Regarding System Output in Situated Computing

Peter Ljungstrand
PLAY Research Group, Interactive Institute
c/o Viktoria Institute, Box 620, SE-405 30, Göteborg, Sweden
peter@viktoria.informatics.gu.se

ABSTRACT

This paper highlights output from systems in situated computing as a promising research area, especially when accompanied by a user oriented approach. A number of aspects of situated output are discussed and illustrated with two applications. PowerView is an alternative interface for PDAs that is particularly suited for situated computing, and the Reminder Bracelet explores new ways of conveying notification events from a device to a user and her environment.

Keywords
Situated output, Context awareness, Notifications, User oriented design

INTRODUCTION

Much of the research on context-aware and situated computing so far has been driven by a technology oriented approach. There have been a number of propositions on how to make creative use of various kinds of sensors for passively capturing different aspects of the surrounding physical environment, as well as various cues from the virtual or computational environment of a user, or rather the user's computer [cf. 3, 8, 9, 10, 11]. The research has mainly been about novel ways of collecting and refining various inputs, and then applying this input in interesting ways, but using fairly standard types of applications. To illustrate, the main difference between a context-aware tourist guide and a desktop based, or even a "standard" PDA based tourist guide, has been how the applications get their input events. In the context-aware case, input events are generated implicitly by monitoring the environment, whereas in the more traditional case the user explicitly and consciously generates the input events.

Focusing on users and applications

Of course there are exceptions [e.g. 6, 9], but many papers on the subject only have a short paragraph considering applications and users, whereas most of the text is spent discussing the technological assets of the system described. This technology research is highly relevant and needed for the situated computing field. I believe, however, that it needs to be complemented to a greater extent by user oriented, applied research. This shifts the leading research question from something like "How can we exploit all these cool sensors?" to "How can we support and augment human activities in a variety of changing situations?" For instance, imagine a project meeting, where your PDA by default only displays information that is related to that very project, of course with the possibility of manual override, by monitoring who else is present, who is not, the meeting time and location, etc.

Tools and appliances

A chef's knife is superb when cooking. A Swiss Army knife also could do the job, but not as easily. Tools optimized for a certain activity tend to be used as an extension of oneself, whereas the focus or attention of the user is on the task at hand, and not on the tool. On the other hand, a Swiss Army knife is versatile, small and mobile, and useable for many other things than cooking. I believe that we must consider similar tradeoffs when designing computationally enhanced tools that support situated and changing activities.

Following current technological and cultural trends, many people seem to be increasingly relying on a multitude of computational devices, may they be mobile, wearable or embedded in the environment [cf. 9]. Both the devices and their users will probably have dynamically changing usage contexts over time. I believe that due to the technical focus of most prior research, this work has mainly considered new ways of system input, and to a somewhat lesser degree, data processing, besides a number of proposed applications. I argue that to fully address the needs of situated computing, one must not only consider novel system *input* and *processing*, but also novel *output*, with a particular *purpose* in mind.

SITUATED OUTPUT

The physical surroundings of a typical PC are relatively static. There is sufficient light, it is relatively quiet, almost constant temperature, no rain or snow, etc. Consequently, there has been no particular need to consider such variables when designing PC applications. Or rather, there have been certain implicit assumptions about the physical use context. However, in certain cases the systems, and consequently the design of their output, has been optimized for a totally different, but specific, physical context. There has been significant research on system output for such specific environments, e.g. inside a car or an airplane cockpit. The important issue here is what variables are allowed to change (the situational aspect) in each of these settings, and what variables are (often implicitly) assumed to be constant.

With mobile devices things have changed, as some of the situational aspects should be allowed to change dramatically. However, many developers still seem hold the same assumptions about the use context as they did when developing desktop office applications. A mobile device can of course also be specialized for a certain task (e.g., a mobile phone), or be a more or less general purpose device (closer to a PDA). Systems and devices that are supposed to be used in a variety of physical contexts and situations typically have a fairly general interface that does not necessarily suit all of the proposed use contexts, which applies to both the physical surroundings and the activity at hand.

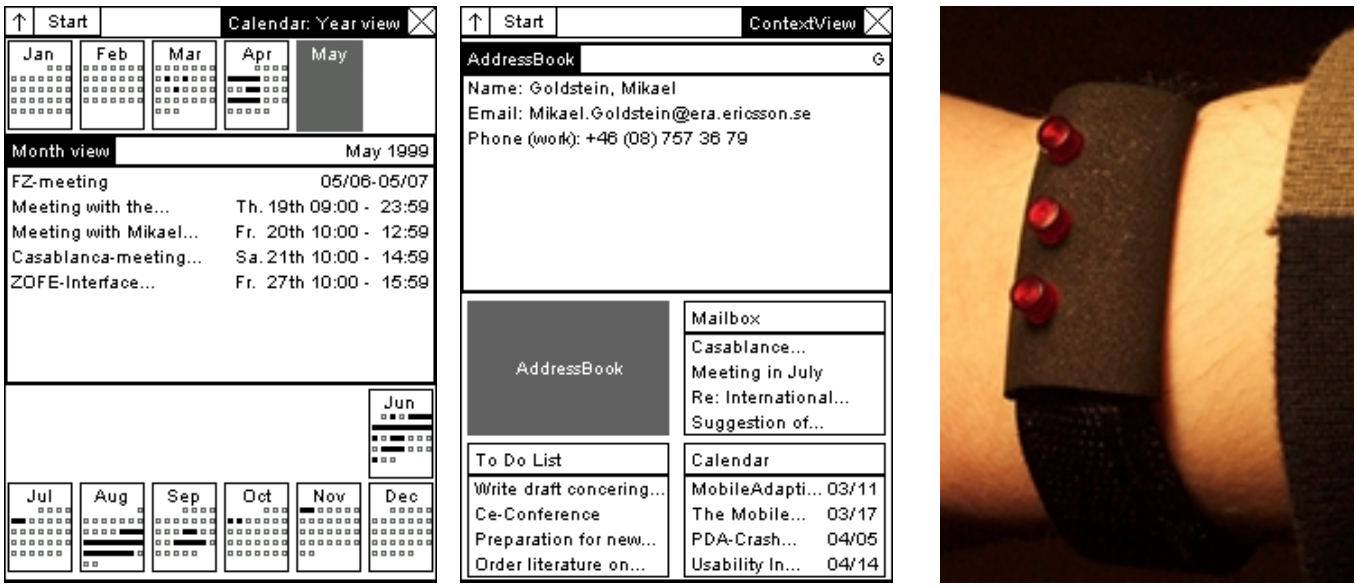

Figure 1: To the left, navigating in the Calendar view in PowerView with focus on the month May 1999. In the middle, a context view showing the address book entry “Mikael Goldstein” in focus, and semantically related information from other ‘applications’ as context. To the right, a closeup view on the Reminder Bracelet on a user’s wrist.

We can distinguish between two basic forms of situatedness here. One is where the system is primarily dedicated to an activity, a person or a role, but where the place or physical environment is allowed to vary. A system can also concern primarily a physical context, a place, location or environment, but with less regard to who the user is or her current activity.

Situated output on small screens

Inasmuch as we use screen-based output from a system, the appearance on the screen should make use of the current context, not only in terms of *what* is being presented, but also *how*. This is of particular importance for mobile devices with small screens. Let us return to the project meeting discussed above for a moment. Imagine that you and the rest of the project group take a break for lunch. As you walk down the street to find a nice restaurant, you continue your discussions, and of course you all bring your PDAs. Since your PDAs are aware of that the other project members are still nearby, they automatically filter out non-project related information. However, the physical context has changed dramatically. Now you are out in bright sunlight, making the screen harder to read. And since you are walking, you have an even harder time trying to read the tiny characters on the screen. In such a case, the system could automatically adapt to the changed physical context by for instance entering a high-contrast mode for the screen (only black and white, no color or grayscales), and a substantially larger font size to compensate for the reading difficulties. A similar adaptive system was described in [11], even though this system did not detect or adapt to other people (or rather their devices) in the proximity.

The scenario above illustrates how a system can adapt to changes in the physical context and then change the (display) output accordingly. But in this case, the activity being supported did not change, at least not from the PDA’s point of view. The application is still trying to support your project meeting, only now while walking down the street.

Information presentation and navigation

As information visualization often is an important issue when using large, high resolution screens, we believe it to be of even greater importance for very tiny screens. To explore some of the properties of small screens, handheld devices and changing use contexts, we have been developing and evaluating a new interface, called PowerView [1]. This is an alternative way of displaying and interacting with a PDA, building on our experience with information visualization and user oriented design. PowerView uses Flip Zooming [5], a focus+context visualization technique, which uses most of the available screen space to display a piece of information that is in the user’s *focus*, while at the same time providing a *context* with smaller, less detailed or abstracted information (see Fig. 1). The rationale is that the context simplifies navigation, as the user gets a sense of how much information there is, what it is about, the current position (e.g., address book entry) within the information set, etc. It also tries to make efficient use of the available screen space. The current prototype integrates four typical PDA applications into one. In addition to allow for navigation within a particular information type (‘application’), such as within an address book or calendar (Fig 1, left), we sometimes display semantically related information from other ‘applications’ as context objects (Fig 1, middle).

PowerView is typically navigated using a stylus, but it also allows for single-handed use of a PDA with the Action Control (a wheel that allows for scrolling and selection) on the side of a Casio Cassiopeia E-11 PDA. This possibility makes PowerView even better suited for a variety of use context. We are currently developing a second generation of PowerView that will take more of the usage context into account, such as for instance detecting other users (devices)

in the proximity, handling of incoming and outgoing phone calls, etc.

Attracting the users' attention

Another aspect of system output in situated computing is how the devices or applications attract their users' attention, especially if the device is not being used actively, i.e. when is not in the foreground of the user's activities [cf. 2]. While some researchers have noted this problem, no one seems to have systematically addressed this issue. For instance, Rekimoto et al. [7] note "that the key design issue for augment-able reality is how the system can gracefully notify situated information." They continue by stating that using beeps and vibrations are not good enough for their application and propose overlaying notification information on a see-through heads-up display, possibly embedded in eyeglasses, or perhaps notification using LEDs on eyeglasses.

Presently, users of mobile devices are notified mainly by attention-demanding sounds and beeps, and to a lesser extent by tactile vibrating devices such as those found in some mobile phones or pagers. Having a multitude of devices that are unaware of each other's states can result in situations that are very hard to interpret, for instance when different notifications coincide. Furthermore, even though notifications generally are intended for the owner of the device, other people nearby can be disturbed, making notifications a social issue as well. People typically address this by disabling alerts sounds in social environments, e.g. business meetings or cinemas. However, this is not a satisfactory solution, as the users actually want to be notified in most cases.

An alternative way of conveying notifications

In an effort to explore alternatives to current notification techniques, we developed the Reminder Bracelet [4]. It is a wristband with embedded LEDs, connected via a thin cable to a user's PDA. It notifies its user of scheduled events stored in the PDA in a subtle and silent manner using light, color and patterns. Initial evaluations have shown that in a number of situations, the Reminder Bracelet was preferred to the alert sound of a PDA or a mobile phone.

While audible signals can be disturbing in some cases, tactile displays are generally non-intrusive and silent. They are also private; for instance a vibrating device, like the ones used in mobile phones and pagers needs to be carried rather close to the body to work well. This means the device has a low degree of publicity. A device such as the Reminder Bracelet falls somewhere in between an alarm sound and a vibrating device in this respect. Let us return to the project meeting again. Imagine that one of the participants suddenly rises and leaves the room, without saying a word. Now, imagine the same situation, but preceded by a ringing signal from his mobile phone. Using notification cues with a higher degree of publicity allows other people present to easier interpret and make sense of the situation at hand, e.g. in terms of causality. In a social context, we need appropriate cues that fit our mental model to understand the behavior of others.

The work on Nomadic Radio [9] shares our interest in providing mobile users with timely notification cues, and highlights the problem of differentiating between various cues from a device. This is especially important when considering displaying cues with different media types, such as audio, tactile or visible cues. However, Nomadic Radio uses a private, audio only environment that needs to be worn at all times, and as the system uses earphones, it has a low degree of publicity.

Conclusions

To summarize, I have tried to show that besides the much needed, technically oriented work on new types of input to context-aware systems, a user oriented approach is needed to fully take advantage of the possibilities that the new technology is providing. In particular, I have focused on situated output from a user perspective, as I believe this is a promising area for further research. While I have considered a number of different use contexts and environments that devices should support in some way, I have not given much though on how to allow for (seamless) transitions between such different situations [cf. 2]. This will be another area for future research.

REFERENCES

1. Björk, S., Holmquist, L.E., Ljungstrand, P. and Redström, J. (2000). PowerView: Structured Access to Integrated Information on Small Screens. To appear in: *Extended Abstracts of CHI'2000* (Short Talk). ACM Press.
2. Buxton, W. (1995). Integrating the Periphery and Context: A New Model of Telematics. In: *Proceedings of Graphics Interface '95*.
3. Golding, A. and Lesh, N. (1999). Indoor navigation using a diverse set of cheap, wearable sensors. In: *Proceedings of ISWC'99*, IEEE Press, pp. 29-36.
4. Hansson, R. and Ljungstrand, P. (2000). The Reminder Bracelet: Subtle Notification Cues for Mobile Devices. To appear in: *Extended Abstracts of CHI'2000* (Student Poster). ACM Press.
5. Holmquist, L.E. (1997). Focus+Context Visualization with Flip Zooming and the Zoom Browser. In: *Extended Abstracts of CHI'97*, ACM Press.
6. Redström, J., Dahlberg, P., Ljungstramd, P. and Holmquist, L. E. (1999). Designing for Local Interaction. In: *Proceedings of the First International Workshop on Managing Interactions in Smart Environments (MANSE'99)*, 13-14 December, Dublin, Ireland. Springer-Verlag. pp. 227-238.
7. Rekimoto, J., Ayatsuka, Y. and Hayashi, K. (1998). Augment-able Reality: Situated Communication through Physical and Digital Spaces. In: *Proceedings of ISWC'98*, IEEE Press.
8. Sawhney, N. and Schmandt, C. (1999). Nomadic Radio: Scaleable and Contextual Notification for Wearable Audio Messaging. In: *Proceedings of CHI'99*, ACM Press, pp 96-103.
9. Want, R., Schilit, B., Norman, A., Gold, R., Petersen, K., Goldberg, D. Ellis, R. and Weiser, M. (1995). *The ParcTab Ubiquitous Computing Experiment*. Technical Report CSL-95-1, Xerox Palo Alto Research Center, March 1995.
10. Salber, D., Dey, A. and Abowd, G. (1999). The context toolkit: aiding the development of context-enabled applications. In: *Proceedings of CHI'99*, ACM Press, pp. 434-441.
11. Schmidt, A., Aidoo, K., Takaluoma, A., Tuomela, U., Van Laerhovem, K. and Van de Velde, W. (1999). Advanced Interaction in Context. In: *Proceedings of HUC'99*, Springer Verlag, pp.89-101.