Mobile Networks and Applications 7, 341–351, 2002

Multi-Sensor Context-Awareness in Mobile Devices and Smart Artifacts

HANS W. GELLERSEN *, ALBRECHT SCHMIDT and MICHAEL BEIGL
TecO, University of Karlsruhe, Vinc.-Priessnitz-Str. 1, 76131 Karlsruhe, Germany

Abstract. The use of context in mobile devices is receiving increasing attention in mobile and ubiquitous computing research. In this article we consider how to augment mobile devices with awareness of their environment and situation as context. Most work to date has been based on integration of generic context sensors, in particular for location and visual context. We propose a different approach based on integration of multiple diverse sensors for awareness of situational context that can not be inferred from location, and targeted at mobile device platforms that typically do not permit processing of visual context. We have investigated multi-sensor context-awareness in a series of projects, and report experience from development of a number of device prototypes. These include development of an awareness module for augmentation of a mobile phone, of the Mediacup exemplifying context-enabled everyday artifacts, and of the Smart-Its platform for aware mobile devices. The prototypes have been explored in various applications to validate the multi-sensor approach to awareness, and to develop new perspectives of how embedded context-awareness can be applied in mobile and ubiquitous computing.

Keywords: context-awareness, sensor integration, mobile computing, ubiquitous computing

1. Introduction

Mobile computing technology has facilitated the use of computer-based devices in different and changing settings. In this article we consider augmentation of mobile devices with awareness of their environment and situation as context. More specifically we look at sensor integration in mobile devices, at abstraction of sensor data to more general characterizations of device situation, and at use of such context to improve user interaction and to support new types of application. Our work is grounded in ubiquitous computing research, and the devices we consider are mobile and computer-based in the widest sense. Specifically, they include smart devices and artifacts in which computing is secondary to a primary purpose, for instance mobile phones and everyday objects augmented with embedded computing.

"Context is what surrounds", and in mobile and ubiquitous computing the term is primarily used in reference to the physical world that surrounds the use of a mobile device. This has also been referred to as physical context, to distinguish it from other types of context in mobile computing, such as conditions in the surrounding systems and network infrastructure [3,7]. Other suggested terminology includes *situational context*, stressing the aim to reflect the situated nature of mobile device use beyond the capture of physical conditions [26,29]. Some confusion occasionally arises from the use of the term context at different levels of abstraction to denote (i) the real world situation surrounding a device, (ii) an aspect of a situation, such as location, and (iii) a specific occurrence of an aspect, such as a specific place. In our work we use the term situation in reference to the real world and the term context for a model of a situation acquired by means of sensing and sensor data processing. We also use the term situational context for multi-faceted characterizations of a situation that typically require substantial analysis and fusion of data from individual sensors.

The use of context in mobile devices is receiving considerable attention in various fields of research including mobile computing, wearable computing, augmented reality, ubiquitous computing and human–computer interaction. A general motivation is that context-awareness can serve to compensate for the abstraction that is required in the first place to make systems accessible in changing environments and situations. The actual utility of context-awareness in mobile systems has been demonstrated in a wide range of application examples, in obvious domains such as fieldwork [17,19] and tourism [3,6,18], as well as in emerging areas like affective computing [15,20]. It has also been shown that context can be applied for different tasks within a mobile device. In overall device operation, context can for example be exploited for context-sensitive resource and power management. In applications executed on mobile devices, context be applied to enable adaptive applications and explicitly context-based services. In the mobile device user interface, context can be used to facilitate a shift from explicit user-driven to implicit context-driven interaction [27,35].

The focus of this article is on how to facilitate context-awareness in mobile devices. Context as considered here originates in the physical surroundings of a device, is captured through sensors, and finally results from analysis of sensor data. Devices may have direct or indirect awareness of context. In the case of indirect awareness, the entire sensing and processing occurs in the infrastructure while the mobile device obtains its context by means of communication. In

---

* Corresponding author, current address: Lancaster University, Computing Department, Lancaster LA1 4YR, UK.

contrast, a device has direct awareness if it is able to obtain context autonomously, (more or less) independent of any infrastructure. We will not further consider indirect awareness in this article, as the reliance on sensing infrastructure limits mobile device use to specifically designed smart environments. Another problem of indirect awareness is the dislocation of context acquisition from context use. For instance, if a device embodies adaptation of its behavior to its temperature, then it is problematic to rely on temperature readings received from the environment.

Mobile devices with direct awareness by definition embody one or many sensors, and models or algorithms for computation of more abstract context from raw sensor data. Most research into context-aware mobile devices has considered the use of single powerful sensors, specifically position sensors and cameras. Position sensors provide access to location which has been shown to be a particularly useful context. Often it is not the position itself that constitutes immediately useful context but additional information that can be inferred, for instance resources registered at a location. Cameras similarly provide access to potentially rich information that can be derived by means of feature extraction and video analysis. While positional context and visual context are powerful for augmentation of devices with some awareness of their situation, they both also have distinct shortcomings. Position is a static description of an environment and does not capture dynamic aspects of a situation. Its usefulness as context also largely depends on precaptured knowledge about locations. Vision on the other hand can be employed to capture activity and other dynamic aspects, but extraction of specific context is computationally expensive and problematic in mobile and uncontrolled environments.

In this article we discuss multi-sensor context-awareness as an alternative approach toward aware mobile devices. In this approach, the single powerful sensor is replaced by a collection of diverse simple sensors, and context is derived from multi-sensor data. We first introduced this approach in the European project TEA on mobile situation awareness, with the aim to provide comparatively cheap technology both with respect to processing requirements and component cost [25]. In TEA and in follow-up projects, we have focused on the integration of deliberately simple and cheap sensors, and on comparatively inexpensive processing of multi-sensor data to obtain context, to meet the constraints typically found in mobile and embedded devices. In this series of projects, we have investigated the approach in different settings to develop an understanding of how multi-sensor context-awareness can be applied in mobile and ubiquitous computing. A variety of device prototypes have emerged from this research and we will discuss three of these in this article: the TEA awareness module used for augmentation of mobile phones; the Mediacup, a coffee cup with embedded awareness of its own state; and the Smart-Its device for augmentation of everyday artifacts. The investigation of these prototypes provides new perspectives of aware mobile devices and insights into the kind of context that can be obtained from multi-sensor embedding. It further contributes an exploration of architectures for embedded awareness, and a rich design experience in trading off context-awareness against energy and memory constraints on micro-controller platforms.

Following this introduction to our work we will reflect on related research in section 2. In doing so, we will keep the focus on mobile devices with direct awareness and look further into both positional and visual context-awareness. In addition we will discuss related work on use of other sensors and types of context in mobile devices, and in particular other research on integration of multiple diverse sensors. In section 3, we will return to the discussion of multi-sensor context-awareness, and provide an overview of our program of research on this concept. This will be followed by a section each on a series of projects, starting with the TEA experience in section 4, followed by section 5 on the Mediacup development, and section 6 on the design of the Smart-Its device. In the final section we will summarize and discuss our overall experience of mobile device augmentation with context-awareness, and point out issues and directions for further research.

2. Background and related work

In distributed computing, the interest in context has emerged in parallel with mobility support for users and devices. Groundbreaking contributions have been the Active Badge system which integrated a positioning technique with a distributed computing infrastructure [14,33], and the Ubiquitous Computing vision that pointed out the importance of location and context for the next era of computing [35]. These contributions were also underlying the ParcTab experiment, an early investigation into context-aware computing in which palm-size personal devices were augmented with locality for mobile access to location-based services [34]. One of the outcomes of the ParcTab work was a taxonomy of context-aware applications that has inspired much further work [24]. However, these pioneering projects all employed indirect awareness with sensors located in the infrastructure, listening for beacons sent out from the mobile devices. In the case of the ParcTab, location information is in fact a by-product of the cell-based communication infrastructure in which the base stations double as device locators. Many other context-aware mobile systems have likewise used cell-of-origin as context for location-based services, for example the GUIDE system deployed in Lancaster on the basis of a wireless LAN [6].

More recent work has increasingly considered the embedding of direct awareness in mobile devices. This has been boosted by rapid advances in sensor technologies, such as piezo-materials, VLSI video, optical gyros and MEMS (Micro Electro-Mechanical Systems) [22]. The trend is leading to sensors in ever smaller packages that increasingly integrate feature extraction in the sensor hardware, enabling powerful sensing at very low cost. In the following sections we will discuss related research on integration of sensors and context-awareness in mobile devices.

2.1. Position sensing and location-based context

Location as position or area in space is a context of particular importance, and has received more attention in mobile computing than any other type of context. Like time, spatial location is an inherent attribute of other types of physical context, and often used implicitly for filtering nearby observations which are more relevant as context from remote and hence less relevant observations [11]. Location is a well understood context and powerful models are available to support querying and processing in a variety of ways. For example, geometric models support location representations on which simple arithmetic can be applied, while symbolic models support the use of set theoretical expressions, such as *being contained in*. However, it is typically not a location as such that is of interest in location-aware mobile computing. Instead, "location becomes a useful indexing device" from which to infer further context [7, p. 292].

Most available position sensing technologies are based on sensor infrastructure rather than sensor integration into the mobile device. This is in particular the case for indoor positioning systems, which are typically based on small radio or infrared cells, or on sensor arrays in the environment for higher accuracy. However, alternatives have been proposed, for example indoor positioning based on visual context (see below). For outdoor positioning, the situation is different. In the GPS system (Global Positioning System) it is the infrastructure that sends out signals while the sensor is located in the client device. GPS sensors have become available in very small package sizes enabling their integration in mobile devices. A variety of projects have shown how integration of GPS with mobile computing devices can support new kinds of application. For example, in the stick-e-notes system GPS is used in conjunction with palm computers to support context-based access to shared notes in fieldwork [2,19].

In contrast to approaches based on positioning technology, one of our explicit objectives is to achieve awareness of context that can not be inferred from location. With the use of diverse sensors we further aim at more direct acquisition of multi-faceted information from the local environment of a device.

2.2. Auditory and visual context

Computer vision is by its very definition concerned with the capture of aspects of the physical world for use in computers. The field is rooted in artificial intelligence and the tradition of building computer models of human capabilities. However, with decreasing cost and size of camera modules and emergence of wearable computers, computer vision is now considered as technology option for mobile context-aware systems. Visual context typically constitutes a rich resource from which more specific context can be derived by means of video analysis and feature extraction. For example, Visual Augmented Memory is a wearable system that embodies a camera and face recognition software to obtain highly specific visual context [9]. In other wearable systems, vision has been used to determine location and to support navigation tasks [31].

More closely related to our work though is that of Clarkson et al. who investigated the recognition of user situations from a wearable camera and microphone [4,5]. The commonality with our work is in the focus on situation rather than location. However Clarkson's work is focused on the recognition technology whereas we are interested in the wider picture of how awareness technology can be integrated in mobile devices, and in particular in less powerful devices than wearable high-end computers.

2.3. Integration of other sensors in mobile devices

Besides technologies for positioning and vision, a range of other sensors have been investigated for augmentation of mobile devices. In many cases this is aimed to facilitate capture of very specific context as input to particular applications. An example is the StartleCam system, based on a wearable computer augmented with bio-sensors for recognition of extreme user situations [15]. Another perspective on sensor integration in mobile interactive devices is to extend the user interface with new interaction techniques. For example, Rekimoto added tilt sensors to a palm computer for one-handed operations [21], and in a similar way we have integrated sensors in a handheld computer for automatic adaptation of display orientation (landscape/portrait) to device orientation [26]. In this line of work, sensors are used for new user interface techniques, whereas our focus in this paper is on sensor integration for awareness of device situation.

2.4. Integration of diverse sensors

An important aspect of our approach is the integration of different kinds of sensor. This aspect has also been investigated in the Smart Badge, a device inspired by the Active Badge but in addition equipped with a large variety of sensors to capture context beyond position [30]. The Smart Badge provides integrated access to diverse types of sensor data, with the research focus on device architecture rather than sensor fusion and inference of generic situational context. The sensor badge is a similar device, however focused on recognition of specific user activities from movement data [8]. Both Smart Badge and Sensor Badge particularly relate to our early work on the TEA awareness module, in that they are conceived as peripheral components to provide other mobile devices with context. The integration of diverse sensors has also been considered by Golding and Lesh for an indoor location technique that does not require any infrastructure [13]. The research prototype integrated sensors for motion, orientation, light, and temperature, and relates to our own work with its emphasis on small, lightweight, low-power and cheap components.

3. Multi-sensor based context-awareness

The combination of comparatively simple sensors is an interesting alternative to the use of single powerful sensors as in position- and vision-based systems. The combination of

Table 1
Real world situations related to sensor data (adapted from [29]).

Situation	Sensor data
User sleeps	It is dark, room temperature, silent, type of location is indoors, time is "night-time", user is horizontal, specific motion pattern, absolute position is stable.
User is watching TV	Light level/color is changing, certain audio level (not silent), room temperature, type of location is indoors, user is mainly stationary.
User is cycling	Location type is outdoors, user is sitting, specific motion pattern of legs, absolute position is changing.

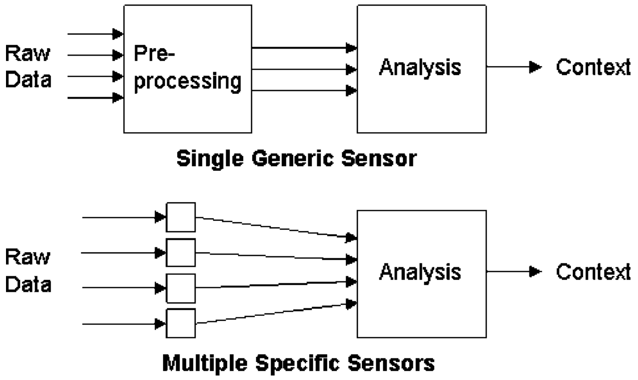

Figure 1. Use of a single generic sensor versus multiple specific sensors for context-awareness (adapted from Van Laerhoven et al. [32]).

multiple diverse sensors that individually capture just a small aspect of an environment may result in a total picture that better characterizes a situation than location- or vision-based context. The rationale for our approach is to advance beyond location-based systems to achieve awareness of context that can not be inferred from position, and we see the diversity of sensors as key. Table 1 lists a few examples of how situations and data from different sensors may relate.

Our work is further motivated by the objective to devise awareness technology that can be embedded in devices with severe resource limitations, and even in simple everyday artifacts. Through the use of multiple diverse sensors we expect to gain access to rich data from which useful context can be inferred with comparatively little computation. As there are multiple sensors, they each only have to contribute a part to the whole picture. This means that preprocessing of sensor data will be more focused than for example in vision, where substantially more processing is required to derive information from just a single data source. Figure 1 illustrates the difference between use of one generic sensor and use of multiple simple sensors.

Over the last years, we have investigated our approach in a series of projects to develop an understanding of its utility for mobile and ubiquitous computing. Our research approach is to build fully integrated sensor-enabled device prototypes and to deploy them in mobile environments and where possible in everyday use. This stresses our interest in advancing both the understanding of how mobile devices can be augmented with awareness, and how aware mobile devices can be applied in the real world. This approach, characteristic for ubiquitous computing research, is aimed at collection of overall design and use experience, with lesser attention to aspects such as optimization of specific sensor fusion and context recognition techniques.

The projects that we will discuss in the remainder of the article cover the development of the TEA awareness module, the Mediacup experience, and our current work on the Smart-Its platform. Each of these projects explores different perspectives of how multi-sensor context-awareness may be applied in conjunction with mobile devices and artifacts:

- TEA – Technology Enabling Awareness: in this initial project we investigated the implementation of multi-sensor context-awareness in a self-contained device that would be available as peripheral or plug-in for mobile host devices. The general application perspective was to supply situational context to a mobile host device to improve the device's service to its user. We explored this perspective in the application domain of mobile telephony.
- Mediacup: in this project we looked at how non-computational artifacts can be augmented with awareness technology. The application perspective is entirely different from that underlying TEA. Awareness is not employed to improve the immediate function of the augmented device, but to create a digital presence for it. New functionality is not expected to emerge in the device itself but in the surrounding system environment.
- Smart-Its: this recently started project is aimed at moving beyond the study of design examples such as Mediacup, toward platforms for aware mobile devices. With this project we also shift our attention from individual devices to ad hoc networking of aware devices and to scenarios of collective awareness.

A general question underlying the research in these projects is what kind of situational context can be obtained from multi-sensor context-awareness. Closely related is the question of how data from multiple sensors can be related effectively to situational context. In the following three sections we will report findings from the individual projects, to be followed by a discussion that will summarize what we have learned so far, drawing some conclusions for further investigation of context-awareness in mobile devices.

4. The TEA context-awareness module

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 user 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 of TEA is to address the kind of context that cannot be derived from location information at all, for example situations that can occur anywhere. Another specific issue investigated in TEA is sensor fusion. The aim is to derive more context information from a group of sensors than the sum of context derived from individual sensors.

4.1. TEA architecture

TEA is based on a layered architecture for sensor-based computation of context as illustrated in figure 2, 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-ended collection of sensors. The data supplied by sensors can be very different, ranging form slow sensors that supply scalars (e.g., temperature) to fast and complex sensors that provide larger volume data (e.g., microphone). The update time can also vary largely 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 serves to reduce the data volume independent of any specific application, and has also been referred to as “cooking the sensors” [13]. 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 idea 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 slightly 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. This is an important requirement in context-aware systems and has motivated the development of architectures such as the Context Toolkit [23].

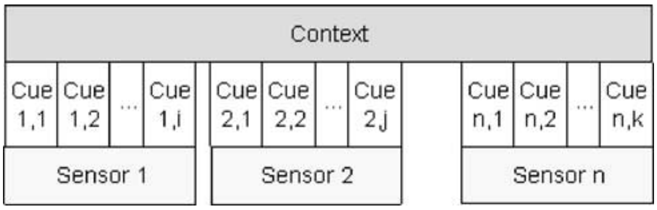

Figure 2. TEA is based on a layered architecture for abstraction from raw sensor data to multi-sensor based context.

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 sensors is combined for calculation of 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 calculating context from cues; rule-based algorithms, statistical methods and neural networks may for instance be used. Conceptually, context is calculated from all available cues. In a rule set however, cues known to be irrelevant may simply be neglected, and in neural networks their weight would be reduced accordingly. The mapping from cues to context may be explicit, for instance when certain cues are known to be relevant indicators of a specific context, or implicit in the result of supervised or unsupervised learning.

4.2. Initial exploration of the approach

To study the TEA approach, we have developed two generations of prototype devices and used them for analysis 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 first generation TEA module 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 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 selected (e.g., user is walking, user is in a conversation, other people are around, user is driving a car, etc.) for data collection at three different sites. The data was 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 many different situations. Over the whole period of time, raw sensor data was recorded for later analysis. The analysis was aimed to identify clusters corresponding to situations that occurred during recording time, e.g., the user is sitting at their desk, walking over to a colleague, chatting, walking back, taking a phone call etc. The aim was further to identify the sensors most relevant to situations, and to develop clustering algorithms.

4.3. Implementation of TEA in a self-contained awareness device

The initial exploration of sensors and their contribution to awareness of real-world situations served to inform development of the second generation device optimized for smaller size and integration with a mobile phone (see figure 3). 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 about 200 Byte RAM available.

The cue extraction algorithms 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 crossings 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 indicative of the noise in the 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.

Figure 3. The TEA awareness device was implemented as self-contained module that can be connected via serial line to mobile host devices. The device has been applied for augmentation of mobile phones with context-awareness.

The prototype is independent of any specific host and at times has been used in conjunction with a palmtop computer, a wearable computer and mobile phones, connected via the serial interface. Primarily however the prototype has been studied in the area of mobile telephony.

4.4. Application in mobile telephony

State of the art mobile phones support so-called profiles which specify settings, such as notification mode, input and output modality, and reaction to incoming messages and calls. Users can define profiles for different situations and specify behavior desired in those situations. In several experiments we have applied the TEA device to automate profile activation, for instance for the situations in-hand, on-table, in-pocket, and outdoors. The device recognized these situations with a certainty of at 87% (and higher, depending on situation), however with a delay of up to 30 s [29]. It has to be noted that the experiments were conducted to assess overall feasibility. More attention to the recognition algorithms would most likely yield still better context prediction. See for instance related work on the use of neural networks for real-time analysis of multi-sensor data [32].

An interesting application domain for context-aware mobile phones as enabled by TEA is the sharing of context between caller and callee in inter-personal communication. For a caller, context may be helpful 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 ring you back later" – a common question in accepting a call). To demonstrate context-enhanced mobile telecommunication, we have used WAP technology to build the Context-Call system. In this system, a call is initiated in the usual way by selecting a number or phone book entry. Instead of directly establishing a connection, the system first looks up the context of the callee, and reports this back 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 this example for interaction in context is provided in [28].

4.5. 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 approach to obtain such context. We have used the approach for obtaining strictly location-independent context that can not be derived from location information. An important finding though is that it is difficult to think of useful context independently of particular application domains, whereas we had initially expected to find widely applicable situational context. As for sensor fusion, our analysis of collected multi-sensor data showed that with our approach context can be derived beyond the sum of context obtained separately from multiple sensors. This initial experience is valuable, but the understanding of sensor fusion will require further research. However, the layered architecture for context computation appears to be a useful foundation. 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 audio, motion, and light sensors contribute to context-awareness in most settings while the utility of other sensors is more dependent on the target application. In addition we found that perception can be improved by using not just diverse sensors but also multiple sensors of the same kind, for example multiple microphones and light sensors with different orientation.

5. The Mediacup: embedding awareness to create Active Artifacts

The Mediacup project is like TEA concerned with multi-sensor integration but introduces a completely different perspective of how to apply awareness technology. The general thinking underlying the integration of awareness into devices is to improve their performance, by enabling them to adapt their resources to an environment, to offer additional functionality, or to improve human–computer interaction. With the Mediacup project we have introduced a new application theme, proposing awareness technology to augment non-computational artifacts with a digital presence to create Active Artifacts [1]. This can be viewed as complementary to the main trend in context-aware mobile devices, in which computer-based devices reach out into the real world through sensors. Activate Artifacts on the other hand are real world objects that reach out into the world of computation [12]. The framework that we have developed for Active Artifacts takes us back to the concern of this article, augmentation of devices and artifacts with sensors and context:

- Autonomous awareness: Active Artifacts have sensors and perception methods embedded to assess their own state and situation independent of any infrastructure.
- Context sharing: Active Artifacts are augmented with communication to make their context available within a local *region of impact.*
- Context use: any application in the local environment can use the information supplied by Active Artifacts as context.

It is important to note that we think of Active Artifacts as everyday objects that retain their original purpose, appearance, and functionality. New functionality is instead expected to emerge in the environment, based on context of a possibly large number of networked artifacts. For an initial exploration of the idea of Active Artifacts we have developed the Mediacup, a coffee cup with technology embedded transparently in a detachable base [10].

5.1. Mediacup – Awareness embedded in coffee cups

The Mediacup is an ordinary coffee mug that is augmented with hardware and software for sensing, processing and sharing context. The Mediacup has evolved through several design iterations, resulting in the implementation shown in figure 4. The design challenge was to integrate multi-sensor awareness technology and communication without compromising properties and use of the cup, for instance avoiding noticeable changes in shape, size and weight.

The Mediacup hardware is based on circular board design to be fitted into the base of the cup. The components integrated on the Mediacup board are controlled by a PIC16F84 microcontroller with 1 MHz, 15K memory for code, 384 byte RAM for data. The sensors include a digital temperature sensor, three ball switches to detect movement and a switch that closes when the cup is placed on a surface. The board further contains an infrared diode for communication, and two large capacitors with 1 F each as energy source. The capacitors are charged wirelessly using a resonant circuit (20 kHz), embedded in the cup's saucer. The board including all components is only 3 mm high.

Figure 4. The Mediacup is an ordinary coffee cup with sensors, processing and communication embedded in the base.

The Mediacup software controls acquisition of raw data from sensors, and performs simple data analysis to compute context. Both acquisition and processing are designed to minimize energy consumption. For example, movement is a parameter that can change fast and frequently and requires sensor readings about every 20 ms but most of the time a cup will actually not be moving at all. Hence, we chose motion detection technology that does not require continued polling but instead triggers an interrupt in the board processor when movement occurs. Detected movement is recorded as event, and a short history of such events is used in a rule-based heuristic to compute context that relates to the use of the cup. Four movement contexts are distinguished: *cup is stationary, drinking out of the cup*, *cup is played with,* and *cup is carried around.* 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 use-related context: *filled up*, *cooled off*, and *current temperature.*

Mediacups broadcast their context together with their IP address as unique ID every two seconds using the upward facing infrared diode. 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 Mediacup environment, i.e. in three rooms and a larger lab space in our office environment. The transceivers are based on the Hewlett-Packard HSDL 1001 IrDA chip and have a footprint of about 1.5 m^2. They are connected through a CAN bus (control area network) and via a gateway to the LAN (local area network) on which collected context is broadcast in UDP packets.

5.2. Experience from design and use

Like TEA, the Mediacup project served to gather extensive hands-on experience with building multi-sensor context-aware devices. While TEA primarily provided insights into issues surrounding sensor fusion and context architecture, the Mediacup provides more insights into embedded systems issues, and the specific challenges in design of transparent technology. Not surprisingly, the embedding of technology in passive artifacts that do not contain any energy source raises issues of power management. In the Mediacup, energy-related concerns have influenced a wide range of design decisions:

- Processing: the microcontroller was set to run with a reduced clock speed of only 1 MHz, reducing power consumption to below 2 mA at 5.5 V in processing mode. Whenever possible, the processor is put in sleep mode with power consumption below 1 μA.
- Motion detection: an earlier version of the Mediacup used the ADXL202 accelerometer (ADXL202) which had to be polled actively. To reduce power consumption, it was replaced by three ball switches connected to the external interrupt pins of the microcontroller. This enables the microcontroller to remain longer in sleep mode when no movement occurs, enabling sleep mode for more than 99% of the time without losing any movement information.
- Temperature: a Dallas DS1621 chip was used to measure temperature (−55 to +125°C) as it consumes only 1 μA in standby mode and 400 μA during very short reading cycles.
- Recharging: nobody would be prepared to change batteries in a coffee cup, or to plug a coffee cup into the mains for a recharge. Therefore, we built rechargeable capacitors into the cup, and a resonant circuit for wireless charging into a stationary plate reminiscent of a saucer. Whenever a user puts their cup on the saucer it gets recharged, and 15 min charge is usually sufficient for about 12 hours cup operation.
- Communication: a low-powered 5 mm infrared diode is used (HSDL4420). The status of the cup is communicated every 2 s to the environment using IrDA-PHY physical layer coding. The IrDA coding is done in software on the microcontroller to have a component less on board. The maximal data rate feasible in software, 19.2 kbps, is used to minimize 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 do not compromise its common use, it does not suffice to minimize and hide the technology. For example, the need to have free line of sight between artifact and transceiver infrastructure should be transparent to the user, which requires careful design of both artifact and infrastructure. Another example that came up with use experience with an early battery-powered prototype was that power provision needs to be completely transparent. We observed that users would not care to check the battery, and to make sure they were recharged. This 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 immediately visible to the user.

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 applications. The common view is to consider sensors as periphery, and applications as the place to make sense of collected data. In the Active Artifacts model that we explored, 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 is also challenged: 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, context becomes a common resource in the environment, and may enable new functionality in other devices and artifacts. In the Mediacup environment for instance, context from coffee cups was not directed at specific applications to start with but once the cups were in place this inspired augmentation of a variety of other artifacts to use Mediacup context. For instance, one user augmented his wrist-worn PC to beep a warning if a cup was picked up and its content was sensed to be too hot. Another use was built into digital door plates, which originally were built to leave notes at doors. In these door plates Mediacup information became used to help infer whether a meeting was in place in the room. These are examples of simple uses but we consider them indicative of the kind of context-based services that can emerge once a framework for collecting and providing context information is in place.

In regard of multi-sensor context-awareness, the Mediacup is an obvious example of tapping into context information that can not be inferred from location. More interesting however is that with the Mediacup we have moved away from the idea of generic situational context to focused context. The context obtained in the Mediacup is focused on capturing the use of the cup, and the level of abstraction relates to how people would perceive a cup and its state ("gone cold" rather than "temperature = 20°C"). The focus on a clear cut domain in which to observe context of course helps frame and apply awareness technology, while the idea of "human readable context" may be helpful in establishing a common level of abstraction to talk about context.

6. Smart-Its: Toward post hoc augmentation of artefacts with sensing and context

The Mediacup work has been inspiring and instructive in pursuit of a better understanding of the role awareness can have in mobile and ubiquitous computing. However, to advance awareness technology for mobile devices we felt a need to move on from single design examples to consideration of platforms. In our recently initiated project Smart-Its we have begun to work on a range of small, embedded devices as platforms for augmentation and interconnection of artifacts. These devices, *Smart-Its*, in general integrate sensing, processing and communication with variations in perceptual and computational capability. Sensors and perception techniques are integrated to facilitate autonomous awareness of an artifact's context, independent of infrastructure, and wireless networking facilitates context sharing and other communication. In comparison to the TEA module, Smart-Its are not thought of as add-on to some host device, but as devices in their own right. In relation to the Mediacup, Smart-Its are conceived as generic devices, designed for flexibility and customizability.

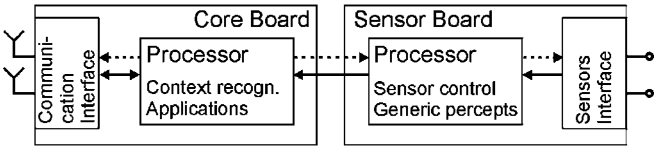

Figure 5. Smart-Its Device Architecture with separate boards for sensors and perception, and for applications and communication.

The Smart-Its devices are based on a modular design with two boards, separating the sensor unit from the core unit. The main components and the data and control flow on the device are illustrated in figure 5. Acquisition of data is allocated on the sensor unit, with a dedicated processor for sensor control and extraction of generic features. Overall device control, application-specific processing, and communication with other Smart-Its is allocated on the core unit. Application-specific processing might for example be computation of artifact-specific context or any other sensor abstraction. The communication interface may support different kinds of network. Generally we assume that all Smart-Its communicate over a shared wireless medium, but some Smart-Its may support additional networks to implement gateways.

One of the first Smart-Its device prototypes that we implemented so far is shown in figure 6. The two boards have a size of about 4 × 5 cm, and are mounted on top of each other. Each board is based on a PIC micro-controller as processing unit. The sensor board is further equipped with a two-axis accelerometer. Another sensor, a simple ball switch, is integrated on the core board and directly connected to an interrupt on the core processor. This enables the device to go into an energy preserving mode when no movement occurs, and to wake up instantly on movement. The core board is further equipped with an RFM interface for wireless communication. The communication is based on detection of Smart-Its within sending range. All Smart-Its use RFM as shared broadcast medium, based on a simple carrier sense collision avoidance protocol (CS/CA).

Enabled by the Smart-Its platform we aim to investigate ad hoc networking and sharing of context, and collective perception based on dynamic sensor fusion across a number of devices. With the Smart-Its Friends system we have built a first demonstrator to illustrate the utility of context sharing [16]. In this system, Smart-Its devices broadcast their movement patterns in their local environment, and if a device

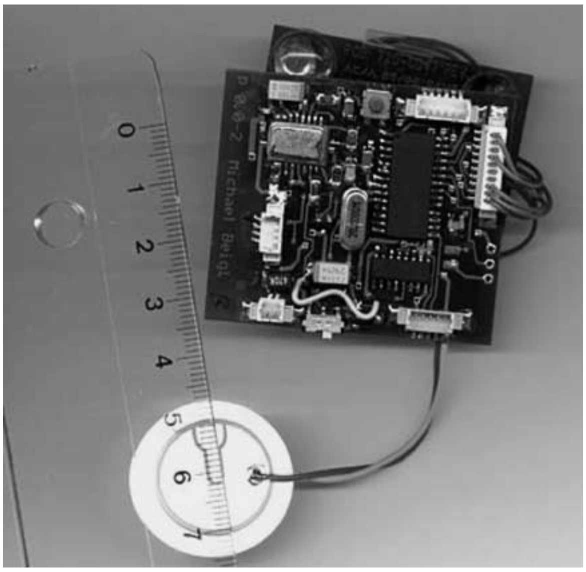

Figure 6. Prototype implementation of the Smart-Its device (scale in centimeters).

receives a movement patterns similar to its own, it establishes an application-level connection. At the user interface, this protocol enables a user to connect two devices by imposing the same movement, i.e. by taking them in one hand to give them a brief shake. While this in itself is a compelling interface technique, it is also indicative of new kinds of applications that may emerge on the basis of context communication.

7. 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 artefacts. 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. In the recently started Smart-Its effort we build on this experience and begin to investigate platforms for aware mobile devices, ad hoc context sharing, and collective awareness across ad hoc connected devices.

We have shown that integration of diverse sensors is a viable approach to obtain context representing real-world situations, and context that captures interaction with everyday artifacts. We have investigated 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. However, while it will continue to be important to work toward application-independent taxonomies of context, our experience also suggests that it is useful to think of “context in context”.

Our work to date was not specifically focused 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. This is among the objectives of the Smart-Its effort, and to some extent reflected in the design for modularity of the Smart-Its device architecture.

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. Our experience highlights substantial challenges for perception techniques to perform in low-end computing environment. In our work, in particular in the Mediacup project, we have carefully crafted sensor control to meet requirements. An important research direction in multi-sensor context-awareness will be perception techniques that can adapt dynamically to changes in sensor input. For instance, we envision scalable perception techniques that perform robustly in conjunction with sensors that are dynamically powered on and off.

Finally, we seek to contribute further to the development of new perspectives of how context-awareness can be applied in mobile and ubiquitous computing. 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 artifacts deeply embedded in everyday life.

Acknowledgements

The Commission of the European Union funded the TEA project in the ESPRIT program (contract 26900), and is now funding work in Smart-Its as part of the initiative The Disappearing Computer in the IST program (contract 25428). We also wish to acknowledge the indirect contribution of all project partners, colleagues, and students who at different times have been involved with some of the work presented.

References

[1] M. Beigl, H.W. Gellersen and A. Schmidt, Mediacups: Experience with design and use of computer-augmented everyday objects, Computer Networks 35(4) (March 2001) 401–409.

[2] P.J. Brown, P.D. Bovey and X. Chen, Context-aware applications: From the laboratory to the marketplace, IEEE Personal Communications 4(5) (1997) 58–64.

[3] K. Cheverst, N. Davies, K. Mitchell and A. Friday, The role of connectivity in supporting context-sensitive applications, in: *Proceedings of the First International Symposium on Handheld and Ubiquitous Computing (HUC99)*, Karlsruhe, Germany (September 1999), Lecture Notes in Computer Science No. 1707 (Springer-Verlag, 1999).

[4] B. Clarkson, N. Sawhney and A. Pentland, Auditory context awareness via wearable computing, in: *Proceedings of the 1998 Workshop on Perceptual User Interfaces (PUI'98)*, San Francisco, CA (November 1998).

[5] B. Clarkson, K. Mase and A. Pentland, Recognizing user context via wearable sensors, in: *Proceedings of the IEEE International Symposium on Wearable Computing (ISWC'00)*, Atlanta, GA (October 2000) pp. 69–76.

[6] N. Davies, K. Cheverst, K. Mitchell and A. Friday, Caches in the air: Disseminating information in the guide system, in: *Proceedings of the*

Second IEEE Workshop on Mobile Computing Systems and Applications (WMCSA'99), New Orleans, LA (February 1999).

[7] A. Dix, T. Rodden, N. Davies, J. Trevor, A. Friday and K. Palfreyman, Exploiting space and location as a design framework for interactive mobile systems, ACM Transactions on Computer–Human Interaction (TOCHI) 7(3) (September 2000) 285–321.

[8] J. Farringdon, A.J. Moore, N. Tilbury, J. Chruch and P.D. Biemond, Wearable sensor badge & sensor jacket for context awareness, in: *Proceedings of the International Symposium on Wearable Computing (ISWC'99)*, San Francisco, CA (November 1999) pp. 107–113.

[9] J. Farringdon and V. Oni, Visual Augmented Memory (VAM), in: *Proceedings of the IEEE International Symposium on Wearable Computing (ISWC'00)*, Atlanta, GA (October 2000) pp. 167–168.

[10] H.W. Gellersen, M. Beigl and H. Krull, The Mediacup: Awareness technology embedded in an everyday object, in: *Proceedings of the First International Symposium on Handheld and Ubiquitous Computing (HUC99)*, Karlsruhe, Germany (September 1999), Lecture Notes in Computer Science No. 1707 (Springer-Verlag, 1999) pp. 308–310.

[11] H.W. Gellersen, M. Beigl and A. Schmidt, Sensor-based context-awareness for situated computing, in: *Proceedings of Workshop on Software Engineering for Wearable and Pervasive Computing*, Limerick, Ireland (June 2000).

[12] H.W. Gellersen, Ubiquitous computing: Where computation and artefacts meet, Informatik/Informatique 2001(4), Swiss Informatics Societies SVI/FSI (October 2001).

[13] A. Golding and N. Lesh, Indoor navigation using a diverse set of cheap wearable sensors, in: *Proceedings of the IEEE International Symposium on Wearable Computing (ISWC'99)*, San Francisco, CA (October 1999) pp. 29–36.

[14] A. Harter and A. Hopper, A distributed location system for the active office, IEEE Network 8(1) (1994).

[15] J. Healey and R. Picard, StartleCam: A cybernetic wearable camera, in: *Proceedings of the International Symposium on Wearable Computing*, Pittsburgh, PA (October 1998) pp. 42–49.

[16] L.E. Holmquist, F. Mattern, B. Schiele, P. Alahuhta, M. Beigl and H.W. Gellersen, Smart-Its friends: A technique for users to easily establish connections between smart artefacts, in: *Proceedings of UBICOMP 2001*, Atlanta, GA (September 2001).

[17] G. Kortuem, M. Bauer and Z. Segall, NETMAN: the design of a collaborative wearable computer system, Mobile Networks and Applications 4(1) (March 1999) 49–58.

[18] S. Long, R. Kooper, G.D. Abowd and C.G. Atkeson, Rapid prototyping of mobile context-aware applications: the Cyberguide case study, in: *Proceedings of the Second International Conference on Mobile Computing and Networking (MOBICOM'96)*, Rye, NY (November 1996) pp. 97–107.

[19] J. Pascoe, D.R. Morse and N.S. Ryan, Developing personal technology for the field, Personal Technologies 2(1) (March 1998) 28–36.

[20] R. Picard and J. Healey, Affective wearables, Personal Technologies 1(4) (December 1997) 231–240.

[21] J. Rekimoto, Tilting operations for small screen interfaces, in: *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST'96)*, pp. 167–168.

[22] P. Saffo, Sensors: The next wave of InfoTech innovation, 1997 ten-year forecast, Institute for the Future (1977).

[23] D. Salber, A.K. Dey and G.D. Abowd, The Context Toolkit: Aiding the development of context-enabled applications, in: *Proceedings of the Conference on Human Factors in Computer Systems (CHI'99)*, Pittsburgh, PA (May 1999) pp. 434–441.

[24] B.N. Schilit, N.I. Adams and R. Want, Context-aware computing applications, in: *Proceedings of the IEEE Workshop on Mobile Computing Systems and Applications*, Santa Cruz, CA (December 1994).

[25] A. Schmidt, K.A. Aidoo, A. Takaluoma, U. Tuomela, K. Van Laerhoven and W. Van de Velde, Advanced interaction in context, in: *Proceedings of the First International Symposium on Handheld and Ubiquitous Computing (HUC99)*, Karlsruhe, Germany (September 1999), Lecture Notes in Computer Science No. 1707 (Springer-Verlag, 1999) pp. 89–101.

[26] A. Schmidt, M. Beigl and H.W. Gellersen, There is more to context than location, Computer & Graphics 23(6) (December 1999) 893–901.

[27] A. Schmidt, Implicit human–computer interaction through context, Personal Technologies 4(2–3) (June 2000) 191–199.

[28] A. Schmidt, A. Takaluoma and J. Mäntyjärvi, Context-aware telephony over WAP, Personal Technologies 4(4) (December 2000).

[29] A. Schmidt and K. Van Laerhoven, How to build smart appliances?, IEEE Personal Communications 8(4) (August 2001).

[30] M.T. Smith, SmartCards: Integrating for portable complexity, IEEE Computer (August 1998) 110–112, 115.

[31] T. Starner, B. Schiele and A. Pentland, Visual contextual awareness in wearable computing, in: *Proceedings of the International Symposium on Wearable Computing (ISWC'98),* Pittsburgh, PA (October 1998) pp. 50–57.

[32] K. Van Laerhoven, K. Aidoo and S. Lowette, Real-time analysis of data from many sensors with neural networks, in: *Proceedings of the International Symposium on Wearable Computing (ISWC 2001)*, Zürich, Switzerland (October 2001).

[33] R. Want, A. Hopper, V. Falcao and J. Gibbons, The Active Badge location system, ACM Transactions on Information Systems 10(1) (January 1992) 91–102.

[34] R. Want, B.N. Schilit, N.I. Adams, R. Gold, K. Petersen, D. Goldberg, J.R. Ellis and M. Weiser, The ParcTab ubiquitous computing experiment, Technical report CSL-95-1, Xerox Palo Alto Research Center (March 1995).

[35] M. Weiser, The computer of the 21st century. Scientific American 265(3) (September 1991) 66–75.

Hans Gellersen holds MSc and PhD degrees in computer science from University of Karlsruhe, Germany. Following his MSc studies he became a research assistant at Karlsruhe's Telematics Institute in 1992. After completion of his PhD in 1996 he continued his research at the Institute as a Director of the Telecooperation Office (TecO), a division set up for collaboration with industry and European research consortia. Since March 2001, Hans Gellersen is a Chair in Interactive Systems in the Computing Department at Lancaster University. His research interest is in ubiquitous computing and human–computer systems that "take the real world into the loop": context-aware computing, situated user interfaces, and augmented everyday artifacts. He initiated the HUC/UBICOMP conference series in 1999, and serves as an editor of the Journal on Personal and Ubiquitous Computing.
E-mail: hwg@comp.lancs.ac.uk

Albrecht Schmidt received an MSc in computing from Manchester Metropolitan University (UK) in 1996, and in 1997 obtained a Master degree in computer science from the University of Ulm (Germany). As a research assistant he was working towards a PhD at TecO, University of Karlsruhe (Germany) from 1998 to 2001. Currently he is in the final year of his PhD at Lancaster University (UK). Albrecht Schmidt is especially interested in situated interaction and context-awareness in the research field of ubiquitous computing.

Michael Beigl is a Senior Research Assistant at the University of Karlsruhe, from where he obtained both his MSc (1995) and PhD (2000) degrees in computing. Dr. Beigl joined the University's Telecooperation Office (TecO) in 1995 to lead various ubiquitous computing projects, and became a Director of TecO in 2001. His research interests evolve around people at the centre of communication and information technology, with specific interest in novel information appliances and artefacts, in mobile and ubiquitous networks, human–computer interaction and in context awareness.