

Adding Some Smartness to Devices and Everyday Things

Hans-W. Gellersen, Albrecht Schmidt and Michael Beigl

TecO, University of Karlsruhe

Vincenz-Prießnitz-Str. 1, 76131 Karlsruhe, GERMANY

Phone +49 (721) 6902-49

{hwg | albrecht | michael}@teco.edu

Abstract

In mobile computing, context-awareness indicates the ability of a system to obtain and use information on aspects of the system environment. To implement context-awareness, mobile system components have to be augmented with the ability to capture aspects of their environment. Recent work has mostly considered location-awareness, and hence augmentation of mobile artifacts with locality. In this paper we discuss augmentation of mobile artifacts with diverse sets of sensors and perception techniques for awareness of context beyond location. We report experience from two projects, one on augmentation of mobile phones with awareness technologies, and the other on embedding of awareness technology in everyday non-digital artifacts.

1. Introduction

It is now widely acknowledged that some awareness of the context in which mobile systems are used can produce added value and foster innovation in many application domains. In mobile computing, the notion of context is generally used in reference to aspects of the environment in which a mobile system operates and to which the system might adapt or respond with appropriate behavior. While context is an open-ended concept, it is commonly associated with straightforward aspects in mobile system environments such as location of users, whereabouts of system components, local availability of resources and such like.

To facilitate awareness of context in mobile systems, some system components have to be augmented with the ability to capture aspects of the system environment, by way of sensing or communicating. Often this is just one component, for example a personal mobile device for context-aware application access; in other systems this may be many components, for example mobile physical objects with location tags to assert overall system context. Either way, each piece of context that enters a distributed mobile system does so through an appropriately augmented system component. Our concern in this paper is how system components can be augmented

appropriately, i.e. how context-awareness can be added to mobile devices and artifacts. The research we report is based on a device-centric view, in which context is primarily associated with a device. For our discussion it is secondary, that context may also be associated with the user of a mobile device, or with applications that may run on the device or elsewhere in a distributed mobile system.

Most of the context-aware mobile systems discussed to date consider location as context, and from a device-centric perspective they are based on adding location-awareness to one or many of their system components. Three general approaches can be distinguished. First there are systems in which components utilize the mobile communications infrastructure to obtain location information, for example the cell-of-origin in cell-based communications. For example, the GUIDE system for tourists in Lancaster employs mobile computers that derive their location from a WaveLAN network [3]. Secondly, components may be equipped with explicit location sensors, i.e. receivers for specific location services, such as GPS. For example, the stick-e-note system for context-aware information access in fieldwork is based on palmtops augmented with GPS receivers [9]. Thirdly, components may be augmented in ways that allows surrounding infrastructure to assert their location. In this case, components strictly speaking have no awareness themselves but it is their augmentation that enables awareness. Examples are name tags in the Active Badge system [6], and the palmsize ParcTab terminals [12], both augmented with infrared diodes that emit signals from which the transceiver infrastructure derives location.

Location is a rich concept, and often it is not the location as such but also information associated with locations that is exploited in location-aware mobile systems. However we would argue that there is more to context than we can capture through location, and our focus in this paper is on augmentation of mobile system components for awareness of context beyond location. More specifically, we investigate the use of diverse sets of sensors in mobile system components for context-awareness. We report experience from two research projects on sensor-based context-awareness, TEA and Mediacup. The TEA project investigates *Technologies for Enabling Awareness* and their application in mobile

telephony [13]. The Mediacup project studies capture and communication of context in everyday environments [2]. The novel issues investigated in these projects are the integration of diverse sensors and perception techniques, and the embedding of autonomous awareness in mobile artifacts.

Diverse sets of sensors and perception techniques are integrated to the end of shifting complexity in context-awareness from algorithmic level to architectural level. This is done by considering deliberately simple sensors and feature extraction methods as opposed to expensive hardware and algorithms. Advanced context-awareness is then achieved through fusion of information obtained from diverse sensors, employing suitable architectures. The approach somewhat contrasts for example vision-based approaches that tend to be compute-intensive, and is geared toward implementation with embedded technologies.

The second issue highlighted in the work we report is the embedding of autonomous awareness in mobile artifacts. It is straightforward to add awareness technology – sensors and perception algorithms – to general purpose computing platforms such as laptops, personal digital assistants and wearable computers. Both the TEA and the Mediacup project however investigate the adding of awareness technology to artifacts that do not provide any platform ready for extension with hardware and software. In the case of TEA, the artifact considered is a mobile phone, which is based on digital technology but still self-contained and not open for extension. In the Mediacup project the challenge is taken further by considering an ordinary coffee cup, representing everyday artifacts. In both projects, artifacts have been augmented and studied in test environments.

In the subsequent sections, we will briefly discuss related work on sensor-augmented mobile artifacts, and then report experience first from the TEA project and secondly from Mediacup work. This will be followed by discussion that sums up our experience with adding context-awareness to mobile artifacts, also pointing out issues and directions for further research.

2. Related work

In a wide range of projects mobile artifacts have been augmented to enable awareness of their location. While three general approaches can be distinguished as discussed in the introduction, artifacts fall actually into two groups. First artifacts that have general-purpose computing platforms ranging from smallest-scale, consider for instance ParcTabs, to high-end wearable PCs. Secondly artifacts explicitly designed for being located such as the Active Badge infrared sender, and the Active Bat ultrasound emitter. Our work in TEA and Mediacup in contrast is concerned with augmenting artifacts that are neither general-purpose computing platforms nor non-functional beyond support of locality.

In handheld computing, there is some related work on adding sensor technologies beyond location to personal mobile devices. For example, Rekimoto added tilt sensors to a handheld to obtain context about the handling of the device [10]. Similarly, we have explored integration of orientation sensors in a handheld computer [14]. In this line of work, the context obtained from sensors is used as user interface extension. This is to be distinguished from context-awareness in mobile computing which is focused on using context to relate a mobile device to its surrounding environment.

While handheld computers generally still remain shielded from their surroundings, a stronger interest in situating devices is pursued in many wearable computing developments. A key motivation for wearable computers is to support their users in improved and proactive ways on the grounds of being permanently with the user. A precondition is a suitable understanding of the user's situation, and in this context a range of projects have investigated sensor integration to obtain information on both user and environment. For example, cameras and computer vision have been integrated with wearable computers for visual context-awareness [16]. While there has been some research into lower-cost vision techniques, this still assumes a suitably powerful computing platform. Beyond vision, the use of other sensors has been explored in a range of wearable computing applications. For instance, the Oregon wearable was equipped with sensors for object presence in a collaborative field engineering application [1], and in the Startlecam application bio-sensors were employed to the end of recognizing extreme user situations [7]. However, these are applications with task focus, and sensor integration is not generalized for wider applicability.

In wearable computing, two projects come close in spirit to our work. Paradiso has investigated sensor integration in footwear with a range of applications [8]. While the project was primarily concerned with enabling shoes as an expressive user interface, this is still related to our Mediacup work as it also augments a non-digital artifact. In both expressive footwear and Mediacup the approach is to obtain information from ordinary use: expressive footwear generates information as the user moves around, and likewise the Mediacup generates information in the course of being used as an ordinary coffee cup. In different ways close to our work is that of Golding and Lesh, who investigated integration of diverse sensors as alternative location technique for indoor navigation [5]. Like we did in the TEA project, they focused on integration of deliberately simple sensors. In their method, multi-sensor data is associated with locations, while in TEA it is associated with a more general notion of context beyond location.

3. TEA - an add-on device for context-awareness

The general motivation underlying the TEA project is to make personal mobile devices smarter. The assumption is that the more a device knows about its user, its environment and the situations in which it is used the better it can provide assistance. The objective of TEA is to arrive at a generic solution for making devices smarter, and the approach taken is to integrate awareness technology – both hardware and software – in a self-contained device conceived as plug-in for any personal appliance which from a TEA perspective is called host. The cornerstones of the TEA device concept are:

- Integration of diverse sensors, assembled for acquisition multi-sensor data independently of any particular application.
- Association of multi-sensor data with situations in which the host device is used, for instance *being in a meeting*.
- Implementation of hardware, i.e. sensors and processing environment, and software, i.e. methods for computing situational context from sensor data, in an embedded device

A specific objective underlying sensor integration is to address the kind of context that can not be derived from location information at all, for example situations that can occur anywhere. While it seems obvious that there is context that can not be inferred from location information, most work in context-awareness has actually served to show how that rich context can be derived from location provided location semantics beyond specification of position are available.

Another specific issue investigated in TEA is sensor fusion. The aim is to derive more context from a group of sensors than the sum of what can be derived from individual sensors.

3.1. TEA architecture

TEA is based on a layered architecture for sensor-based computation of context as illustrated in figure 1, with separate layers for raw sensor data, for features extracted from individual sensors ('cues'), and for context derived from cues.

The sensor layer is defined by an open array of sensors including both environmental sensors for perception of the real world and logical sensors for monitoring of conditions in the virtual world, for instance logical state of the host device. The data supplied by sensors can be very different, ranging from slow sensors that supply scalars (e.g. temperature sensor) to fast and complex sensors that provide a large amount of more or less structured data (e.g. a camera or a microphone); also the update time varies from sensor to sensor.

The cue layer introduces cues as abstraction from raw sensor data. Each cue is a feature extracted from the data stream of a single sensor, and many diverse cues can be derived from the same sensor. This abstraction from sensors to cues is generic, i.e. independent of any specific application. This process of preprocessing sensor data has also been referred to as cooking sensors [5], and serves to reduce the amount of data substantially before further abstraction. Just as the architecture does not prescribe any specific set of sensors, it also does not prescribe specific methods for feature extraction in this layer. However, in accordance with the philosophy of shifting complexity from algorithms to architecture it is assumed that cue calculation will be based on comparatively simple methods. The calculation of cues from sensor values may for instance be based on simple statistics over time (e.g. average over the last second, standard deviation of the signal, quartile distance, etc.) or on somewhat more complex mappings and algorithms (e.g. calculation of the main frequencies from a audio signal over the last second, pattern of movement based on acceleration values). The cue layer hides the sensor interfaces from the context layer it serves, and instead provides a smaller and uniform interface defined as set of cues describing the sensed system environment. This way, the cue layer strictly separates the sensor layer and context layer which means context can be modeled in abstraction from sensor technologies and properties of specific sensors. Separation of sensors and cues also means that both sensors and feature extraction methods can be developed and replaced independently of each other.

The context layer introduces a set of contexts which are abstractions of real world situations, each as function of available cues. It is only at this level of abstraction, after feature extraction and data reduction in the cue layer, that information from different sensor is fused in the process of calculating context. While cues are assumed to be generic, context is considered to be more closely related to the host device and the specific situations in which it is used. Again, the architecture does not prescribe the methods for calculation of context from cues; rule-based algorithms, statistical methods and neural networks may for instance be used. Conceptually, context is calculated from all




Figure 1. TEA is based on a layered architecture for abstraction from raw sensor data to multi-sensor-based context.

available cues. In a rule set however, cues known to be irrelevant may simply be neglected, and in neural network their weight would be reduced accordingly. The context calculation, i.e. the reasoning about cues to derive context, may be described explicitly, e.g. when cues are known to be relevant indicators of a certain real world situation, or implicitly in methods that learn context from example data.

The context layer hides lower interfaces from applications, which are based on the context interface. In the application, context can then be associated with reactive behaviour.

3.2. Initial exploration of the approach

To study the TEA approach, we have developed two generations of prototype devices and used them for exploration of multi-sensor data, and for a validation of TEA as add-on device for mobile phones. In parallel to development of the first prototype we have also conducted scenario-based requirements analysis to investigate our assumption that there is useful context for personal mobile devices that can not be derived from location but from multi-sensor input. In this analysis, a range of scenarios were developed for both mobile phones and personal digital assistants (PDA), and it was found that the potential for context beyond location was higher in communication-related scenarios than in typical PDA applications which led us to focus further studies on the domain of mobile telephony.

The TEA device was developed in two generations. The first generation device was developed for exploration of a wide range of sensors and their contribution to context-awareness. It contained common sensors such as microphone, light sensor and accelerometers but also



Figure 2. The current implementation of the TEA awareness device has about the size of a mobile phone battery pack.

sensors for example for air pressure, certain gas concentration and so on. With several implementations of the device, large amounts of raw sensor data were collected independently at different sites for further analysis of multi-sensor fusion following two strategies:

- Analysis of the contribution of a sensor or group of sensors to perception of a given context, i.e. a specific real-world situation: For this study a number of situations that we considered relevant for personal mobile devices were defined (e.g. user is walking, user is in a conversation, other people are around, user is driving a car, etc.). Then data was collected for each of these situations, with independent data collection at three different sites. The data was then subjected to statistical analysis to determine for each sensor or sensor group whether its inclusion increased the probability of recognizing situations.
- Analysis of clusters in collected multi-sensor data: Here the strategy was to carry the device over a longer period of time so it accompanies a user in different situations. Over the whole period of time, raw sensor data was recorded and to be later analyzed to identify clusters corresponding to situations that occurred during recording time, e.g. situations such as user is sitting at her desk, walking over to a colleague, chatting, walking back, engaging in a phone conversation and so on. This process was aimed at identifying the sensors relevant to situations, and at development of a clustering algorithm supporting awareness of situations of interest.

3.3. Prototype implementation and validation

The initial exploration of sensors and their contribution to awareness of typical real-world situations served to inform development of the second generation device optimized for smaller packaging, and shown in figure 2. The device integrates two light sensors, two microphones, a two-axis accelerometer, a skin conductance sensor and a temperature sensor. The sensors are read by a micro-controller, that also calculates the cues and in some applications also the contexts. The system is designed to minimize the energy consumption of the component. The micro-controller (PIC16F877) has a number of analog and digital inputs and communicates via serial line with the host device. The calculation of cues and contexts is very much restricted due to the limitations of the micro-controller. Programs have to fit into 8K of EEPROM, and have only 200 Bytes of RAM available.

The feature extraction algorithms to generate the cues have been designed to accommodate these limitations. Data that has to be read with high speed such as audio is directly analyzed and not stored. Typical cues for audio that are calculated on the fly are the number of zero crossing of the signal in a certain time (indicator of the frequency) and number of direction changes of the signal

(together with the zero crossings this is an indicator of the noise in the audio signal). For acceleration and light basic statistical methods and an estimation of the first derivative are calculated. Slowly changing values – temperature and skin conductance – are not further processed in the cue layer (the cue function is the identity). The contexts are calculated based on rules that were extracted off-line from data recorded with the sensor board in different situations.

The prototype is independent of any specific host and has been used in conjunction with a palmtop computer, a wearable computer and mobile phones. Primarily however the prototype is being applied in the area of mobile telephony. State of the art mobile phones support so-called profiles to group settings, such as notification mode, input and output modality, and reaction to incoming messages and calls. Users can define profiles for different situations (e.g. home, meeting, car, etc.) and specify behavior desired in those situations. The TEA device has been added to a mobile phone to automate activation of such profiles which otherwise have to be activated manually by the user. The approach was validated in an experiment, in which the TEA device was used to control a small set of typical profiles [13].

3.4. Application in mobile telephony

An interesting application domain for context-aware mobile phones as enabled by TEA is the sharing of context between caller and callee. For a caller, context may be helpful for instance to assess whether it is a good time to call (in fact, “is it a good time to call” is quite commonly asked when a phone conversation is initiated), and for a callee it may help to assess importance of an incoming call (“is it important or can I phone back later” – a common question in accepting a call). To study context-enhanced communication, we have implemented the WAP-based application “context-call”. In this application, a call is initiated as usually by entering the number of the callee. The application however does not establish the call straightaway but instead looks up the context of the callee and provides this information to the caller. The caller is then prompted to decide how to proceed – for example whether to use a voice service or a short message service. A detailed discussion of the application is provided in [15].

3.5. Discussion of TEA experience

Our experience gathered in the TEA project supports the case for investigation of context beyond location, and for fusion of diverse sensors as an approach to obtain such context. We have used the approach for obtaining strictly location-independent context such as “in a meeting”, “in a conversation”, “user is walking” which can not be derived from location information. As for sensor fusion, our analysis of collected multi-sensor data showed that with our approach context can be derived beyond the sum of

what can be obtained from individual sensors. This initial experience is valuable, however it is clearly not sufficient to derive any methodology for systematic application of sensor fusion for context-aware applications. However, what we find generalizable is the layered approach to perception. The two-step abstraction first from sensors to cues and then from cues to context proved to be a suitable strategy for the perception process as such, and in addition it also supports architectural qualities such as modularity and separation of concerns.

In TEA, extensive experience was gained with a wide range of sensors and their integration. From this experience we can derive some indication as to which sensors are of particular interest for the overall objective of capturing real world situation. We found that in particular sensors for audio, movement and light provide contributions to awareness in most settings while most other sensors have rather specific applications in which they are valuable. In addition we found that perception can be improved by using not just diverse sensors but also multiple sensors of the same kind, in particular microphones and light sensors with different orientation. More generally, it was found that placement substantially influences the contribution of a sensor to multi-sensor based awareness. In some ways, this challenges the approach of tightly packing sensors. In the context of augmenting personal mobile devices, an alternative would be disaggregation and distribution of sensors for instance on the user’s body or clothing, assuming a body area network for data collection.

Last not least, it should be noted that our experience also extends to the exploration of practical applications with commercial prospect such as the context call we briefly discussed. The community is currently debating what the killer application of context-awareness might be, and based on our research we would suggest that if there is a killer application it will be in the area of interpersonal communication.

4. Mediacup – embedding awareness technology in everyday artifacts

The Mediacup project was conducted in parallel to TEA, and while also investigating embedded awareness technology it is motivated differently. TEA is about making artifacts smarter, i.e. to improve the functionality the artifact offers their user. In contrast, the Mediacup project is about using artifacts to collect context information transparently, i.e. without changing the function and use of the artifact. The core idea is that by embedding awareness technology in the everyday things people use we can obtain context on everyday activity so to speak at the source. This approach assumes a distributed system in which some artifacts are augmented to collect context information, while other artifacts are computationally augmented to use such context.

4.1. Aware artifacts model

The context-awareness model investigated in the Mediacup project is based on the following concepts:

- Artifacts are augmented with an awareness of their own local context. To this end artifacts are equipped with sensors but also with a processing environment and software for autonomous calculation of artifact-specific context from sensor data.
- Artifacts broadcast their context in their local environment. To this end aware artifacts are augmented with basic communication capabilities. In the presence of many aware artifacts, context broadcast establishes a context information spaces of a certain local scope.
- Any applications, appliances or information artifacts in the environment can use the locally available context, without further knowledge of the artifacts from which the context originate.

4.2. Mediacup – Awareness embedded in coffee cups

For exploration of the aware artifacts model we have augmented coffee cups representing non-digital everyday artifacts with awareness technology. The Mediacups, as we call the augmented mugs, contain hardware and software for sensing, processing and communicating the state of the cup as context information.



Figure 3. The Mediacup is an ordinary coffee cup with sensors, processing and communication embedded in the base.

The current implementation of the MediaCup is shown in figure 3 on the right. It is the result of several design iterations that have been carried out over the last two years. The goals of the hardware development were to provide an ordinary cup with sensing capabilities, processing power, and communication. The design challenge was to provide these additional features without changing the basic properties (shape, size, and weight) of the cup noticeably and without compromising everyday use (ensuring robustness, and maintenance-free use). The current version of the MediaCup hardware comprise a digital sensor for temperature, three metal ball switches to detect motion, a switch that detects when the cup is placed on a surface, an infrared diode for communication, and a microcontroller (PIC16F84) as processing unit. The power is stored in two 1 F GoldCaps, which can be wirelessly charged using a resonant circuit with 20kHz. The PCB is laid out circular so it can be placed in the cup base. The board with all components mounted is only 3mm high.

The Mediacup software controls acquisition of raw data from sensors and on top of that computes cup-specific context. The process of sensor reading and abstraction is designed to minimize energy consumption. Movement is a parameter that can change fast and frequently, but in most cases a cup will be not moved. To detect movement by sensor polling would have required readings about every 20ms; to avoid this, the motion detectors are connected to the interrupt pins of the processor, triggering readings only when changes have occurred. Detected movement is recorded as event, and a short history of such events is used in a rule-based heuristic to detect more abstract events with a cup-related meaning; these are *cup is stationary*, *cup is moving*, *drinking out of the cup*, and *fiddling with the cup*. In contrast to movement, temperature is a parameter that is changing slowly in the real world. Also, the adaptation speed of the sensor is very slow, and therefore it is read only every two seconds. The tracked temperature information in conjunction with some motion information is used to compute further cup-related context: *filled up*, *cooled off*, and *current temperature*.

Mediacups broadcast their context together with their unique ID (i.e. their IP address) every two seconds using the infrared LED which faces overhead. The communication range is about two meters with an angle of 45°. The cup information is collected through an overhead transceiver infrastructure installed in the usage environment of the cups, i.e. 4 rooms in our office environment. The transceivers are based on HP's HSDL 1001 IrDA chip and have a footprint of about 1,5m². They are connected through a CAN bus (car area network) and a gateway to the local ethernet, in which collected context is broadcast in UDP packets.

4.3. Experience from design and use

Like TEA, the Mediacup project served to gather extensive experience with sensor-based context-

awareness. However while TEA primarily provided insights into issues surrounding sensor fusion and context perception architecture, the Mediacup provides substantial experience on different issues, i.e. on the embedding of awareness technology in 'unpowered' artifacts, on issues surrounding transparency of technology, and on a paradigm shift in use of sensors for context-awareness. Not surprisingly, the embedding of technology in artifacts that are not powered themselves raises issues of power management. Our experience from iterative design of the Mediacup is that power concern become a central issue that influence a wide range of design decisions:

- **Processing.** The used microcontroller runs with a reduced clock speed of only 1 MHz; this reduces the power consumption to below 2mA at 5.5V in processing mode. The processor is switch to sleep mode (power consumption below 1 μ A) whenever possible.
- **Motion detection.** In one of the early versions of the cup an accelerometer (ADXL202) was used. To reduce the power consumption and to make it feasible to wake up the electronic from sleep mode whenever the cup is moved without active polling the accelerometer was replaced by three ball switches. These switches are connected to the external interrupt inputs of the microcontroller. This makes it feasible to put the microcontroller more than 99% of the time in sleep mode without losing the information if the cup is moved.
- **Temperature.** A Dallas DS1621 chip was used to measure temperature (-55 to $+125$ $^{\circ}$ C). It consumes 1 μ A in standby mode and 400 μ A during the short reading cycles.
- **Recharging.** Nobody wants to change batteries in a coffee cup or plug in the coffee cup for recharging everyday. So two design option arise. First fitting in a battery that runs for the live time of the cup or second recharging the cup with no additional attention of the user. We went for the second choice: our approach is to provide a saucer with the cup that can be placed on the table and that is connected to a power line. Whenever the user puts the cup on the saucer the cup is wirelessly recharged. Away from the saucer the MediaCup electronics run about 12 hours with the 2-farad capacity.
- **Communication.** For communication a low-powered 5mm infrared diode is included (HSDL4420). The status of the cup is communicated every two second to the environment using the IrDA-physical layer coding. The IrDA coding is done in software on the microcontroller to save an additional component. The data rate is set to the maximum that is feasible to implement in software in the current design (19.2kbit/s) to reduce the time the diode has to be powered.

Exploration of the Mediacup also gave insights into issues of transparency. If an artifact is to be augmented in ways that don't compromise its common use, it does not suffice to minimize and hide the technology. For example, the requirement of free line of sight between artifact and transceiver infrastructure has to be transparent to the user: the design has to ensure that in common use free line of sight will be given. Another example, that came up with use experience with an early battery-powered prototype, was that power provision needs to be transparent. We observed that users would not care to check the battery, and to make sure they were recharged. This is was not expected but in hindsight is not surprising: the fact that the battery ran flat did – by design - not influence the artifacts use and only had effects that were not visible to the user. In the current prototype this issue was addressed by introducing the wirelessly chargeable Gold Caps, which are charged whenever a cup is placed on its accordingly augmented saucer.

Beyond the practicalities of transparent embedding of awareness technology in everyday artifacts, the Mediacup also provides early experience with a paradigm shift in how we perceive and design sensor-enhanced aware application. The traditional view is to consider sensors as periphery, and applications as the place to make sense of collected data. In the aware artifacts model as explored with the Mediacup, the notion of sensor periphery is replaced by a notion of what we might call sensory appliances. The making sense of sensory data is decentralized and shifted to the source of the data. The notion of context-aware application as embodying sensor integration is also replaced: in an environment such as explored in the Mediacup project there is no application that would explicitly take input from a set of sensor or sensory artifacts; instead there small specialised applications and appliances that consume some context. For instance, in the Mediacup environment, digital door plates which originally were built to leave notes at doors, were augmented to indicate that a meeting is in place whenever a co-location of filled coffee cups was derived from context in its local environment. This is not an important or far reaching application but it is indicative of the kind of context-based services that can emerge once a framework for collecting and providing context information is in place.

5. Discussion and conclusion

In the TEA and Mediacup projects we have gathered substantial experience with sensor-based context-awareness and embedding of awareness technology in mobile artifacts. We have gained important insights into sensor fusion for awareness of situational context, into architectural issues, into embedded design of awareness technology, and into a new perspective on context-enabled environments and applications.

We have shown that integration of diverse sensors is a viable approach to obtain context that represents complex real-world situations, and context that captures interaction with everyday artifacts. We have to some extent investigated generic approaches for deriving context from sensor data, and our experience suggests that some degree of abstraction, i.e. the calculation of cues, can be implemented independently of specific applications. In fact, we expect that future generations of sensors will provide general-purpose cues besides the raw sensor data. Our work also indicates the value of sensor fusion, however our experience is too limited to attempt any generalization to generic fusion methods.

Our work to date was not specifically focussed on architectural issues. The TEA architecture and the aware artifacts model though explore issues of modularity, separation of concerns, and the coupling of context acquisition and context consumption. It will be important future work to further investigate these issues and to develop principles for the architectural design of multi-sensor context-aware systems.

Embedded design of awareness technology gives rise to the old discussion of trading off performance for cost, with the most critical cost being power consumption. However our experience highlights substantial challenges for perception techniques to perform in low-end computing environment. In our work, in particular in the Mediakup project, we have carefully crafted sensor control to meet requirements. An important research direction in the area of multi-sensor based perception will be to embody sensor control to some extent, or to embody adaptation to changing sensor properties. We envision scalable perception techniques that perform robustly in conjunction with sensors that are dynamically powered on and off; this would introduce a notion of quality of service in perceived context, which we also expect to become an important research direction.

Finally, we believe our work contributes to development of new perspectives on application of context-awareness, in which acquisition and use of context disseminate further into everyday activity. The aware artifacts model is a first exploration in this direction, studying a shift from context-aware applications with sensor periphery to dynamic systems of specialized appliances and artifacts, some of which are augmented to capture context while others are augmented to use context.

5. REFERENCES

1. Bauer, M., Heiber, T., Kortuem, G. and Segall, Z. A Collaborative Wearable System with Remote Sensing. *Proceedings of the International Symposium on Wearable Computing*, Pittsburgh, Pennsylvania, 19-20 October 1998.
2. Beigl, M. and Gellersen, H.W. MediaCups: Experience with Design and Use of Computer-Augmented Everyday Objects. To appear in *Computer Networks, Elsevier Publishers*.
3. Cheverst, K., Davies, N., Mitchell, K. and Friday, A. The Role of Connectivity in Supporting Context-Sensitive Applications. *First International Symposium on Handheld and Ubiquitous Computing (HUC99)*, Karlsruhe, Germany, Sept. 1999, Lecture Notes in Computer Science No. 1707, Springer-Verlag.
4. Gellersen, H.W. and Beigl, M. Ambient Telepresence: Colleague Awareness in Smart Environments. *Proceedings of Managing Interactions in Smart Environments (MANSE'99)*, Ireland, December 1999, Springer-Verlag London, p. 80-88.
5. Golding, A. and Lesh, N. Indoor Navigation Using a Diverse Set of Cheap Wearable Sensors. *Proceedings of the IEEE International Symposium on Wearable Computing (ISWC99)*, San Francisco, CA, October 1999.
6. Harter, A. and Hopper, A. A Distributed Location System for the Active Office. *IEEE Network* 8(1), 1994.
7. Healey, J. and Picard, R. StartleCam: A Cybernetic Wearable Camera. *Proceedings of the International Symposium on Wearable Computing*, Pittsburgh, Pennsylvania, 19-20 October 1998, pp. 42-49.
8. Paradiso, J.A., Hsiao, K.-Y., and Benbasat, A. Interfacing the Foot: Apparatus and Applications. *CHI 2000 Extended Abstracts*, The Hague, The Netherlands, April 2000, ACM Press, p. 175-176.
9. Pascoe, J., Morse, D.R. and Ryan, N.S. Developing Personal Technology for the Field. *Personal Technologies* 2(1), March 1998, p. 28-36.
10. Rekimoto, J. Tilting operations for small screen interfaces. *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST'96)*, p. 167-168.
11. Salber, D., Dey, A.K. and Abowd, G.D.. The Context Toolkit: Aiding the Development of Context-Enabled Applications. In the *Proceedings of the 1999 Conference on Human Factors in Computing Systems (CHI '99)*, Pittsburgh, PA, May 15-20, 1999.
12. Schilit, B.N., Adams, N.L. and Want, R. Context-aware computing applications. *Proceedings of the Workshop on Mobile Computing Systems and Applications*, Santa Cruz, CA, Dec. 1994.
13. Schmidt, A., Aidoo, K. A., Takaluoma, A., Tuomela, U., Van Laerhoven, K., Van de Velde, W.. Advanced Interaction in Context. *First International Symposium on Handheld and Ubiquitous Computing (HUC99)*, Karlsruhe, Germany, Sept. 1999, LNCS 1707, Springer-Verlag.
14. Schmidt, A., Beigl, M. and Gellersen, H.W. There is more to context than location. *Computer & Graphics* 23 (1999) 893-901.
15. Schmidt, A., Takaluoma, A. and Mäntyjärvi, J. Context-Aware Telephony over WAP. *Personal Technologies* 4 (4) (Springer Verlag London). ISSN 0949 2054.
16. Starner, T., Schiele, B. and Pentland, A. Visual Contextual Awareness in Wearable Computing. *Proceedings of the Int. Symposium on Wearable Computing (ISWC'98)*, Pittsburg, 20 Oct. 1998, pp. 50-57.