

Cover Page

Title of submission: Meeting in Quiet: Choosing Suitable Notification Modalities for Mobile Phones

Category of submission: Design Research Study

Name and full contact address (surface, fax, email) of the individual responsible for submitting and receiving inquiries about the submission: Daniel Fallman, Umea Institute of Design, Umea University, SE-90187, Umea, Sweden, +46 90 786 6697, daniel.fallman@dh.umu.se

Meeting in Quiet: Choosing Suitable Notification Modalities for Mobile Phones

Daniel Fallman

Umea Institute of Design
Umea University
Umea, SE 90187 SWEDEN
daniel.fallman@dh.umu.se

Björn Yttergren

Umea Institute of Design
Umea University
Umea, SE 90187 SWEDEN

Abstract

This project starts out from one of those awkward moments when one's mobile phone sets off with a loud noise in situations one would definitely not want it to; e.g. at the opera or in an important meeting. The design challenge is how a mobile phone, as a computer device that appears in several different human contexts, could come to adapt its behavior to the situation. Is there a way for a mobile phone to know, or at least have a clue, which notification modality would be the most appropriate?

Our prototype design detects if a user is within a certain distance from another user. If so, the two users' mobile phones automatically set up an ad-hoc network and exchange information about notification modality. It makes use of a near-field communication technology using magnetic fields rather than traditional radio waves.

Keywords

Mobile Phones, Notification Modality, Context Aware, Mobility, Design, Near-field Communication, Magnetic Link

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Copyright © 2005 AIGA | The professional association for design.

Introduction

Every once in a while, one forgets to switch off the notification sound of one's mobile phone when in a meeting, a seminar, at the opera, or just simply involved in a face-to-face act of communication. These are all examples of situations in which a loud noise is just not the right way for the cell phone to notify its user that someone wants to communicate. This project originates from one such awkward moment; with the design challenge of how a mobile phone, as a computer device that appears in several different human contexts, could come to adapt its behavior accordingly. Is there a way for a mobile phone to "know", or at least have a clue, when it is not appropriate to loudly buzz?

Today, it is the sole responsibility of the owner of the mobile phone to make sure the notification modality is properly configured for the situation in which both the device and the human user is located. While this provides the owner of the device full control, it seems that many users forget about or neglect the configuration possibility and simply leave the sound switched on at all times. Even though some people may consider a constantly ringing mobile phone as a part of their lifestyle, a part of what they are—others simply forget about the configuration because they are involved in other, more important, activities.

Quite often, the mobile phones that go off at meetings, funerals, marriage counseling, Tai Chi sessions, and at the opera are not only annoying to other attendees—but also, and perhaps mostly so, to the unlucky owner of the mobile phone, the one having to quickly look for his or her phone in bags and pockets, pushing the buttons frantically. It is not uncommon that such sudden and unexpected outbreaks put the current

social occasion to a halt and move all attendees' involvements away from what was previously dealt with to the annoying sound coming from somewhere in the vicinity. The meeting comes to a stop, perhaps never to recover completely; people move from being involved in a beautiful opera aria to being annoyed by a disturbingly weird ring tone.

With examples like these in mind, there seems to be a need for mobile information technology—that unlike traditional interactive systems such as stationary phones and desktop computers are brought into many different social contexts [7, 8]—to pay attention to what is going on around them.

Purpose and Scope

Through a number of cycles of scenario-based design—some philosophical, some practical—the aim of this project has been to explore different approaches for assisting users in choosing suitable notification modalities for their mobile phones. The scope of this part of the project has been to creatively explore a wide range of alternative conceptualizations as to how users could be assisted in choosing suitable notification modalities for their mobile phones. As a final step in the design process, a prototype system of one such assistant technology has been implemented.

In this paper, we will first present a short overview of the literature in the area of context-aware computing. Second, we will present some of the scenarios that we have developed for assisting users in choosing a suitable notification modality. Third, we will introduce and discuss in detail our prototype implementation, where mobile phones are equipped with technology that uses a magnetic field to determine distance to

other users and which sets up an ad-hoc network for phone-to-phone communication.

Mobility, Context, and Involvement

The current understanding of the concept of mobility in Human-Computer Interaction (HCI) and Computer-Supported Collaborative Work (CSCW) literature seems to pre-assume that mobility is more or less solely about freeing people from geographical constraints [12]. To come to terms with this limiting view, some authors have argued that the concept of mobility must be enriched by also including the contextual interactions people establish and the social character of that kind of interaction [7, 8, 10, 12, 16].

According to such a view, mobility can be regarded in terms of three related dimensions: spatial mobility, temporal mobility, and contextual mobility. Spatial mobility is clearly the most common way of conceptualizing mobility in the HCI field, where the term mainly denotes corporeal geographical freedom. Here, a mobile device such as a mobile phone is an IT product occupying some defined physical space in a particular geographic location. In our project, we would ask in relation to this understanding of mobility: Are there spaces which call for certain notification modalities?

If spatial mobility mainly concerns questions of 'where,' then temporal mobility answers to questions of 'when,' i.e. considering the dimension of time. Here, a mobile phone becomes an IT product that is used at a particular point in time. In relation to our project, are there temporal aspects that need to be taken into consideration when assisting users in choosing

notification modalities? Do the modalities vary throughout the day?

From the point of view of a contextual understanding of mobility, a mobile phone would be an IT product appearing within a specific loop of human activity. We ask: how can a particular human activity or social occasion come to influence or determine the choice of notification modality?

The importance of context to IT products is not unique to research in mobile devices. It has already been reflected in various ways in interactive systems research related to HCI. Some of the fields which have been especially concerned with contextuality are ubiquitous and pervasive computing [13, 15]; tangible computing [11]; social computing [5]; wearable computing [2]; augmented reality [9]; and obviously the previous efforts in Artificial Intelligence [6].

Several attempts have been made to prototype context-sensitive devices which display some kind of awareness of their location, their physical context, other devices in the spatial and temporal vicinity, and who their user is [3], e.g. the Xerox PARCTab [13], a system aware of its location from which some simple but allegedly valuable services are provided. While the focus of these efforts so far has been that of designing specific devices and applications, a more theoretical understanding of the role of context—and thus involvement—in mobility has however not been presented [4]. Hence, aside from spatial and temporal understandings of mobility, there are only a few attempts at providing frameworks for thinking about mobile use of interactive systems based on issues of context [8, 12]. We have previously argued that any

notion of mobility applied in mobile HCI must not come to adopt a naïve Cartesian model where mobility is solely thought of in temporal and spatial terms. While these are the dimensions highlighted by pervasive rhetorical figures such as ‘anytime, anywhere’, we believe it is important that we do not conceal or forget the situatedness and involvement of human action where the physical and social setting, the context, of human action is of highest importance. But as what we generally mean by ‘context’ is made up of a complex weave of things like artifacts, meaning, and language, it has historically been intrinsically difficult to grasp for any computational system [6, 16].

The Process

With the idea in mind that there is also a social and involving character to mobility—not just a matter of corporeal freedom—we set out to develop a range of alternative use models. With the term ‘use model’ we mean the basic notion or idea for how human-computer interaction should proceed, in this case for assisting users in choosing suitable notification modalities for their mobile phones.

Scenario Design

Below, we present some of these alternative conceptualizations:

Individual Choice

The use model currently applied by more or less all mobile phones on the market, as discussed above, provides the individual user with full control and responsibility for choosing notification modality. Each user has to decide whether the current situation is one in which sound is acceptable (by his or her standards), or if it is an occasion where it is not. In the latter case,

the user has to consciously turn the sound off each time.

While this provides the owner of the device with complete control, it appears as if users often forget about or simply neglect the configuration possibility and leave the sound switched on at all times.

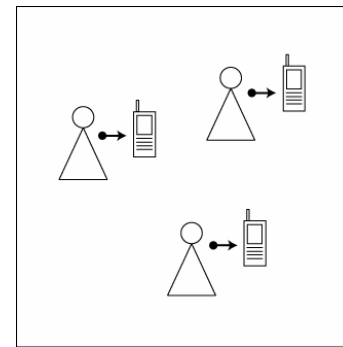


Figure 1: The individual choice use model

Everyone Else's Choice

One simple but yet quite thought-provoking alternative would simply be to switch control away from the user. What if everyone else using their phones could decide about your mobile phone's notification modality, while you had no say yourself? You could then control everyone else's setting, but have no say what-so-ever regarding your own mobile phone's settings.

While it appears highly unlikely that this use model would ever become commercially available, it works well as a thought-provoking experiment. Unlike the current situation, it takes the responsibility away from the individual user. If one's phone would set off during a meeting for instance, the owner of the device would

not have to feel ashamed for forgetting about switching the sound off. Quite the opposite, she could complain to the meeting crowd that none of them had been so kind as to put her phone to a more suitable meeting modality!

Hence, with this model the responsibility changes from individual user to the group of people that are involved in some shared human business. Whether a certain occasion should not be disturbed by loud ringing tones, or if in fact it could here becomes a social decision rather than an individual choice. Hence, not only is the actual activity of changing the setting moved away from the individual user, but—perhaps more important—so is the decision about whether or not certain situations call for certain behaviors.




Figure 2: The everyone else's choice use model

Location

Another use model, with a Cartesian heritage, is to use geographical location as a sign of certain types of human behavior [10]. Here, beacons or some kind of positioning system could be used to charge specific rooms or locations with computational meaning. A

beacon could inform one's mobile phone—without effort on the part of its user—that the room it just entered is a meeting room, and the phone could automatically change its setting. Likewise, when the user (and the device) leaves such a location, it would automatically switch back to full sound mode again.

While fairly easy to implement, this use model has some inherent weaknesses, most of which have to do with its naivety. The argument is that social occasions which call for certain human behaviors are not usually definable by geographical positions alone, even if it would be the most convenient for an interactive system. This 'human perspective' is in this sense anything but mysterious or something requiring continental philosophers to be involved in the process—it is simply to say that entering an empty meeting room is not the same as entering a meeting.

Also, this use model would not deal with social occasions that are completely disconnected from a geographical position. Typical ad-hoc meetings, if you for instance would bump into a couple of friends in a café, would not be recognized by such a system.




Figure 3: The location-based use model

Time

Another idea springing out of a Cartesian way of thinking would be a use model which would rely on a specific position in synchronous time, or—in other words—the time of day. Many mobile phones today allow their users to synchronize their desktop computer calendar with that of the mobile phone. The idea for this use model would be to have the phone act according to this built-in calendar. Hence, if a user in her calendar has a meeting scheduled between 10:00 AM and 11:00 AM, the phone automatically switches off the sound notification at 10:00 AM and back again at 11:00 AM.

Similarly to the location-dependent use model, a time-based system would be easy to implement, but suffer some important drawbacks. First, it would require the user to actively provide information to the calendar system about all meetings and other events where sound would not be appropriate, i.e. 'pre-program' the system. Second, a time-based pre-programmed system would not be able to respond to ad-hoc meetings or other sudden changes in social occasion that call for a

certain behavior from a human point of view (such as a quickly arising argument with one's fiancée about some issue of apparent importance to her).

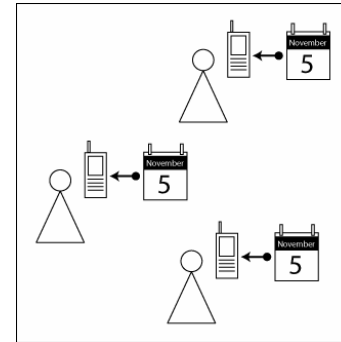


Figure 4: The time-driven use model

Combination Model

To make the two Cartesian systems more reliable, they can of course easily be combined without much added computational complexity. Hence, a system which cares both for geographical location (you are in a meeting room) as well as keeping track of time (you also have a meeting now in your calendar) provides a greater deal of certainty to the system as to decide to switch from sound to some other notification modality more suitable to meetings.

If you stop by the same meeting room later that day at say 3:00 AM, and you do not have a meeting in your calendar, then the system could decide not to change your settings. While a combination system would make the system more reliable, at least when it comes to guessing when not to change over to quiet mode, it would nevertheless still be unable to cope with ad-hoc meetings and all other sudden changes in social

occasion that are not pre-planned or tied to a specific physical location known in advance to host such human behavior.




Figure 5: A use model based on the combination of time and location

Involvement

But how could such occasions be dealt with from a computational perspective? Or can they at all? While there are no definite limits as to what complex computer systems can simulate, it seems safe to argue that computers cannot ‘understand’ a certain situation in the same sense as (most) human beings would understand the same situation, and act accordingly [6]. While approaches to develop such ‘thinking systems’ have been explored previously, perhaps most prominently so within old-school artificial intelligence—now largely abandoned, computational system cannot understand the complex patterns of cultural and social rules and conventions that make up a specific social occasion [6, 8, 10, 14].

But if we let go of the idea of having a mobile phone completely understand a human situation, is there a

way in which it can get a clue about what is going on around it and what its user is involved in (or not involved in)?

This notion was the starting point for the involvement use model. It can be used as an add-on to the combination model presented above, particularly capable of dealing with ad-hoc meetings.

The idea here is to use physical proximity to other users as the clue for the system to recognize involvement. If a user is within a certain distance from another user, the two users’ mobile phones set up an ad-hoc network between them and exchange information. In its simplest form, close proximity to another user would tell both users’ phones to switch to silent mode.

If one of the user’s phone is in silent mode for some reason, the other adopts the same behavior. If several users are in the close vicinity, their phones are also ‘infected’. Hence, in this scenario, appropriate notification modality would spread like a virus from phone to phone. Without requiring specific user activity, phones would automatically go to silent mode if their users would enter into close proximity of other users, and likewise, if a user’s distance to the group expands (e.g. leaving the meeting room, saying goodbye on the street and walks away, etc.) the system would go back to sound notification.

There are obviously a lot of real-life situations where this basic notion fails to operate properly, e.g. on the platform waiting for a train with others, where you in fact would want to hear your phone. Another such situation would be if you are expecting a very urgent

call, or if there is a person in need, etc. But the argument is that these use models need to be regarded as basic starting points for discussion, not as blueprints for implementation.

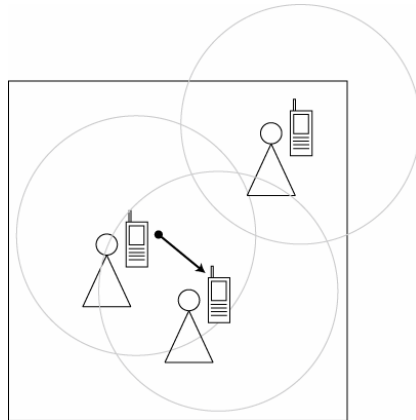


Figure 6: The basic involvement use model

In a more advanced form, it would be possible for this use model to interact with the combination model (location + time) to deal specifically with ad-hoc meetings and other kinds of sudden changes in social occasion that are not pre-planned or tied to a specific physical location.




Figure 7: The involvement use model in combination with the location + time use model

Design Solution

From these basic use models, we choose to look into the idea of *involvement* in more detail and provide a prototype implementation of such a system. How could this use model actually be implemented using the technology we have available today?

The approach we have chosen to quickly prototype and achieve a functional system—to show the technical feasibility of this use model as well as be able to perform user testing—consists of a number of different parts:

Communication between devices

First, we designed and developed a custom-made circuit which forms a transceiver that uses magnetic link technology to communicate with other, similar transceivers, using a custom-made protocol.

A key technical question here was what means of communication should the system apply. While most wireless communication is accomplished by propagating an RF plane wave through free space—ordinary radio waves basically—we chose to make use of a related but in some aspects very different kind of technology: near-field magnetic link. Unlike ordinary radio waves, this solution rather makes use of a non-propagating ‘quasi-static’ magnetic field, the physics of which is described in detail by Aura Communications Technology, Inc. [1]. Here, a modulated magnetic field is generated by a transducer element, where the magnetic field remains relatively concentrated on a small area around the transmitting device. The power in ordinary radio waves, i.e. a propagating RF wave, in the far field rolls-off as one over the distance from the source squared ($1/r^2$). This should be compared to a magnetic field, such as the one generated in our prototype, where the roll-off is $1/r^6$ [1].

For most typical wireless applications such a hefty roll-off would be a considerable disadvantage on the part of magnetic communication systems, because of its limited range. When a large range is needed, propagating RF communication remains the primary choice. In close proximity systems however, a sharp roll-off behavior is on the contrary a substantial advantage. The strong attenuation over distance sets up a ‘cloud’ or a ‘bubble’ of around 150 centimeter which is highly predictable and—again different from typical RF waves—relatively unaffected by the physical surroundings [1]. While plane waves are significantly affected by their surroundings, e.g. attenuated by the human body, the ‘bubble’ set up by a magnetic field quite reliably follow the $1/r^6$ behavior—regardless of

metal objects, conductive materials, people, etc.—and the magnetic field passes relatively unimpeded.

Yet another benefit of the magnetic link technology in close proximity systems is that its very predictable roll-off behavior allows several closely located users to transmit and receive on the same frequency with no interference [1]. This character makes it an even more appropriate technology for our specific purposes, as we would typically want to allow more than two users to be located simultaneously within the same ‘bubble’.

In our implementation, the ‘bubble’ around each transceiver is around 150 centimeters. When a similar transceiver enters into such a cloud (i.e. two users of the system are within 5 feet of each other), the two transceivers sets up an ad-hoc network using the magnetic field as the communications link and a custom-made protocol over which data is sent. The sharp roll-off behavior, as discussed previously, makes the bubble’s edges highly accurate (within 1-2 cm).




Figure 8. A picture of the custom-made circuitry containing the transceiver device

Link to the Mobile Phone

To allow each transceiver to communicate with a mobile phone, we have used QTEC 2020 phones with Microsoft PocketPC Phone Edition. Each phone is equipped with the custom circuitry in a small auxiliary box, connected to the phone through a standard RS232/USB link.

Software on Mobile Phone

Custom-made software runs in the background on the mobile phones and translates commands coming in from the box and changes the phone's notification modality settings accordingly.

Prototype Limitations

The current prototype system only implements the involvement use model (see 'Scenario Design' above) as introduced above for three mobile phones. Hence, the prototype system neither takes into account the user's geographical position (the location model) nor

does it look into the user's calendar (the time model). Primarily, limited project funds and a short time frame have prevented us from making a more complete system (add the combination model on top of the involvement model). As the combination model, from a computational perspective, would be a fairly straightforward implementation, we have focused the design process to center on the more unexplored area that the involvement model constitutes.

Discussion and Conclusion

As a point of departure, this project took those awkward moments when one's mobile phone sets off with a loud noise in situations one would definitely not want it to; e.g. at the opera, in an important meeting, or just when involved in a delicate argument with a loved one.

The design challenge has been how a mobile phone, as a computer device that appears in several different human contexts such as these above, could come to adapt its behavior to the situation? Is there a way for a mobile phone to "know", or at least have a clue, when it is not appropriate to loudly buzz?

To deal with these issues, we have taken a step back and introduced, analyzed, and discussed the way in which the current understanding of the concept of mobility in HCI and CSCW literature seems to pre-assume that mobility is more or less solely about freeing people from geographical constraints. Our argument has been that doing so may lead to design solutions that are negligent to a social perspective, to what we as human users are involved with and in. Our argument is that this is generally far more important to

us as people than our current location in space and time.

To try to break out of the Cartesian perspective, not just in our argumentation but also when it comes to actual design, we introduced another way of thinking about mobility; a contextual understanding of mobility that stresses that mobile information technology devices such as mobile phones are products that first and foremost appear within a specific loop of human activity—and only in a secondary sense hold a specific geographical locations at a specific time. While the latter properties are much simpler to deal with from a computational perspective, we have tried to explore the contextual model also through prototype design in order to find out if it also is a computationally sound perspective.

Our strategy when it comes to interactive systems design—i.e. when social and contextual notions need to take a computational shape—has been to use physical proximity to other users as the clue for the computational system to recognize human involvement. While this obviously does not provide the computational system with a human type of ‘understanding’ of social occasions, it at least provides a clue—clearly limited but still a clue—to what its user might be involved in.

Our prototype design detects if a user is within a certain distance from another user. If so, the two users’ mobile phones automatically set up an ad-hoc network in between them and exchange information about notification modality. If this is combined with various kinds of more traditional input, e.g. geographical location (such as in a meeting room) and time-based information (such as the user having a meeting in her

calendar at the same time), this system could be fairly reliable in assisting the user in choosing a suitable notification modality for various occasions.

The concept of human involvement to guide a computational system, as described by Fallman [7, 8], is clearly a far more complex matter than simply measuring the distance between two points. The accuracy of the system would benefit from more contextual input, such as what kind of activity in which the user is involved. Is it primarily physical (e.g. chopping wood) or mental (e.g. mental arithmetic)? Is the user sitting or standing? Is the user talking? What is the level of ambient light? Is it a loud or quiet environment? Which users are facing each other? This list can be made much longer, but the point is made. Input like this could contribute to a more accurate system, but the tradeoff would be an immense increase in the level of complexity—and the system would lose some of its simplistic beauty. Maybe the strength of this system lay not in its accuracy, but in fact in its simplicity?

The simple implementation, the cheap hardware and thin software which is behind this system also makes it seemingly more realizable to implement commercially. A more accurate system, tracking for instance eye-gaze, altitude, temperature, lighting conditions, and sound level, would also imply a more expensive as well as a far more unwieldy solution.

Despite the relative lightness of the system, some readers have probably already started questioning the possibility of realizing this system in a commercial product. Market wise, there seems to be many obstacles for a system like ours to take off. For

instance, a couple of new standards in technology and communication protocols would need to be agreed upon; all new devices would have to incorporate this extra functionality, increasing production cost; old phones still in circulation and use would undermine the whole system; and so on.

This study has not dealt primarily with these issues, despite their obvious importance. We have rather focused on exploring alternatives to the currently monopolized use model when it comes to notification modalities for mobile phones. In doing so, we have nevertheless come to deal with one of the key questions also when it comes to commercial potential: Would users be ready to give up some control over their own devices to be part of a larger system, where there is a chance that the group as a whole works better in some respect than single users by themselves? This question seems to have relevance for many of the established ideas and ways of working within interaction design, human-computer interaction, and computer-supported collaborative work. Stretching it, the question even has a political dimension.

In light of this discussion, it is important to realize that the goal of the design process presented in this paper has not primarily been to generate a product which fits and would make a smooth appearance on the current market. Quite the opposite, actually. We believe that it is one of the key strengths of design research that it has the power to come up with tangible artifacts and functional demos that come to act as demonstrations of new concepts and ideas that go a bit further than the drawing table. This kind of design research makes the possible visible and tangible. It shows people how things could be designed; how things could look,

sound, feel, and work. It presents people with alternatives.

Future Research

By doing so, i.e. presenting people with alternatives, a number of new and interesting questions seem to emerge, which in themselves are worthy of future research. What is the effect on people when their mobile phones start determining their own behavior, and how will people respond? Is there a risk that social occasions are 'technologized', that people start to behave according to their mobile phones? Also, how would a system like this come to change they way people work, their working styles?

In addition, one important question for interaction design and human-computer interaction is of course what happens when interaction in this way moves away from the screen and the keyboard to become a social negotiation process?

If nothing else, it is our hope that design research, like the system presented in this paper and the ideas behind it, that seeks to question the sometimes implicit frameworks that seem to freeze our perspectives and guide our designs, at least will come to influence both consumers and producers, as well as other researchers, so that some of the ideas might lead to new ideas that might eventually trickle its way through a full product design process.

References

1. Aura Communications Technology, Inc. <<http://auracomm.com/>>
2. Abowd, G. D., Dey, A. K., Brotherton, J., Orr, R. J.: Context Awareness in Wearable and Ubiquitous

- Computing. First International Symposium on Wearable Computers, IEEE Computer Society Press, Los Alamitos CA (1997) 179—180
3. Davies, N., Blair, G., Cheverst, K., Friday, A.: Supporting Adaptive Services in a Heterogeneous Mobile Environment. In Cabrera, L.-F. & Satyanarayanan, M. (Eds.) Proceedings of the Workshop on Mobile Computing Systems and Applications (Mobile '94, Santa Cruz, CA, Dec.) IEEE Computer Society Press, Los Alamitos CA (1994) 153—157
 4. Dix, A., Rodden, T., Davies, N., Trevor, J., Friday, A., Palfreyman, K.: Exploiting Space and Location as a Design Framework for Interactive Mobile Systems. Transactions on Computer—Human Interaction, Vol. 7, No. 3, Sept. (2000) 285—321
 5. Dourish, P.: Where the Action Is: The Foundations of Embodied Interaction. MIT Press, Cambridge MA (2001)
 6. Dreyfus, H.: What Computers Still Can't Do, MIT Press, Cambridge MA (1992)
 7. Fallman, D.: From Freedom to Involvement: On the Rhetoric of Mobility in HCI Research, Proceedings of HICSS-38, Hawaii International Conference on System Sciences, January 3-5, Hilton Waikoloa Village, Big Island, Hawaii: IEEE (2005).
 8. Fallman, D.: In Romance with the Materials of Mobile Interaction: A Phenomenological Approach to the Design of Mobile Information Technology, Doctoral Thesis, ISSN 1401-4572, RR.03-04, ISBN 91-7305-578-6, Umea University, Sweden: Larsson & Co:s Tryckeri (2003).
 9. Feiner, S., MacIntyre, B., Hollerer, T., Webster, A.: A Touring Machine: Prototyping 3D Mobile Augmented Reality Systems for Exploring the Urban Environment, First International Symposium on Wearable Computers, IEEE Computer Society Press, Los Alamitos CA (1997) 74—83
 10. Harrison S., Dourish, P.: Re-Place-ing Space: The Roles of Place and Space in Collaborative Environments. In Proceedings of ACM Conference on Computer-Supported Cooperative Work CSCW'96, (Boston, November) ACM Press, New York NY (1996)
 11. Ishii, H., Ullmer, B.: Tangible Bits: Towards Seamless Interfaces between People, Bits and Atoms. In Pemberton, S. (Ed.) Proc.of the ACM Conference on Human Factors in Computing Systems (CHI97) (Atlanta GA, March) ACM Press, New York NY (1997) 234—241
 12. Kakahara, M. & Sørensen, C.: Mobility: An Extended Perspective. In Sprague, R. (Ed.) 35th Hawaii International Conference on System Sciences, HICCS-35 (Big Island, Hawaii), IEEE, Los Alamitos CA (2002)
 13. Want, R., Schilit, B. N., Adams, N. I., Gold, R., Petersen, K., Goldberg, D., Ellis, J. R., Weiser, M.: An Overview of the ParcTab Ubiquitous Computing Experiment. IEEE Personal Communications. 2, 6, Dec. (1995) 28—43.
 14. Waterworth, J. A.: Spaces, Places, Landscapes and Views: Experiential Design of Shared Information Spaces. In Munro, A., Höök, K., & Benyon, D. (Eds.) Social Navigation of Information Spaces. Springer-Verlag, London UK (1999)
 15. Weiser, M.: The Computer of the 21st Century. Scientific American, 265, 3 (1991) 66—75.
 16. Winograd, T.: Architectures for context. Human-Computer Interaction, 16(2-3) (2001).

Acknowledgements

We would like to thank Interactive Institute for funding parts of this study and for allowing us to use their premises for testing and development. Thanks also to Brendon Clark, Microchip, and Coilcraft.