 901 87 Umeå, Sweden; email: mikael.wiberg@informatik.umu.se; S. Whittaker, Sheffield University, Sheffield, S1 2TN, UK; email: s.whittaker@shef.ac.uk

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or [permissions@acm.org](mailto:permissions@acm.org).

© 2005 ACM 1073-0616/05/1200-0356 \$5.00

ACM Transactions on Computer-Human Interaction, Vol. 12, No. 4, December 2005, Pages 356–387.

Additional Key Words and Phrases: Availability management, availability negotiation, awareness, interaction asymmetry, interaction management, interruptions, lightweight negotiations, mobile phones

---

## 1. INTRODUCTION

Conversational interruptions are a fact of working life. Many studies of organizational behavior have shown that workers spend large amounts of their time engaged in brief conversations [Hudson et al. 2002; Mintzberg 1973; Panko 1992; Whittaker et al. 1994]. About 90% of these brief conversations are unplanned [Kraut et al. 1993; Whittaker et al. 1994] which means that they are interruptive to at least one party [Rouncefield et al. 1995]. For example, Sproull [1984] found that managers spent 80% of their day in brief communications and that their activities were interrupted 21% of the time.

Studies of both face-to-face and technologically-mediated communication show that managing conversational interruptions is a complex process that involves careful negotiation [Fish et al. 1992; Isaacs et al. 2002; Nardi et al. 2000; Schegloff 1968; Tang et al. 1994, 2001a, 2001b]. For example, people are unreceptive to attempts to initiate communication at certain times. Hudson et al. [2002] found that managers judge themselves to be unavailable for communication 12% of the time. And many attempts to initiate opportunistic communications fail. Whittaker et al. [1994] found that 60% of work phone calls fail to reach their intended recipient. Together, these studies suggest that conversational initiation is problematic because negotiation is necessary before an impromptu conversation can begin and because such negotiation can fail.

There are many reasons why participants may reject attempts to initiate impromptu conversation. One major reason is cognitive overload. By agreeing to initiate a new conversation, recipients are forced to defer their current activity or conversation and the context of that current activity may be lost. A study of interruptions showed that interrupted parties fail to return to their prior activity 45% of the time [O’Conaill and Frohlich 1995]. A critical factor here is the *asymmetry* between initiator and recipient [Nardi et al. 2000; Tang et al. 2001a, 2001b]. If negotiation is successful and an impromptu conversation takes place, then the initiator is able to discharge their immediate communication responsibilities, thereby reducing their cognitive load. But for recipients, agreeing to a conversation means that they not only have to take on someone else’s communication goals, they also have to remember the content and status of their original interrupted activity. As a result, a recipient’s cognitive load is greatly increased. Other overload problems arise when incoming calls overlap with current conversations, interrupting or disrupting those conversations [Nelson et al. 2001; Pering 2002].

This suggests that recipients have to learn to manage availability. And recent developments in wireless technology exacerbate this by making people more contactable. Office workers are away from their desks for substantial periods of time, with figures ranging from 50–90% depending of the nature of the work [Bellotti and Bly 1996; Whittaker et al. 1994; Wiberg 2001]. In the past,

this would have meant that they were harder to contact. However, new technologies such as mobile phones and SMS mean they are now contactable even when they roam the office or are engaged in meetings. Various telecommunications companies have touted the value of *anytime, anywhere* wireless access, allowing immediate response to customers, clients, and co-workers [Wiberg and Ljungberg 2001].

One important characteristic of many technology-mediated communications is that, in contrast to face-to-face interaction, neither initiator nor recipient has information about the other's current situation. Face-to-face settings provide a rich set of visual and auditory information about the recipient's activities, allowing one to infer how interruptible they are [Dourish and Bellotti 1992; Fish et al. 1992; Kraut et al. 1993; Tang et al. 1994, 2001a, 2001b; Whittaker et al. 1994]. Such information is clearly absent when initiator and recipient are not co-present.

Our most pervasive communication technology, the telephone, also has little direct support for availability management. It is hard with the phone to shield oneself from unwanted interruptions [Brown and Perry 2000]. Switching the phone off or letting calls go through to voicemail may mean that important calls are not responded to immediately. And attempts to dynamically manage availability by switching the phone on and off are exceedingly laborious. Furthermore, the phone does not provide explicit support for call filtering. When available, caller ID offers the recipient information about the identity of the caller, but not about purpose of the incoming call. And from the call initiator's perspective, there are no reciprocal awareness mechanisms to provide information about the recipient's potential responsiveness, allowing initiators to determine an appropriate time to call [Milewski and Smith 2000; Nardi et al. 2000; Tang et al. 2001a, 2001b].

Together these observations suggest that we need to refine our empirical and theoretical understanding of *availability management* [Tang et al. 2001a, 2001b; Whittaker et al. 1997; Wiberg 2001]. By doing so, we should be able to devise general principles for the design of improved technology for availability management. Note that such designs can be agnostic about the characteristics of the ensuing interaction as to whether this involves computer-mediated or human-computer interaction. In either case, availability management is a pressing problem for which solutions are urgently required, and we would expect our empirical results and technology to generalize to both situations.

We follow the standard HCI method of ethnographically-motivated iterative system design [Rosson and Carroll 2002], and the structure of this article mimics this method. To better understand availability management processes, we first carried out an empirical study of how conversational initiation occurs in face-to-face situations. On the basis of this study, we generated novel design principles for availability management. We implemented an availability management system built according to these principles and present data from an empirical study in which people used the system to negotiate their availability. We also present an application of our design principles to a mobile setting where its use will experience critical availability management problems. We conclude with a discussion of our results, their design implications, and also




Fig. 1. Research site for the naturalistic study.

describe the theoretical implications, of our findings for conversational models. By following this method, we are able to (a) present novel observational data, (b) derive new design principles, (c) implement and test a novel system.

## 2. A NATURALISTIC STUDY OF AVAILABILITY MANAGEMENT

To better understand availability management processes, we conducted a naturalistic study [Solso 1998] of people's attempts to establish face-to-face interaction in an office building during a three-month period. We collected data using a web camera with motion detection, allowing the camera to start recording as soon as somebody moved along in the hallway. Thus, each attempt to initiate interaction could be monitored. We informed people beforehand that such recording was taking place. The research site where the observations were made is illustrated in Figure 1:

The office environment consists of a hallway with offices, a copy and printer room, and a conference room. Next to the conference room is a door leading to an office environment with 12 workstations. There is also a lecture room next to the hallway. There are a lot of people moving around this building during office hours. As Figure 1 shows, the camera captures people about to enter the office, people passing by the office door in the hallway, and people in the copy and printer room. Although our observations were limited to recording from only a single location, that location nevertheless allowed us to record one private space (A's office) and two public spaces (the hallway and the copy and printer room). It is well known that people tend to be more open to impromptu conversations in public spaces [Whittaker et al. 1994]. By recording conversations in both private and public spaces, we hoped to sample different styles of conversation and to include different individuals. However, it is clear that there are limits in the data we collected because the private space observations were focused around A, an academic.

During the three months of observations, we recorded 120 initiations. We first present some general statistics about initiations. First, these were brief, lasting 7 seconds, measured from the time of the initial request to the agreement to start interacting. This brevity was reflected in the fact that they took only 2.1 conversational turns, where turns were defined to include gestures and

Table I. Distribution of Participants in Our Availability Negotiations

No. of Participants	No. and Frequency of Recorded Conversations
2	207 (82%)
3	31 (12%)
4	9 (4%)
5	5 (2%)
Never more than 5 participants	Total: 120 videos, 252 participants total

other non-speech sounds. They were also generally successful, with 93% leading directly to the establishment of interaction. In the 7% of unsuccessful cases, conversation was deferred rather than aborted completely, with participants agreeing to a later time to begin the conversation. These conversations were also typically dyadic with an average of 2.1 people being involved, the exact distribution is shown in Table I.

We also observed the extent to which availability negotiation placed attentional demands on the recipient. Participants had some difficulty in disengaging attention from the activity they were engaged in before the interaction request. During the initiation phase, they switched their attention back to that interrupted activity an average of 2.9 times. Attention switches were defined as actions taken to note the status of a current activity before accepting an interruption, for example, looking back at the computer screen or looking back at papers on the desk. They also included attempts to construct reminders about the status of the interrupted activity, allowing people to reinitiate the activity more easily after handling the interruption. Examples of reminders included taking a quick note about current work, or moving an important paper to a salient location, for example, putting it on the keyboard, to guarantee it would be returned to. The fact that such attention switches occurred during negotiations shows both that interruptions are disruptive to current activity, and also that recipients make some attempt to minimize disruption by attempting to retain the context of interrupted activities.

To give the reader a feel for availability management, we present three illustrative examples. The examples are complete, transcribed video-recorded episodes of negotiations. After each utterance we present interpretive and contextual information.

*Example 1.* Awareness, a brief negotiation, and a short conversation.

A IS SITTING IN FRONT OF HIS COMPUTER TYPING A DOCUMENT WHEN HIS COLLEAGUE B WALKS INTO VIEW. B NEEDS TO TALK TO A ABOUT THE RESCHEDULING OF A PROJECT MEETING.

1. B: [WALKS UP TO THE OFFICE DOOR, SIGNALING HER APPROACH BY STAMPING HER FEET RATHER MORE LOUDLY THAN NORMAL] (B has signaled her intent to interact but can see that A is busy, so waits for A's attention to initiate interaction)
2. A: [LOOKS UP AT HER, LOOKS BACK ON THE SCREEN TO NOTE WHERE HE WAS, LOOKS BACK AT B AGAIN] 'Yes?' (The response indicates that he is available for interaction. The conversation now begins.)
3. B: 'Yeah, it's about the meeting tomorrow, could we do that

after lunch?’’ (No context supplied about the meeting, and a rather quick question since B observed that A was currently involved in another activity)

4. A: ‘‘Sure.’’ [THEN QUICKLY LOOKS BACK AT THE SCREEN] (No continuation indicates the interaction is over and that A needs to continue his work)
5. B: [B NOTICES A’S GLANCE AT THE SCREEN AND RESPONDS] ‘‘Great, see you!’’ (B leaves the room).

In this example, the whole interaction only lasts for about 15 seconds mainly involving nonverbal communication. This negotiation is quick, lightweight, and effective. It takes only two turns, and involves the recipient saying one word: yes. It also relies on visual awareness information, and timing is critical as B waits for A to look up before she talks. Note that the recipient A agrees to talk *before he is aware of the topic under discussion*, which is arranging next day’s meeting as mentioned in line 3.

The problems of disengaging attention and the role of nonverbal signals are shown in the next example.

*Example 2.* Negotiation involving a short postponement of upcoming interaction.

A IS SITTING IN FRONT OF HIS COMPUTER READING AN EMAIL. F COMES BY TO ASK WHETHER A WOULD LIKE TO GO FOR A CUP OF COFFEE.

1. B: ‘‘Coffee?’’
2. A: [WITHOUT TAKING HIS EYES OFF THE SCREEN] ‘‘Yeah, that would be great!’’ [THEN HOLDS UP HIS RIGHT HAND SO THAT THE BACK OF HIS PALM COVERS HIS CHEEK] (The gesture indicates that he is interested but needs a few seconds)
3. B: [WAITS]
- A: [4 seconds later] [A LOWERS HIS HAND AND TURNS HIS HEAD TOWARDS B] ‘‘OK, now I’m ready’’ [the negotiation ends and A follows B out of the room].

Again the negotiation is extremely brief, requiring very few turns and exploiting situational awareness information. Here, however, the recipient uses gesture to briefly defer interaction—signaling “just a moment” in response to the interaction request. The lowering of the hand and the turning of his head towards B indicates his availability even before answering, “OK, now I am ready”. So although the interaction is successful in initiating conversation, A ensures that it does not begin until he has reached an appropriate stopping point in his prior activity. So an important part of the negotiation here is about exactly when the conversation should start. The need to defer conversation, albeit briefly, illustrates some of the cognitive problems that recipients experience in trying to switch attention.

We now present an excerpt from a longer hallway conversation between two people A and B, showing how a potential third participant determines whether the time is appropriate for an interruption.

*Example 3. Availability check without interruption.*

A AND B ARE DISCUSSING AN ONGOING PROJECT IN THE HALLWAY OUTSIDE A'S OFFICE. C PASSES BY ON HIS WAY TO THE COPY AND PRINTER ROOM.

1. A: '“Wouldn't it be nice if we could also involve them in this project...”'
  2. B: '“Yeah, but, on the other, I don't think that...”' [B THEN LOOKS AT C AND THEN TURNS HIS HEAD BACK TO A AGAIN] (Indicates he noticed C but it was not an appropriate time to involve C in the discussion)
  3. C: [APPROACHES A AND B BUT NOTICES THAT THEY ARE BUSY IN CONVERSATION AND CONTINUES INTO THE PRINTING ROOM]
  4. B: [CONTINUES THE CONVERSATION] '“as I said, I don't think that it's good if the project is too big. It will be a lot of overhead work with project coordination and administration.”'
  5. A: '“Yeah, maybe we can talk about this tomorrow?”'
  6. B: '“OK”' [THEN B GOES BACK TO HIS OFFICE]
  7. C: [A COUPLE OF MINUTES LATER, C COMES BACK FROM THE COPY ROOM AND ENTERS A'S OFFICE] '“Are you busy now?”'
  8. [A LOOKS UP FROM HIS COMPUTER] '“no..., please have a seat”' (invitation to start a conversation)
- C: [C SITS DOWN IN THE CHAIR BEFORE STARTING TO TALK] '“so...what do you think about my proposal?”' (A and C have discussed a project proposal a couple of days earlier, and the full conversation now follows)

This illustrates C's careful monitoring of another potential participant A's availability before attempting to establish interaction. C determines that his initial attempt to engage A is inopportune and defers this until A and B are finished. This kind of awareness seems to be an important part of the lightweight and effortless management of availability, interruptions, and sessions in face-to-face settings. Again the point at issue is exactly *when* it is opportune to begin the conversation. Note again that, as in Example 1, the recipient does not wait to determine the topic of the conversation before agreeing to talk (line 8).

Taken together, these data suggest four critical properties of availability management.

- (1) *Negotiation*. Even though almost all interactions concluded with the participants agreeing to communicate, it is obvious that this conclusion is arrived at only after negotiation. As Examples 2 and 3 show, initiators can't just barge in and begin speaking: they have to take into account the recipient's current activity and problems with context switching. The conversation can only begin when the recipient signals that they have disengaged from their prior task. These observations about the importance of negotiation are consistent with findings from face-to-face communication [Schegloff 1968, Whittaker et al. 1994], video-mediated communication [Bly et al. 1993; Fish

et al. 1992; Tang et al. 1994], and instant messaging [Isaacs et al. 2002b; Nardi et al. 2000]. Note also that negotiations do not generally include the initiator stating the purpose of the interaction. The important issue seems to be to negotiate when to talk, not what to talk about. This is because participants can often infer the likely conversational topic from context and the identity of the initiator [Whittaker et al. 1994].

- (2) *Awareness*. All examples show the importance of awareness information about other participants' current activities. This information critically determines whether and when communication begins. For instance, in Example 3, C initially did not attempt to negotiate A's availability because he could see A was engaged in another conversation. Again, this is consistent with work on face-to-face and technology-mediated communication [Fish et al. 1992; Isaacs et al. 2002; Nardi et al. 2000; Tang et al. 1994, 2001a, 2001b; Whittaker et al. 1994].
- (3) *Brevity*. In all cases, negotiations were extremely brief (lasting 7 seconds on average). Other work has shown that entire opportunistic conversations only last about 2 minutes [Kraut et al. 1993; Whittaker et al. 1994] so conversational initiation must be brief compared with this. Brevity is made possible, as the examples show, by having access to awareness information about the activities of others.
- (4) *Attentional Disengagement and Cognitive Load*. All the examples show that recipients had problems in context-switching from their prior activity. On average, they switched attention back to that activity 2.9 times, and they carried out activities to help them remember prior task context.

One assumption in generating these principles is that our observations of face-to-face conversations will generalize appropriately to technology-mediated settings [Whittaker 2002]. But face-to-face and phone conversations may have different characteristics because the phone allows more general public access, making it more difficult for participants to infer the purpose of the call from caller ID. Nevertheless, our observations generally confirm findings from other research examining the initiation of technology-mediated conversations suggesting such generalization is justified [Fish et al. 1992; Isaacs et al. 2002a; Nardi et al. 2000; Tang et al. 1994, 2001a, 2001b].

### 3. DESIGNING THE NEGOTIATOR: SUPPORTING AVAILABILITY NEGOTIATIONS ON MOBILE PHONES

We now describe a system design that is intended to meet these four requirements. We developed a system called The Negotiator that embodies the three main design characteristics following directly from our naturalistic requirements: (a) support for negotiation about when to talk, (b) awareness information about when a recipient is available for conversation, (c) lightweightness both to ensure that negotiations are brief and to reduce attentional overhead.

#### 3.1 The Negotiate Talk Time Model

To make the negotiation process as lightweight as possible, we chose a familiar metaphor related to timing an activity, that is, a simple timer. We believe the






Fig. 2. The temporal negotiation loop.

timer is appropriate since it has an interface that is very simple (the only thing you do is to turn the knob around) and unambiguous (i.e., you set a time with a timer; you do nothing else with it). Both simplicity and clarity are critical design properties for a device that needs to minimize cognitive overhead. Furthermore the central function of the timer is to control the timing of an event which is consistent with our participants' focus on negotiating when to talk.

So, those wishing to initiate a conversation adjust the timer to indicate to the recipient when they want to talk. And responding to an incoming interaction request is just a matter of selecting an appropriate time and sending back the reply as illustrated in Figure 2. If the suggested time isn't appropriate, the initiator can respond with another suggested time, creating a negotiation loop. Note that the negotiation loop is about when to speak and not what to speak about. This is consistent with our observational data.

The Negotiate talk time model consists of two people (A and B). A is someone who wants to get in touch with B. In this model A is labeled the initiator, and B is labeled the recipient (of an interaction request). The model does not include the interaction that follows the negotiation. Thus, B is not a recipient of the content of the interaction, only a recipient of an interaction request. Negotiation is necessary to find a time for the interaction that is suitable for both parties. Thus, the overall goal is to for the initiator and recipient to find a suitable time by negotiating their availability until they reach an agreement.

How does this model support the three design requirements we derived from the naturalistic study, i.e., negotiation, awareness, and lightweightness?

The concept of negotiation is illustrated in Figure 3 by the arrows that bring A and B's suggestions about an appropriate time for the interaction closer to each other. Participants suggest when they want to talk by selecting a point on the timeline shown between A and B. Selecting a point at the left end indicates a request to talk soon, and selecting one on the right indicates a later request time. The numbers indicate the sequenced phases of the negotiation. A makes a request to interact immediately (i.e., the arrow from A to B (1)). But B finds the time inappropriate so he/she responds with a suggestion to postpone the conversation until much later (2). A finds the proposed time too late so he/she suggests an intermediate time (3) between (1) and (2). B still thinks the time is inappropriately soon, so he/she suggests a time somewhere in between (3) and (2), namely (4). A concludes the negotiation by confirming to B that he/she will get back to him/her at that time.




Fig. 3. Scheduling Interaction in the negotiation loop.

In the above model, awareness is realized by passing suggestions of timing back and forth. These provide implicit information about both the busyness of the recipient and the urgency of the initiator's projected call. Thus, the recipient of an interaction request becomes aware that initiators urgently seek contact with him/her simply by them repeatedly requesting an immediate conversations. At the same time, the timing of the recipients' responses indicates their availability to the initiators. We also provide information about the identity of the initiator (using caller ID). As we have seen, caller identity is a good (if not foolproof) clue to the subject matter of the call.

We also offer support for lightweightness. Instead of requiring participants to provide detailed contextual information about their availability or the urgency of their communication request, the model relies on the simple transmission of information about the projected time for a call. Not only is this consistent with our naturalistic data, it should also be supportable by a simple UI. This should expedite the negotiation process and support brief negotiations. Similarly, by providing a very simple interface and operations, we hope to minimize the interference with users' prior tasks and allow recipients of interruptions to smoothly switch context.

### 3.2 Design Reasoning and Implementation of the User Interface

Our Negotiate talk time model is very different from other systems supporting conversation initiation and availability management. The Negotiate talk time model differs from other attempts to support initiation that rely on profiles, for example, allowing users to indicate their availability using a fixed set of states, such as, busy, available, and so on [Milewski and Smith 2000; Tang et al. 1994]. One problem with such an approach is that participants often forget to set or reset their profiles and as a result, that profile information is usually not representative of their current state [Milewski and Smith 2000; Tang et al. 1994]. More importantly, it seems that one's availability may be highly situated, depending on the identity of the current caller and the nature of their request [Nardi et al. 2000; Tang 2001a, 2001b], making it difficult for participants to set their availability for all projected calls. The situatedness of availability also argues against approaches that finesse negotiations of availability by

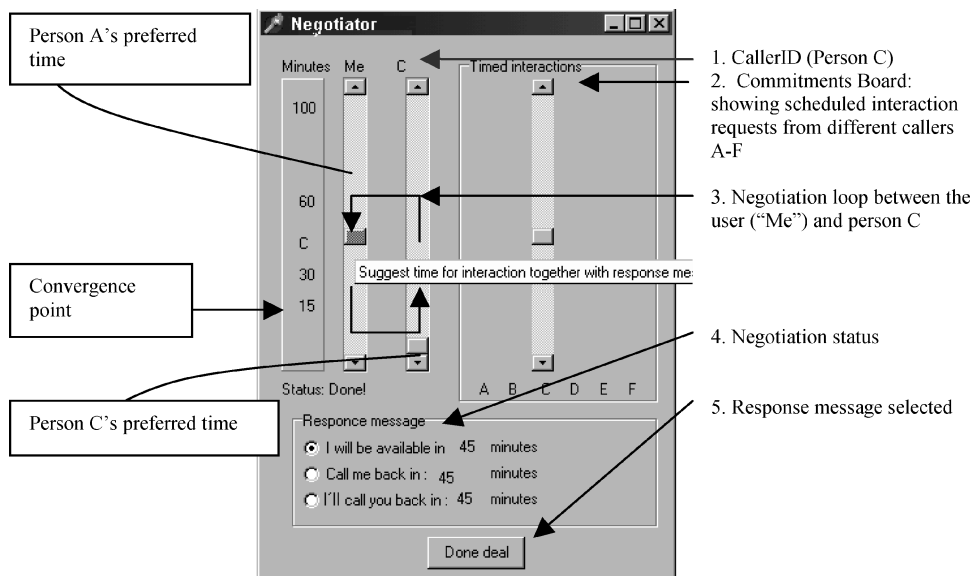


Fig. 4. The Negotiator user interface.

analyzing the recipients' visual or speech behaviour to automatically determine their availability [Sawhney and Schmandt 1999; Vertegaal et al. 2002]. Such systems would have to be highly sensitized to contextual parameters and recipient preferences to successfully determine whether a recipient will take a call. Our model also differs from systems like Instant Messaging that provide implicit information about the recipient's current activity by transmitting their recent keyboard activity [Nardi et al. 2000; Isaacs et al. 2002a, 2002b; Tang et al. 2001a, 2001b]. Such systems provide no explicit system support for negotiation [Nardi et al. 2000], although they can provide information about impending initiation attempts [Isaacs 2002a; Tang 2001a, 2001b]. Finally, our design does not allow initiators to supply the purpose of their call (e.g., via SMS) as our observational data showed that the negotiation was about when to talk not what to talk about.

A critical target use for the technology was to handle requests for conversations during meetings. We therefore chose a graphical interface and input technique since people already operate their mobile phones and Blackberries unobtrusively in meetings. We believe that it would be inappropriate to manage availability by voice commands as these could disrupt the meeting especially since speech technology remains error prone. Figure 4 illustrates the user interface. The goal is to make both participants' current suggestions for a time to talk obvious to highlight the differences between these suggestions, hopefully allowing rapid negotiation to an agreed upon time and to show various conversations that have already been scheduled. We have implemented versions of the Negotiator that run on laptop computers and smart phones, such as Qtc 1010 Pocket PC running PocketPC phone edition.

In Figure 4, the initiator's Caller ID (1) indicates who is calling, (in this case, the number is identified as person C). If the suggested time is appropriate, the recipient agrees to take the call at the proposed time by pressing the agree (done deal) button. If it isn't appropriate, the recipient negotiates a more appropriate time (3) by altering the scrollbar on the left to suggest an alternative time, choosing a response message (5), and replies to the initiator. In the example, the initiator finds the recipient's suggestion ("I'll call you back in 45 minutes") to be satisfactory and confirms this by pressing done deal. The recipient sees the acceptance (4) and clicks on the done deal button to put the caller on the Commitments Board showing agreed upon upcoming interactions with different people (2). The availability negotiation is lightweight requiring just four simple UI actions: altering the scrollbar to an appropriate time, choosing a response message, sending the suggestion by pushing a button, and clicking the done deal button to put the deal on the Commitments Board (2).

We were confronted with some complex design decisions in endeavouring to meet our overall goal of supporting lightweight negotiation. One issue was whether response messages should be graphical or text-based. We concluded that a graphical representation (e.g., a picture of a timer with a time set on it) could be ambiguous, that is, it could mean "I'll call you back in 30 minutes" or "I am available in 30 minutes". This is related to the *responsibility handover* problem which concerns whether the initiator or recipient wants to be responsible for making the actual call. We opted for flexibility here, allowing participants to send messages that indicated they were taking responsibility for the call "I'll call you back in 5 minutes," in contrast to "Call me back in 10 minutes" (which is the initiator's responsibility) or "I am available in 20 minutes" (nobody's responsibility). To reduce the cognitive load of remembering to initiate agreed upon calls, we expected recipients to prefer not to take responsibility for making the call and to prefer the "Call me back in X minutes" or "I am available in Y minutes" messages. The default mode was set to the delegate alternative, that is, the "I am available in Y minutes" message, and the program automatically alters the radio buttons to that alternative after each completed negotiation loop.


A final note concerns the seeming complexity of the interface. Although there appear to be multiple graphical elements to keep track of, it is important to consider what recipients need to focus on and what needs to be done to be able to respond. As Figure 5 shows, the recipient needs only to direct his/her attention to one single area of the user interface to handle an incoming request and start the negotiation. The recipient's scrollbar is labeled "Me" and the initiator's scrollbar is labeled with the Caller ID (in this case C). If the recipient is unable to take the call immediately, he/she alters his/her scrollbar to send an deferring message "I'll call you back in 10 minutes" and returns to his/her current activity. However, our concerns about UI complexity meant that a major objective of our study was to determine whether our interface was lightweight enough for this cognitively demanding situation.

#### 4. EXPERIMENTAL STUDY

We now describe an evaluation of the system. The experimental study had four main goals: (a) to determine how well the interface supports lightweight

**Active field/Explicit part**

The recipient of an incoming call uses this scrollbar. By altering the scrollbar the recipient can easily respond to an incoming interaction request.

**Passive field**

This field is passive for the recipient. It presents alternative call initiation times. By having this presented next to the recipient's scrollbar he/she can see the initiator's suggestion at a glance and decide if the time is appropriate. If the suggested time is inappropriate the recipient can simply modify his/her scrollbar again to suggest an alternative time.

Fig. 5. The Negotiation loop implemented in the user interface.

negotiations; (b) to collect objective data about how people negotiate availability and manage interruptions, specifically to test hypotheses both about users' strategies for minimizing cognitive load while handling interruptions and about the effects of their prior real-world experience of handling interruption; (c) to record subjective comments about the experience of using the software to manage availability to probe very situational aspects of dealing with interruptions; and (d) together suggestions about how to improve the technology. We used an experimental study to evaluate the implementation because this allowed us to observe people under conditions where we had control over the frequency and nature of interruptions.

#### 4.1 Users

Forty people took part in the study. They were allocated equally to experimental and control conditions. Experimental users had to carry out a primary task that required their full attention, while, at the same time, carrying out a secondary task of handling repeated requests for urgent conversations under conditions that simulated natural usage. Half the experimental users used the Negotiator to manage these interruptive requests, while the other half relied on pen and paper. Controls were not subjected to interruptions; they simply had to carry out the primary task. The exact breakdown of users was therefore: 10 Negotiator, 10 Pen and Paper, and 20 Control. All participants used mobile phones in their everyday work and social life, and we collected detailed information about the two experimental groups' usage patterns, work activities, and methods for dealing with phone calls when busy (see Table II). The Controls provided a baseline as they were uninterrupted, and the Pen and Paper group provided data about how people schedule conversations using familiar 'technology'.

#### 4.2 Method

Each of the three groups underwent a slightly different procedure. For the Control group, we explained the primary task which they then carried out. For the Pen and Paper group, we explained the primary task in the same way, but then told them that they would also have to carry out a secondary task of scheduling various impromptu interaction requests using pen and paper.

Table II. Participant Demographics, Activities, Phone Usage and Everyday Interrupt Strategies for Two Groups Receiving Interruptions

Participant	Male/ Female	Average # Calls/ Week	Average # Meetings/ Week	Average # Busy Hours?	Real-Life Strategies for Handling Mobile Phone Interruptions When Busy (Offline, online, VoiceMail, "go offline", etc).
<b>N1</b>	F	38	4	15	Turned off + voicemail
<b>N2</b>	M	15	4	20	Turned off or left phone at the office. Very occasionally took phone into meet- ings when expecting an urgent call.
<b>N3</b>	M	35	5	30	Turned off + Voice mail or silent signal.
<b>N4</b>	F	20	2	12	Vibrator signal + no ringer, or turned off. No voice mail. Sometimes needed to dis- connect incoming calls when in meetings.
<b>N5</b>	M	25	5	15	Turned off + voice mail.
<b>N6</b>	M	5	10	15	Turned off or left phone in office.
<b>N7</b>	F	80	4	25	Turned off.
<b>N8</b>	F	10	5	15	Turned off or leaving the phone at the office.
<b>N9</b>	F	30	8	7	Vibrator signal + no ringer, or turned off. Uses voice mail. Sometimes needed to disconnect incoming calls when in meetings.
<b>N10</b>	M	4	3	20	Turned off + voice mail.
<b>C1</b>	M	5	3	15	Turn off + voice mail
<b>C2</b>	F	6	5	20	Turn off.
<b>C3</b>	F	12	2	8	Turned on, vibrator, no ringer.
<b>C4</b>	M	3	8	10	Turned off + voice mail
<b>C5</b>	F	20	5	15	Turned on, Vibrator, no ringer
<b>C6</b>	M	5	3	5	Turned off + voice mail
<b>C7</b>	M	10	6	15	Turned off.
<b>C8</b>	M	12	2	20	Turned on, vibrator or silent signal
<b>C9</b>	F	15	6	15	Turned off.
<b>C10</b>	F	4	4	10	Turned off, voice mail

The procedure for the Negotiator group was more complex, involving three phases: (a) a training phase in which we instructed users about the software allowing them to explore its features; (b) an experimental phase where they carried out the primary task while using the Negotiator to schedule impromptu conversations; (c) a follow-up interview where we asked questions about their experience of using the software and about its overall design.

During training, we explained each of the Negotiator's features and demonstrated its use in several simulated negotiations. We did not proceed to the experimental phase until we were sure that the users understood the system. Users experienced little difficulty in learning the system and this took about 10 minutes on average.

Both the experimenter and the users were in the same room for the Negotiator and Pen and Paper conditions. This allowed the experimenter both to observe the user in detail and to answer any pressing user questions, although it does decrease the ecological validity of the setting, since, in real life, caller and recipient are not in the same room.

### 4.3 Experiment

*Primary Task.* The experimental setting was intended to simulate a naturalistic work setting with constant interruptions. All users carried out a primary task which was to watch complex videos in order to be able to answer detailed questions about them later. Users watched the videos in 3 separate 15-minute sessions, interspersed by two 10-minute breaks, for a total of 65 minutes. It is important to note that the breaks were not real breaks from the experiment but part of the procedure that potentially allowed a user the chance to take a call without missing something important in the video. By providing these breaks when there was no primary video-watching task, we could determine whether users preferred to take calls immediately or whether they wanted to defer calls and respond to the callers during breaks in the video. The users were informed that the breaks were part of the experiment and that no videos would be shown them. They were also told that they could take notes about the videos which were demonstrations of 5 different technologies from the CHI99 conference. They were not allowed to rewind or review the videos in any way. Immediately after the experiment, we asked them 10 questions about the contents of the videos, two questions about each of the 5 technologies they had seen. Details of the videos and questions are presented in Table III. As demonstrated in the table, the questions required users to pay close attention to the details of the videos as they were shown. We stressed to participants that this was the task that they should focus on.

*Secondary Task.* In addition to attending to the videos, both Negotiator and Pen and Paper users had to carry out a secondary task of scheduling a series of urgent impromptu conversations. While watching the videos, they received a number of requests for phone calls. At various points in the video, the experimenter contacted the user and asked them if they could talk on the phone right then. Users then had to schedule these calls using the Negotiator or Pen and Paper so as to minimize their effect on the primary task, that is, to schedule these during a break or after the experiment. Users were told that phone calls could be of differing importance, ranging from issues like “Your house is on fire!” (i.e., an emergency) to typical everyday coordination topics like “Can you pick up the children after school?” (i.e., an important, but not urgent issue). They were not told the importance of each call in advance, and the only information provided was Caller ID. Most importantly, they were told not to schedule any calls during the videos themselves and to avoid scheduling two calls at the same time. They were told that, if they wished, they could use the breaks for such conversations. However if they were too busy in the breaks or just needed the break to relax for a couple of minutes, they were told that that could postpone calls for very long periods (e.g., until some time after the experiment session) rather than taking them in the next break. Negotiations continued until both participants agreed on a time to talk. The experimenter varied his responses depending on the user’s response. In general, he tried to suggest a time that was intermediate between the time suggested by the user and the current time. So, for example, if the user wanted to talk in 10 minutes, he suggested 5 minutes. If the user suggested 20 minutes, he replied with

Table III. Details of Videos and Follow-Up Questions for the Primary Task

<b>5 video presentations from the CHI 1999 video program and 10 follow-up test questions (two per presentation)</b>
<b>Presentation 1:</b>
Druin, A., Montemayor, J., Handler, J., McAlister, B., Boltman, A., Fiterman, E., Plaisant, A., Kruskal, A., Olsen, H., Revett, I., Schwenn, T. P., Sumida, L., & Wagner, R. (1999). Designing PETS: A Personal Electronic Teller of Stories, <i>Proceedings of ACM CHI 99 Conference on Human Factors in Computing Systems</i> (Vol. 1, pp. 326–329).
Follow-up questions:
1. What methods were used in this study?
2. What did they label the design group?
<b>Presentation 2:</b>
Lindeman, R. W., Sibert, J. L., & Hahn, J. K. (1999). Towards Usable VR: An Empirical Study of User Interfaces for Immersive Virtual Environments, <i>Proceedings of ACM CHI 99 Conference on Human Factors in Computing Systems</i> (Vol. 1, pp. 64–71).
Follow-up questions:
3. What was the main difference between pointing in the air vs pointing at a flat board?
4. What was the main advantage with this approach?
<b>Presentation 3:</b>
Toyoda, M., & Shibayama, E. (1999). Hyper Mochi Sheet: A Predictive Focusing Interface for Navigating and Editing Nested Networks through a Multifocus Distortion-Oriented View, <i>Proceedings of ACM CHI 99 Conference on Human Factors in Computing Systems</i> (Vol. 1, pp. 504–511).
Follow-up questions:
5. What was the main problem with traditional interfaces that the Hyper Mochi Sheet tried to overcome?
6. What was meant by a “Predictive Focusing Interface”?
<b>Presentation 4:</b>
Hinckley, K., & Sinclair, M. (1999). Touch-Sensing Input Devices, <i>Proceedings of ACM CHI 99 Conference on Human Factors in Computing Systems</i> (Vol. 1, pp. 223–230).
Follow-up questions:
7. What was the motivation behind making the icons invisible when the user took his/her hand off the mouse?
8. What was claimed to be the advantage with this approach?
<b>Presentation 5:</b>
Mynatt, E. D., Edwards, W. K., LaMarca, A., & Igarashi, T. (1999). Flatland: New Dimensions in Office Whiteboards, <i>Proceedings of ACM CHI 99 Conference on Human Factors in Computing Systems</i> (Vol. 1, pp. 346–353).
Follow-up questions:
9. How did they split an information object on the whiteboard?
10. What happened to an information object that was drawn towards a corner or side of the whiteboard?

10 and so forth. If, however, the user suggested extremely long deferrals, for example, 80 minutes, he would reiterate the request to talk now. Users could also renegotiate times if it happened that they were unable to talk at a time that they had agreed upon earlier with the experimenter.

For both Pen and Paper and Negotiator conditions, we wanted to simulate calls from multiple participants. Although it would have been desirable to have 5 different people make these calls, it would have been impractical to arrange for 5 different experimenters to negotiate different call times. A single experimenter therefore assumed 5 different identities (A, B, C, D, E) when making



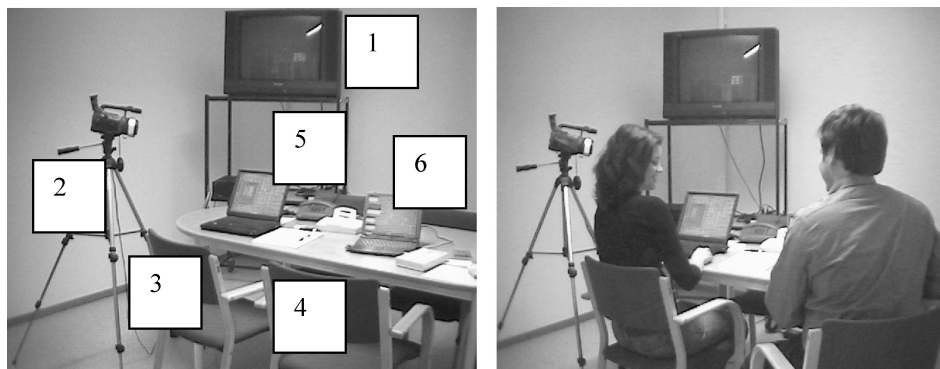


Fig. 6. Experiment setting.

the calls. He prefaced each call by introducing himself. After each negotiation was complete, users had to remember the identity of each caller (e.g., person A) and the time when they agreed to talk to that person.

The Negotiator interface provides Caller ID information about who is calling, and the Commitments Board records the times when calls have been agreed upon. In order to increase cognitive load, the experimenter's attempts to initiate conversations were timed to coincide with the more complex parts of the video. Participants received an average of 7.9 interrupts in the entire 65-minute session. Our original intention had been for each participant to receive an equal number of interruptive requests but because some negotiations became protracted, we were unable to balance these numbers exactly. When they occurred, actual conversations were brief. The conversation was mainly to check whether users remembered the identity and time of the call, since our main focus was on negotiating when to talk and not on supporting the conversations themselves.

Figure 6 shows the setup for the experiment for the Negotiator group. The experimenter sat beside the user as illustrated in Figure 6 (right). The video player (1) was used to present the primary task videos, and a video camera (2) was used to capture the users' speech and their use of the system. The user was placed to the left (3) and the experimenter to the right (4) to enable direct observations of the subject. The experimenter made incoming calls to the phone on the table (5). Two laptop computers (6) connected via a WLAN were used to negotiate an appropriate time for interaction between the experimenter (in one of his 5 different guises) and the user. We used the laptop implementation as this provided more straightforward logging capabilities.

The set up for the Negotiator and the Pen and Paper groups was intended to simulate an interruptive work environment: the video simulated having to attend to a meeting or being engaged in focused work while being interrupted by external calls and having to defer the interrupting calls to a more convenient time.

*Measures.* We recorded the following data. For all users on the primary task, *memory for the videos*, we recorded users' responses to each of the 10 questions. For the secondary task (carried out only by the Negotiator and Pen and

Paper groups), negotiating times to talk, we recorded whether the outcome of the negotiation was successful. We defined unsuccessful negotiations as when (1) users completely failed to respond to the experimenter's request for a conversation; or (2) they took the call while watching the video, thus interrupting their primary activity; (3) they inadvertently scheduled two calls in to overlap. We also logged information about each phase of the negotiation using the Negotiator including (a) the delegational type of negotiation response, that is, whether the user generated the "I'll call you back in X minutes", "Call me back in X minutes" or "I am available in X minutes"; (b) the time they suggested at each phase of the interaction, that is, whether they suggested they wanted to talk/were available/expected a call in 5 minutes, 10 minutes, 80 minutes etc.

#### 4.4 Negotiator Post-Test Questionnaire

There were 10 post-test questions given to the Negotiator group addressing (a) the design of the tool; and (b) users' experience in the experimental situation, in particular their ability to focus on the primary task and to handle interruptions.

- 1) Did you feel comfortable using this tool? (if yes, why? / if no, why not?)
- 2) What were its main advantages?
- 3) What were its main drawbacks?
- 4) Would you like to use it on a regular basis?
- 6) Did it have any unnecessary features?
- 7) Did you feel that you were able to focus on the presentation?
- 8) How did the incoming phone calls affect you and your attention?
- 9) How did the tool affect you and your attention?
- 10) Any other comments?

#### 4.5 Hypotheses

We had four main experimental hypotheses, relating to (a) support for lightweight negotiations, (b) cognitive load, and (c) users' prior experience of handling work-related interruptive calls.

4.5.1 *Lightweightness and Support for Negotiation.* The Negotiator is designed to support negotiation. It should also allow users to carry out secondary task negotiations while retaining focus and performance on the primary task. This led to two hypotheses:

- Lightweightness.* The straightforwardness of using the interface should mean that the Negotiator group should perform as well on the primary task as the Control group (who were not interrupted).
- Support for Negotiation.* The Negotiator group should perform better on the primary task than the Pen and Paper group because of the negotiation support provided by the interface. They should also perform better on the secondary task, having fewer failed availability negotiations and completing negotiations faster and in fewer exchanges.

**4.5.2 Cognitive Load.** Users in the Negotiator and Pen and Paper groups were under conditions of extreme cognitive load in the experiment. They not only had the primary task of remembering the contents of the video, they also had to schedule various incoming interruptive calls, keeping track of the identity of the caller and the time of each successfully scheduled call as well as the details of any current unresolved negotiations about calls, while taking care to ensure that they did not schedule two calls simultaneously.

We hypothesized two main ways that users could reduce their cognitive load. First, they could delegate responsibility to the experimenter for making the calls either by explicitly asking him to make the call himself “Call me back in 10 minutes” or by implicitly by stating when they were available: “I will be available in 10 minutes”. Delegating responsibility meant that they did not have to remember the identity of the caller but only the time of the call (to avoid scheduling two calls in overlap). Second, they could defer taking the calls for very long periods, for example, until much later in the experiment, rather than taking the call in the next break. This meant that they could focus on extracting information from the video and using the breaks to revise their notes about the videos to ensure that they had fully understood them. This led to two hypotheses:

- Delegation.* We expected that participants who engaged in a delegation strategy (by allocating responsibility for making the calls to the experimenter) would perform better on the primary task of remembering the videos. They would also make fewer errors in scheduling calls in the secondary task.
- Deferral.* We expected that participants who engaged in a deferral strategy (by postponing calls for very long periods rather than taking them in the next break) would perform better on the primary task of remembering the videos. They would also make fewer errors in the secondary task of scheduling calls.

**4.5.3 Prior Experience.** We also expected that users’ prior experience with using their phones in busy situations would help them in the experimental setting.

- Experience.* We expected that people who received more real-life mobile phone calls or had more hours that they were unavailable would be better at dealing with interruptions, performing better on the primary task and making fewer scheduling errors on the secondary task.

## 4.6 Results

**4.6.1 Lightweightness and Support for Negotiation (see Section 4.5.1).** These hypotheses were generally confirmed. Scores on the primary task (memory for video content) were 75% for Negotiator, 59% for Pen and Paper and 78% for Controls. Analysis of Variance (ANOVA) with experimental condition as the independent variable and memory scores as dependent variables showed as expected that there was a main effect of condition ( $F(2,37) = 8.4$ ,  $p < 0.001$ ) with Tukey planned comparisons showing a difference between Negotiator and Pen and Paper ( $p < 0.01$ ) but not between Control and Negotiator ( $p > 0.05$ ).

This analysis takes into account the fact that there were different numbers of subjects in the different groups.

The results suggest that Negotiator users were able to protect the primary task, remaining focused on the video even in the presence of multiple requests for conversation and performing as well as uninterrupted Control users. We also looked at whether primary task performance with the Negotiator degraded over the course of the experiment. Recall that users were presented with 5 different video clips in the experiment. We might expect memory for video content to degrade for later videos as users experienced extended cognitive load. However there were no overall differences in memory scores across tasks on a one way analysis of variance ( $F(4,45) = 0.43$ ,  $p > 0.10$ ), again indicating that the Negotiator offered effective support for handling interruptions.

On the secondary task, negotiating availability, comparisons of Negotiator and Pen and Paper groups showed the following.

- a) Negotiations took significantly fewer exchanges with the Negotiator. An ANOVA comparing the number of exchanges taken to complete each negotiation for each interface indicated that people using the Negotiator took an average of 2.8 exchanges compared with 3.1 for Pen and Paper ( $F(1,18) = 11.2$ ,  $p < 0.001$ ).
- b) The time taken for each individual exchange was significantly shorter with the Negotiator than with Pen and Paper. Negotiator users responded faster to each element of the negotiation, within 9.3s, whereas the Pen and Paper users took 18.6s ( $F(1,18) = 32.3$ ,  $p < 0.0001$ ).
- c) There were no differences between conditions in terms of negotiation success ( $F(1,18) = 0.8$ ,  $p > 0.05$ ). However, the respective success rates were: 91% and 90%, and the lack of difference between the two groups may be the result of a ceiling effect. Indeed, it may have been that the Pen and Paper group were focusing so much on succeeding at negotiating availability, that this distracted them from the primary task.

What can we conclude from this? Our predictions were generally confirmed. For the primary task, the success of the Negotiator is striking: users performed as well as Controls and much better than the Pen and Paper group. They were also generally successful at secondary task negotiations and did this both faster and in fewer exchanges than the Pen and Paper group. The lightweightness, brevity, and success of negotiations are consistent with our observations of natural availability negotiations, suggesting that the Negotiator successfully meets our original design criteria.

**4.6.2 Cognitive Load (Section 4.5.2).** We analyzed the response choices of each experimental user in order to identify for each negotiation the frequency with which they generated a delegation response by asking the experimenter to call them back. However, people who delegated did not make fewer errors on the negotiation task ( $r(18) = 0.39$ ,  $p > 0.05$ ) nor there a difference on the primary task ( $r(18) = 0.17$ ,  $p > 0.05$ ).

We Analyzed the experimental groups' choices of the time that they wanted to talk. There was no evidence that people who deferred (by proposing much

later talk times) performed better on either primary or secondary tasks ( $r(18) = 0.17$ ,  $p > 0.05$  and  $r(18) = 0.06$ ,  $p > 0.05$ ).

**4.6.3 Prior Experience (Section 4.5.3).** We found some evidence that people's real world experience affected their task performance. Those who received more phone calls during the week performed better on the video memory task ( $r(18) = 0.45$ ,  $p < 0.05$ ), although not on negotiation ( $r(18) = 0.05$ ,  $p > 0.05$ ). Neither primary or secondary task performance correlated with users' number of busy hours. Finally, we looked at the relationship between laboratory performance and users' real-world strategies for dealing with interruptions when they were busy. We found that users who in real life turned off their mobile phones when they were busy performed better on the memory task than those who left them on ( $F(1,18) = 2.2$ ,  $p < 0.05$ ), indicating that those who multitasked in real life performed less well on the primary lab task.

Together these results suggest little evidence for the cognitive overload hypothesis but lend some support for prior real-world experience as being a determinant of task performance. Contrary to our expectations, delegation of responsibility for making the call did not reduce users' memory load and allow them to focus more on the videos and deferring conversations did not seem to help with the primary task of remembering the videos or with negotiation. Instead users' performance seemed to relate most closely to their real-world experience and strategies for dealing with calls. These findings are important because they show the validity of our experimental set up in mirroring users' real-world experiences.

**4.6.4 The Cognitive Load Paradox: Why Do Users not Defer or Delegate?** Contrary to our expectations, our statistical results suggest that neither delegation nor deferral helped performance with the lab task. We next explored why this was the case.

To analyze delegation, we first looked at the types of responses that users generated. One major surprise was that users did not generally delegate responsibility for making the call to the experimenter. Instead, on 64% of occasions, they used the "I'll call you back in X minutes." response. Of the remaining (delegational) responses, 29% were "Call me back in X minutes." and 7% were "I am available in X minutes." We also examined each user's most frequent (i.e., modal) response and overall, only one user out of 20 (5% of the users) used a delegational response most frequently. These data are striking because the Negotiator interface was explicitly designed to promote delegation: after each negotiation, the interface defaulted back to the "I am available in X minutes" response and users had to make a deliberate choice to change this.

Why did users take responsibility for calls themselves when this seemed to increase their cognitive load? We examined the interview data to look for reasons for this behavior. The main reason users gave was that it allowed them to remain in control. Since they were somewhat unsure of exactly how busy they were going to be at the agreed time, by taking the responsibility for initiating the call themselves, they were able to ensure that they wouldn't inadvertently be in the middle of another task when that time arrived. One user commented: "Since

I have the opportunity to reply and say when I am available I want to be sure that I am actually available at that time". By assuming responsibility, users feel more confident that they can start the call at a time that is convenient for them. Feeling in control may also be important given the overall complexity of the experimental situation. A second reason people gave for not delegating related to experience with other technologies such as email or voicemail. Here there is an exchange of communication responsibility: the norm is that the initiator makes a first attempt to communicate, and it is then the responsibility of the recipient to respond.

To analyze deferral, we looked at the times that users requested for the conversation. Most users did not defer conversations (which might reduce immediate cognitive load) but preferred to take them as soon as possible, usually during the next break. Recall that users could defer conversations for very long periods (e.g., until after the videos had finished). This longer-term deferral would have been one way to allow them to focus on the video, using the breaks to go over their notes to ensure that they could remember it. Nevertheless, the vast majority of recipient responses involved proposing to hold the conversation within the next few minutes. Fully 47% of responses were requests to have conversations within the next 5 minutes.

Why do users prefer to take calls in the short term when it might be more efficient to defer them? Post-test comments showed that some users believed that they could reduce their memory load by taking conversations as soon as possible. Most felt that it was easier for them to discharge conversations quickly (e.g., by holding them in an upcoming break) rather than trying to postpone them. This way they felt that they avoided stacking up requests and having to schedule, and remember their commitments for multiple conversations. In doing so, they incurred short-term memory costs, while reducing long-term costs.

But a second more important reason seemed to be social, about not risking being interpreted as rude. Most users exploited the system to defer interruptions when they were busy but once the immediate task was done, they felt obliged to take the call as soon as possible. "One advantage with this system is that it allows you to decide when to take a call. You can choose to call somebody back directly when you're done with your work without risking being interrupted all the time with other issues".

Further analysis of negotiation times suggested another interesting phenomenon: users avoided interim scheduling. Although users' preferred strategy (64% of the time) was to hold conversations as soon as possible after completing their current activity, there was some evidence for a long-term deferral strategy with 22% of responses proposing to hold the conversations 20 minutes or more from the contact time. Only a few responses (14%), however, requested conversations in the 5–20 minute interval.

Why did users prefer primarily to schedule immediately, or secondly in the long term, but not during interim periods? This again seemed to relate to control. Users said that they tended to avoid interim times for conversations (i.e., 5–20 minutes) because they weren't at all sure about their availability during this time. Their preference was to take conversations as soon as possible but

if they knew that they weren't available immediately, they preferred to schedule for the long term (when they knew they were definitely available) rather than interim periods. They were less confident that they would be available during these interim times, and they didn't want to run the risk of suggesting an interim time and later find that they were unavailable during that time (either because the video was too complex or they were in the middle of a scheduling negotiation) and then have to renegotiate a time to talk. Overall, they were more confident about both their immediate and long-term availability, leading them to prefer these times. One user drew an analogy with real-life scheduling: "If I'm in a meeting I can't usually tell exactly when that meeting will end. It would be embarrassing if I'd promised to be available in 20 minutes and then couldn't keep my promise. It's much easier to have control over the present time or schedule things to take place much later during the day. You simply need time in between activities and that frame is hard to schedule".

#### 4.6.5 *Feedback on the Negotiator System and Suggestions for Redesign.*

Several users underscored the importance of lightweightness of the interface along with the requirement that it not distract them from their main activity. One user said, "It is extremely important not to lose your focus on the issue at hand [the primary task], if the interface isn't easy to use you kind of flip after 15 minutes". Another user requested a reduction in the number of actions in the negotiation loop. "If you can't respond with just one action, then you kind of lose track of what's happening [in the video]." However, the objective data for both the primary and secondary tasks show that the Negotiator was largely effective and hence met this lightweightness requirement. Users were able to both remember the videos and hold rapid, largely successful negotiations about their availability: "I think that the biggest advantage with this system is that I can rapidly find an appropriate time for both me and the person calling me to talk about things."

Another user pointed out that without the system she would have had to make herself unavailable (the predominant strategy people use in real life). However, she suggested two changes, one was to be able operate the system with her left hand (so she could continue taking notes while handling the incoming call), and secondly that there be a simple defer negotiation button (e.g., "I'll get back to you with a time to talk as soon as I can") that could be pressed when one was extremely busy. Such a button would reduce cognitive load as users wouldn't have to decide on a response type or a time for their response [Nelson et al. 2001; Pering 2002]. Other users suggested that more information about the topic of the call might help with negotiations. We already provide Caller ID, and one extension might be to allow the initiator to provide a brief SMS message about the topic of the call, although this design would run counter to our observations that natural availability negotiations seldom mention the content of the projected conversation. Overall, our interface seems to provide a good trade off between sending prior information about what a conversation will be about (e.g., using free form SMS) versus having little or no information (the current situation with today's mobile phones).

A major problem users experienced was remembering commitments. Recall that these were depicted on the Commitments Board showing specific times. While this seemed to help with scheduling, not all users liked this: “The illustration of commitments made to different people as scrollbars wasn’t a good idea, I don’t think. . . It was just stressful to see these things”. In particular, remembering commitments was a problem for the (admittedly small) group of users who chose to use long-term deferral as their scheduling strategy. One user suggested that, once commitments started to pile up, system defaults should dynamically change so as to reduce the user’s cognitive load. The default response message could change from “I’ll call you back.” to “Call me back.” Another problem with the Negotiator was that commitments are represented relative to the time that they were made, and they did not change visually as time passed. Several users pointed out that the interface should visually represent when a commitment was about to become active. Such a signal would prevent them (a) from overlooking a commitment, or (b) from arranging another call at that time.

Another possible redesign of the user interface might be to have all commitments presented on a single timeline with information about how soon from a given time each will occur. Each new incoming request (“call me in 5 minutes”) would appear as a pending button at a particular point on the timeline, for example, at the 5 minute mark. The user could then decide whether this was an appropriate time by looking at other possibly competing commitments on the timeline. They could then either accept the proposal, (leading the button to change status from pending to committed) or renegotiate it by moving it to another point on the timeline (where its status would still be pending). Another problem for users was that the Negotiator could, on occasions, be distracting. When users were in the middle of a negotiation, they sometimes ended up glancing at the system display instead of the video in order to see whether the scrollbars had changed position and whether they needed to enter a new response.

Despite these reservations and suggestions for redesign, it was clear nevertheless that the Negotiator was superior to Pen and Paper, with many Pen and Paper users experiencing major problems in recording and remembering both the identity and suggested times of the conversations they had arranged with others. Both the objective and interview data show that the Negotiator simplified both of these processes and the overall negotiation. One radical alternative to the Negotiator might be free-form SMS, allowing negotiations to take place in text. However, Pen and Paper users experienced difficulties in handling incoming interruptions manually with pen and paper, a technology they were extremely adept with. Managing interruptions with free-form SMS might be even worse due to the lack of support for rapid data entry on mobile phones. Free, form SMS would also run counter to our initial observation that participants wanted to negotiate when to talk and not what to talk about. Composing and processing SMS messages might also detract from the primary task.

Other users suggested that they would only use the Negotiator in limited work settings (e.g., during project meetings, board meetings, hearings,



seminars). They thought it unlikely that they would negotiate times to talk with their friends or relatives as this would be considered antisocial behavior. They saw the system as primarily allowing them to respond to people that they would otherwise not be able to respond to, for example, to answer a mobile phone during a board meeting. Users were also concerned that other participants in a meeting might consider it rude to be engaged in detailed negotiations (e.g., by obviously manipulating a mobile phone). Of course, the lightwightness and brevity of Negotiator interactions mean the distraction should be relatively minor.

If this interface became widely adopted, it might reduce overall disruption by allowing participants to deal with potentially urgent calls implicitly rather than having to answer the phone during a meeting. This brings up another issue. Some participants felt that the whole process of negotiation might be avoided by just taking the call immediately. They found that waiting for the results of pending negotiations was highly distracting. However, this suggestion ignores the social cost and disruption to others of taking a call in a meeting.

*4.6.6 PocketPC Phone Implementation.* Building on our empirical and design findings, we implemented the Negotiator on a Qteq 1010 smartphone, running PocketPC phone edition.

We applied our three criteria of negotiation, lightwightness, and context in the new design. A primary focus for our redesign was to ensure lightwightness. One observation made during the experiments was that users had to monitor and interact with the different stages of the negotiation and that this was potentially distracting. For a mobile setting, we thought it better to reduce distraction by simplifying the negotiation loop while still allowing users to signal their availability status and context. In the previous version of the Negotiator, the entire negotiation loop was explicitly represented showing incoming calls, responses, and replies to responses. In our mobile version, we opted for a reduced and more implicit representation of negotiation. We supported lightwightness and context by enabling busy users to quickly send a preformulated response message indicating their current availability. Figure 7 indicates how a response message can be sent quickly to a caller.

The default response is preset to “I will call you back in 0h25m”. The user can send that message in a lightwight way by pushing only two buttons: first the red hang-up button (to access the telephone network) and then send (to send the message as a flash SMS) (see Figure 8(A) and 8(B)). Flash SMS bypasses message notification allowing rapid responses to negotiation: the message content pops up on the receiver’s phone display automatically instead of being put in their SMS inbox and announced as “You’ve got 1 new SMS message”. To send a preformulated response message, the user selects his/her choice from the drop-down menu, decides upon a suggested time (see Figure 8(A)), and pushes the send-button (see Figure 8(B)). We also allow users to customize their own messages, for example, “I am extremely busy, try again in: XhXm”. In Figure 8(C), a user has typed in “I am at the movies. . . I will call you back in: XhXm”. To add a customized response message to the list of preformulated alternatives, the user then selects “Add message” from the menu (see Figure 8(D)).

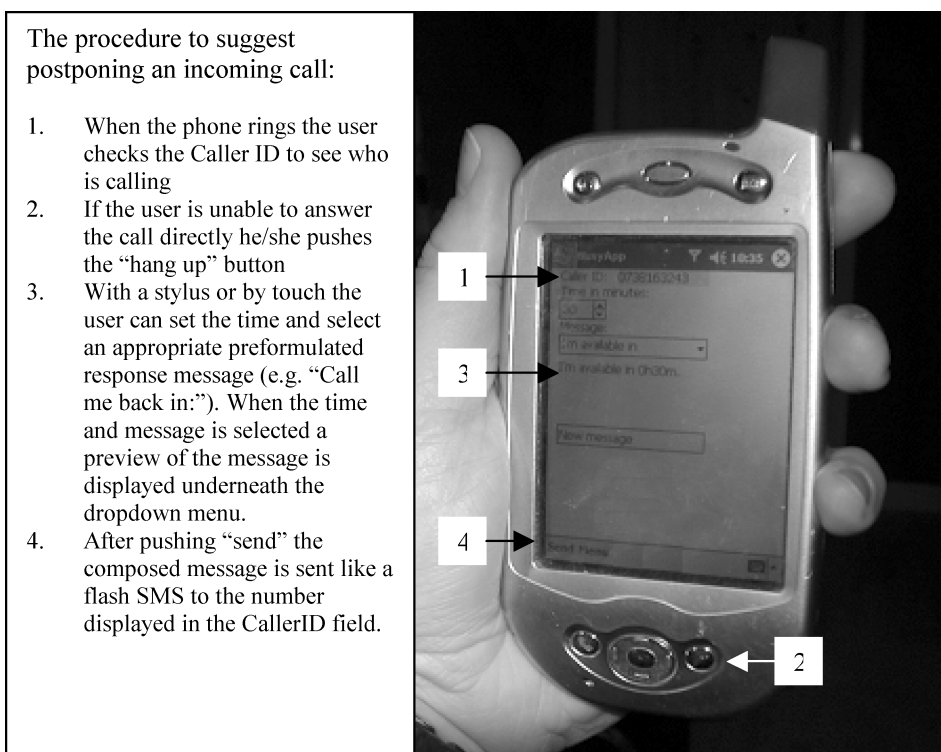


Fig. 7. Implementation of Negotiator on a Qteq 1010.

Further explicit negotiations are not supported in the mobile UI. If a recipient replies “I will call you back in 15 minutes,” and it turns out to be critical for the caller to get in contact immediately, for example, in case of emergency, then we expect the caller to call straight back. Thus, our new and more implicit negotiation loop includes this alternative, a very direct way of negotiating availability based on current practices and expectations.

We have not yet conducted any formal user studies of this version of the system. However, our preliminary experiences in using the mobile version suggest that it is useful. We intend to conduct both detailed experimental studies as well as long-term contextual studies to learn more about designing lightweight support for mobile availability management.

## 5. CONCLUSIONS

We now summarize this exploratory study of availability management.

- (1) We first identified the problems of availability management and dealing with interruptions in work settings, noting that current technologies such as the phone do not provide effective ways for participants to address these problems. We also noted that problems of availability management are likely to increase with the more widespread use of mobile phones and distributed workgroups.

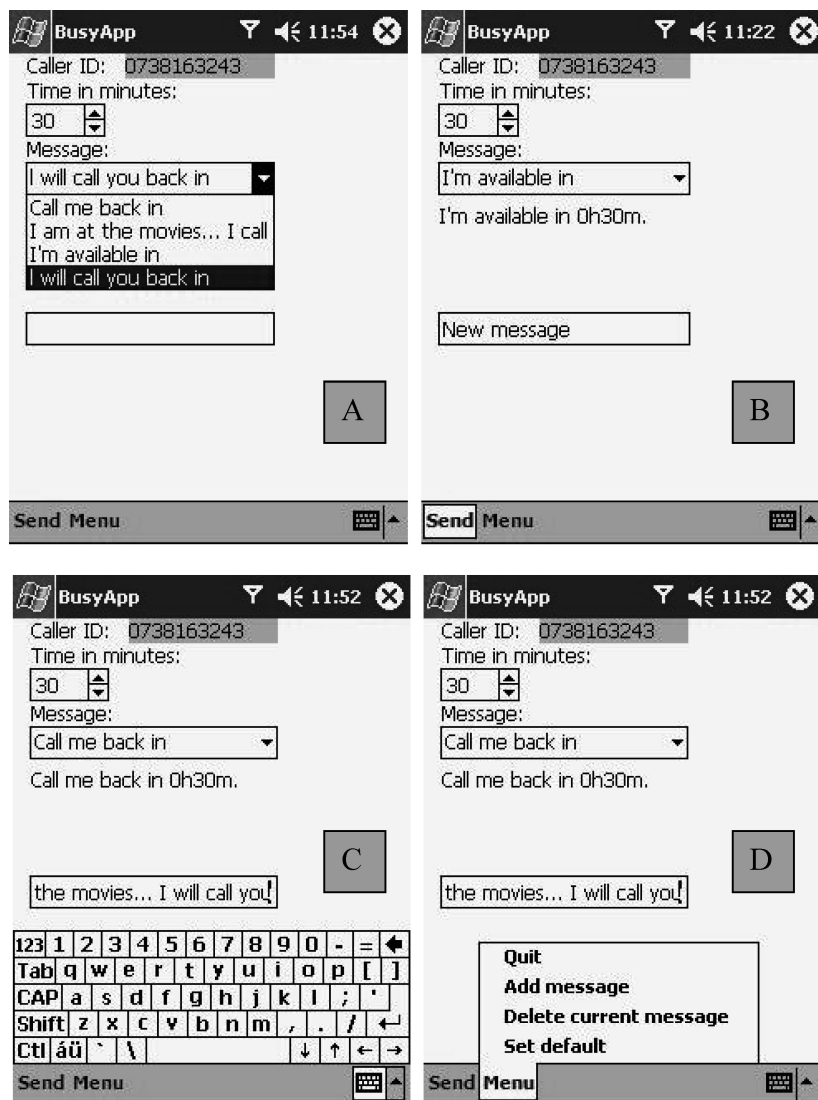


Fig. 8. Screenshots of Negotiator running on the Qteq 1010 PocketPC.

- (2) To understand the problems of availability management, we carried out a study of how interruptions are handled in face-to-face situations. We found four characteristics of availability management, namely, negotiation, dependence on awareness information about the availability of others, brevity, and cognitive effort that are needed to shift from one's current activity to the interruption.
- (3) On the basis of these observations, we designed and implemented a technology, the Negotiator, that addresses these four main design requirements.

- (4) We carried out an experimental study of interruption management using this technology. People were able to use the Negotiator effectively to schedule conversations while focusing on an intellectually demanding activity, indicating the interface satisfied our initial design criteria. Contrary to our expectations, we found that people prefer to take responsibility for returning calls rather than delegating them, and they prefer to schedule calls as soon as possible rather than deferring them until a time when they are definitely free. Interview analysis suggests the cognitive and social reasons why people adopt these strategies.

Our experimental results about deferral are consistent with findings from prior observational workplace studies showing that, in most contexts, users prefer to take interruptions as soon as possible [Hudson et al. 2002; Mintzberg 1973; Sproull 1984]. Users' justification for this strategy was consistent in all these studies, (and with what we found here); people prefer to take interruptions now, incurring the cost of disrupting their current activity in order to avoid the future overhead of having to schedule and remember later commitments to talk. The results for delegation are also consistent with these findings. Here, users felt a social obligation to return calls and a need for being polite rather than delegating them even though it required more effort to do this.

However, there is an alternative interpretation of these findings which focuses on organizational as opposed to individual efficiency [Grudin 1994; Kakihara et al. 2002; Palen 1999]. While dealing immediately with an impromptu interruption may temporarily disrupt a particular user, there may be organizational advantages. For example, a coworker may be prevented from making progress on a critical project task because she needs information that only the recipient can supply. By answering her quickly, the recipient may incur a personal disruption, but the overall project will make more rapid progress. An alternative way of viewing this may be in terms of communication reciprocity [Whittaker et al. 2002b]: being responsive to others' requests even when they are personally disruptive may lead the same people to respond quickly when one has urgent requests of one's own. Future work needs to explore the trade-offs between individual and organizational costs and benefits [Grudin 1994; Kakihara et al. 2002; Palen 1999] for impromptu conversations.

The experimental results also show an interesting pattern for timing impromptu interactions: people prefer either to deal with interruptions immediately or to defer them to the longer term. They do not want to schedule for interim times frames. This seems to be because users have good insights into their current availability (I know whether I am busy right now) and long-term availability (I know that I will be free tomorrow afternoon). In contrast, they are poor at projecting when current activities will finish, making it extremely difficult to schedule communications for the period immediately after their current task. This would suggest utility for technologies that enable users to temporarily defer important conversations until they are no longer busy, for example, by putting these on a call list that later unobtrusively reminds them to return a call.

Another design possibility we have already described is that all interaction requests be placed on a single scheduling timeline along with commitments, allowing users to avoid overlapping commitments. Users could also combine this with information from their current work schedules (e.g., their electronic organizer or smart phone) allowing them to coordinate calls more easily with other activities. Other research on email and voicemail showing that people often fail to reply to outstanding messages suggests a critical need for improved general technologies to support scheduling and tracking of communication commitments [Bellotti et al. 2003; Duchenaut and Bellotti 2001; Mackay 1988; Whittaker et al. 1998, 2000, 2002a, 2002b].

Other future designs might also incorporate more automatic awareness information in order to provide the initiator with more information about the recipient's current activities. For example, the system could automatically detect whether the recipient is already in conversation with another person, using either visual [Vertegaal et al. 2002] or verbal information [Sawhney and Schmandt 1999]. Other simpler techniques might involve access to the recipient's calendar to determine availability [Tang et al. 1994]. One risk of relying exclusively on automatic analyses to reject certain incoming calls is that you might miss emergency calls. This makes it imperative for the recipient to actively decide whether to take each call. Nonetheless, automatically-derived activity information might be useful in a different way. It could allow initiators to confine their conversational requests to times that are less disruptive to the recipient in the same way that information about a recipient's keyboard activity promotes more effective Instant Messaging [Isaacs et al. 2002a, 2002b; Tang et al. 2001a, 2001b]. Alternative future support for negotiation might be to develop richer and more flexible sets of preprogrammed responses to incoming calls, for example, "I am in a meeting. I could talk in 5 mins." "I am on another call. I will call back in 10 mins.", an approach that has been taken by Nelson et al. [2001] and Pering [2002] and which we have implemented in our mobile application. One potential disadvantage with this approach, however, is that the more flexible the set of preprogrammed responses, the more complex it is to decide among these when users are already cognitively overloaded. This implies that the response set needs to be carefully designed and evaluated to avoid its complexity detracting from the primary task. The mobile phone version of the Negotiator is preset to the three different types of response messages used in the experiment. It allows users to add personal responses to this basic list if they feel they can manage extra complexity.

There is also a need for more empirical research on availability management. There are clearly many different types of work that people carry out in offices. It would be useful to have more observational data from different situations, workplaces, and types of workers. Here we only looked at academic settings. We also need additional controlled experiments. In our laboratory study, all initiation requests were of equal status. Furthermore, users were always completely focused on their primary task, whereas in reality they may often be engaged in relatively unimportant and hence interruptible tasks. In future experiments, we plan to investigate how people manage their availability when requests are of different status and when the primary task is interruptible. Other situations

where availability management is likely to be very different are when participants are communicating in order to socialize or to have fun [Nardi et al. 2000]. Finally, we need evaluation of different prototype technologies in real working situations.

We also need to have more empirical research that enriches our understanding of different types of interruptions, and the contexts in which they occur. Some issues relate to anticipation, that is, an expected call should be less disruptive than a completely unanticipated one. Other issues concern one's primary activity: interruptions received during a work break may be less disruptive (although possibly more annoying) than calls received during primary work time. Finally, there are questions about setting. Are impromptu conversations in an open plan office construed as interruptive in the same way that external calls are?

There are important general theoretical implications to our results. As other studies have shown, there is a need to develop new communication models to account for various unique features of technologically mediated conversations. One underresearched characteristic of mediated communications is that these often involve multiple concurrent communication threads. This gives rise to major problems in initiating new interactions [Hudson et al. 2002; Nardi et al. 2000; Nelson et al. 2001; Isaacs et al. 2002; Tang et al. 2002b; Wiberg 2001], task management [Ducheneaut and Bellotti 2001; Whittaker and Sidner 1996; Whittaker et al. 1997, 1998, 2000, 2002a, 2002b], and preserving the context of existing interactions [Ducheneaut and Bellotti 2001; O'Conaill and Frohlich 1995; Whittaker et al. 1997; Wiberg 2001]. Current communication theories tend to focus on the internal characteristics of established interactions rather than the problems of initiating and maintaining multithreaded interactions that we have highlighted here. We therefore need both more investigation of these phenomena and more work developing theories to account for them.

#### ACKNOWLEDGMENTS

Thanks are due to Victor Kaptelinin for valuable comments on earlier versions of this article. Thanks are also due to Tomas Lundqvist for the co-development of the Negotiator system. Finally, thanks are due to the participants in the user test and all the members of Interaction Theory Lab for valuable discussions throughout this project.

#### REFERENCES

- BELLOTTI, V. AND BLY, S. 1996. Walking away from the desktop computer: Distributed collaboration and mobility in a product design team. In *Proceedings of ACM Conference on Computer Supported Cooperative Work*. Cambridge, MA. ACM Press, 209–218.
- BELLOTTI, V., DUCHENEAUT, N., HOWARD, M. A., AND SMITH, I. E. 2003. Taking email to task: The design and evaluation of a task management centered email tool. *ACM Conference on Human Factors in Computing Systems (CHI'03)*, 345–352.
- BLY, S., HARRISON, S., AND IRWIN, S. 1993. Media spaces: Bringing people together in a video, audio and computing environment. *Comm. ACM* 36, 28–45.
- BROWN, B. AND PERRY, M. 2000. Why don't telephones have off switches? Understanding the use of everyday technologies; A research note. *Interact. Comput.* 12, 6, 623–634.

- DAVENPORT, H. AND BECK, J. 2001. *The Attention Economy: Understanding the New Currency of Business*. Harvard Business School Press.
- DOURISH, P. AND BELLOTTI, V. 1992. Awareness and coordination in shared workspaces. *Conference Proceedings on Computer-Supported Cooperative Work*. 107–114.
- DUCHENAUT, N. AND BELLOTTI, V. 2001. Email as habitat. *Interactions*. ACM Press, New York, 8, 5, 30–38.
- FISH, R., KRAUT, R., ROOT, R., AND RICE, R. 1992. Evaluating video as a technology for informal communication. In P. Bauersfeld, J. Bennett, and G. Lynch Eds. *Human Factors in Computing Systems (CHI'92)*. 37–48.
- GRUDIN, J. 1994. Groupware and social dynamics: Eight challenges for developers. *Comm. ACM* 37, 1 (Jan.).
- HUDSON, J., CHRISTENSEN, J., KELLOGG, W., AND ERICKSON, T. 2002. I'd be overwhelmed, but it's just one more thing to do. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: Changing Our World, Changing Ourselves* (April).
- ISAACS, E., WALENDOWSKI, A., AND RANGANTHAN, D. 2002a. Hubbub: A sound-enhanced mobile instant messenger that supports awareness and opportunistic interactions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (April).
- ISAACS, E., WALENDOWSKI, A., WHITTAKER, S., SCHIANO, D., AND KAMM, C. 2002b. The character, functions, and styles of instant messaging in the workplace. In *Proceedings of the Conference on Computer-Supported Co-operative Work*.
- KAKIHARA, M., SØRENSEN, C., AND WIBERG, M. 2002. Fluid interaction in mobile work practices. Tokyo Mobile Roundtable (May) Tokyo, Japan.
- KRAUT, R., FISH, R., ROUTH, R., AND CHALFONTE, B. 1993. Informal communication in organizations: form, function and technology. In *Groupware and Computer-Supported Co-operative Work*. R. Baecker, Ed., Morgan Kaufmann, 287–314.
- KRISTOFFERSEN, S. AND LJUNGBERG, F. 1999. An empirical study of how people establish interaction: Implications for CSCW session management models. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI'99)*. 1–8.
- MACKAY, W. 1988. More than just a communication system: Diversity in the use of electronic mail. In *Proceedings of Computer Supported Cooperative Work*. 344–353.
- MILEWSKI, A. AND SMITH, T. 2000. Providing presence cues to telephone users. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work* (Dec.).
- MINTZBERG, H. 1973. *The Nature of Managerial Work*. Harvard University Press.
- NARDI, B., WHITTAKER, S., AND BRADNER, E. 2000. Interaction and Outeraction: Instant Messaging in action. In *Proceedings of Conference on Computer Supported Cooperative Work*. 79–88.
- NELSON, L., BLY, S., AND SOKOLER, T. 2001. Quiet calls: Talking silently on mobile phones. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM Press.
- O'CONNAILL, B. AND FRÖHLICH, D. 1995. Timespace in the workplace: Dealing with interruptions. In *Proceedings of Human Factors in Computing Systems (CHI'95)*. ACM Press, 262–263.
- PALEN, L. 1999. Social, individual and technological issues for groupware calendaring systems. In *Proceedings of Human Factors in Computing Systems (CHI'99)*. ACM Press, 17–24.
- PANKO, R. R. 1992. Managerial communication patterns. *J. Organizat. Comput.* 2, 1, 95–122.
- PERING, C. 2002. Taming of the ring: Context specific social mediation for communication devices. In *Conference Extended Abstracts on Human Factors in Computer Systems* (April).
- ROUNCFIELD, M., VILLER, S., HUGHES, J., AND RODDEN, T. 1995. Working with constant interruption: CSCW and the small office. *The Informat. Soc.* 11, 4, 173–188.
- ROSSON, M. B. AND CARROLL, J. M. 2002. *Usability Engineering: Scenario-Based Development of Human-Computer Interaction*. Morgan Kaufmann, San Francisco.
- SAWHNEY, N. AND SCHMANDT, C. 1999. Nomadic Radio: Scalable and contextual notification for wearable audio messaging. *Proceedings of CHI'99*.
- SCHEGLOFF, E. 1968. Sequencing in conversational openings. *American Anthropologist*. 70, 1075–1095.
- SOLSO, R., JOHNSON, H., AND BEAL, K. 1998. *Experimental Psychology: A Case Approach*, 6th Ed. Addison Wesley Longman Inc.
- SPOULL, L. 1984. The nature of managerial attention. *Advances Inform. Process. Organizat.* 1, 9–27.

- TANG, J., ISAACS, E., AND RUA, M. 1994. Supporting distributed groups with a montage of lightweight interactions. In *Proceedings of Conference on Computer Supported Cooperative Work (CSCW'94)*. ACM Press, 23–34.
- TANG, J., YANKOLOVICH, N., BEGOLE, J. B., VAN KLEEK, M., LI, F., AND BHALODIA, J. 2001b. ConNexus to awarenex: Extending awareness to mobile users. *Human Factors in Computer Systems (CHI'01)*, 3, 1, 221–228.
- VERTEGAAL, R., DICKIE, C., SOHN, C., AND FLICKNER, M. 2002. Designing attentive cell phone using wearable eye-contact sensors, In *Conference Extended Abstracts on Human Factors in Computer Systems* (April).
- WHITTAKER, S. 2002. Theories and methods in mediated communication. In *The Handbook of Discourse Processes*. Graesser, A., Gernsbacher, M., and Goldman, S. Ed. Erlbaum.
- WHITTAKER, S. AND SIDNER, C. 1996. Email overload: Exploring personal information management of email. In *Proceedings of Conference on Computer Human Interaction (CHI'96)*. ACM Press, 276–283.
- WHITTAKER, S., DAVIS, R., HIRSCHBERG, J., AND MULLER, U. 2000. Jotmail: A voicemail interface that enables you to see what was said. In *Proceedings of Conference on Human Computer Interaction (CHI'00)*. ACM Press, 89–96.
- WHITTAKER, S., FROHLICH, D., AND DALY-JONES, O. 1994. “Informal workplace communication: What is it like and how might we support it?” In *Proceedings of ACM Conference on Human Factors in Computing Systems*.
- WHITTAKER, S., HIRSCHBERG, J., AMENTO, B., STARK, L., BACCHIANI, M., ISENHOUR, P., STEAD, L., ZAMCHICK, G., AND ROSENBERG, A. 2002a. SCANMail: A voicemail interface that makes speech browsable, readable and searchable. In *Proceedings of Conference on Human Computer Interaction (CHI'02)*. ACM Press.
- WHITTAKER, S., HIRSCHBERG, J., AND NAKATANI, C. H. 1998. All talk and all action: Strategies for managing voicemail messages. In *Proceedings of Conference on Computer Human Interaction (CHI'98)*. ACM Press.
- WHITTAKER, S., JONES, Q., AND TERVEEN, L. 2002b. Managing long term conversations: Conversation and contact management. In *HICCS'02*.
- WHITTAKER, S., SWANSON, J., KUCAN, J., AND SIDNER, C. 1997T. eleNotes managing lightweight interactions in the desktop. *ACM Trans. Comput.-Hum. Interact.* 4, 2, 137–168.
- WIBERG, M. AND LJUNGBERG, F. 2001. Exploring the vision of “anytime, anywhere” in the context of mobile work. *Knowledge Management and Business Model Innovation*. The Biztech Network, Brint Press.
- WIBERG, M. 2001. In between mobile meetings: Exploring seamless ongoing interaction support for mobile CSCW. PhD thesis, Department of Informatics, Umeå University, Sweden.

Received December 2003; revised October 2004; accepted March 2005 by Susanne Bodker