UBICOMP2002

ADJUNCT PROCEEDINGS

PETER LJUNGSTRAND AND LARS ERIK HOLMQUIST (EDS.)

FOURTH INTERNATIONAL CONFERENCE
ON UBIQUITOUS COMPUTING
SEPTEMBER 29 - OCTOBER 1, 2002
DRAKEN CINEMA AND CONFERENCE CENTER
GÖTEBORG, SWEDEN

PREFACE

This volume is the adjunct proceedings of the UbiComp 2002 international conference on ubiquitous computing, collecting extended abstracts for posters, doctoral consortium, and video submissions. At this point in time, we already know that this year's conference will be by far the biggest UbiComp yet, and the number of submissions in the different categories reflected this.

Posters are a vital part of an academic conference, since they provide an opening for late-breaking and controversial work. They also provide an opportunity for researchers to discuss and test new ideas in an informal setting. We accepted a total of 27 posters for UbiComp 2002, spanning every imaginable area of UbiComp research and development, from telematics to user studies, from protocols for music sharing to ubiquitous computing in education. The poster abstracts are the first and by far the most substantial part of this volume.

Doctoral students are important for the continued growth and development of the UbiComp research field. This year, we arranged the first UbiComp Doctoral Consortium, chaired by Anind Dey of Intel Research Berkley. The abstracts submitted by the doctoral consortium students comprise the second part of the adjunct proceedings.

The video program was an innovative new part of this year's conference, where we invited researchers and practitioners to submit videos that represented UbiComp research, new and old. We received a very high number of submissions and the video jury spent many hours viewing and commenting the videos. While the video program is mostly a "live" event, in the form of a screening of all selected videos at the UbiComp conference, several authors also chose to supply written abstracts. These comprise the final part of this volume.

Göteborg, Sweden, September 2002

Peter Ljungstrand, Posters and Video program chair

Lars Erik Holmquist, General Chair

This volume is published by the Viktoria Institute, Göteborg, Sweden, September 2002.
Cover design by Sus Lundgren, PLAY Studio, Interactive Institute.
Printed by TeknologTryck, Elektroteknologsektionen Chalmers, Göteborg.

TABLE OF CONTENTS

EXTENDED ABSTRACTS FROM THE UBICOMP 2002 POSTER PROGRAM

EXTENDED ABSTRACTS FROM THE UBICOMP 2002 DOCTORAL CONSORTIUM

EXTENDED ABSTRACTS FROM THE UBICOMP 2002 VIDEO PROGRAM

A Paper-Based Ubiquitous Computing Healthcare Environment

Magnus Bång, Erik Berglund and Anders Larsson
Department of Computer and Information Science
Linköping University
S-581 83 Linköping, SWE
magba@ida.liu.se

ABSTRACT
Clinicians use material objects such as paper documents, folders, and sticker notes to support cognitive and collaborative tasks. In this paper, we draw design implications from a medical workplace study and present a prototype for a paper-based ubiquitous computing case management system. LINDA2 and its digital pens provide a tangible paper interface that allows clinicians to use traditional paper forms as input devices. We discuss the pros and cons of paper-based input and report on two distributed user-interface approaches that are used to overcome problems of feedback in the paper interface.

Keywords
Ubiquitous computing, collaborative work, distributed cognition, paper interfaces, healthcare environments.

INTRODUCTION
In recent years, it has become apparent that material objects play important roles in supporting collaborative work [1]. For example, studies of Air Traffic Control have shown that the physical tools (e.g., flight strips) are crucial in coordinating and aligning the efforts of a team of collaborating individuals. Furthermore, professionals structure their physical environments to suit the cognitive demands of the work situation. For example, they offload cognition to the environment [2,3] by placing out things in space to track work-in-progress (i.e., they form external memory systems).

Unfortunately, very few computer systems designs have been presented that support and preserve the physicality of work environments to support cognitive and collaborative practices. Particularly interesting from a ubiquitous computing standpoint is, consequently, how we can design tools that maintain the physicality of work and also provide the advantages of the computer.

This paper presents LINDA2, a prototype of a case management tool with a physical paper interface and the underlying requirements that guided the design. The approach enables clinicians to combine paper-based practices with the advantages of digital media for data collection, storage and processing of information.

The specific aim of this paper is to illustrate how we can develop paper-based systems that bridge the physical and virtual worlds to support work practices.

WORKPLACE STUDY & REQUIREMENTS
We studied how clinicians worked with material objects in a middle-sized Swedish hospital [4]. Administrative work was paper-based. For example, *patient folders* were created for each patient that contained forms used to accumulate information such as which drugs that had been administered to the patient. To coordinate collaborative activities in the healthcare team, clinicians arranged spatially patient folders on a desk. A folder's position on the desk signified a patient's clinical status and her rank in the workflow.

The arrangement of objects in space supported several cognitive and collaborative functions in the workplace. For example, the arrangement of folders on the desk aided the team in maintaining a *shared awareness* of the current clinical situation (i.e., it displayed the current problem state). The arrangement also functioned as a *physical memory* that clearly offloaded individual memory tasks of clinicians. Additionally, nurses placed laboratory results and sticker notes visible in the folders for the team members; practices that *directed attention* to important information and tasks. These practices are difficult to support in a traditional computer system.

It became clear that the paper-based work needed to be maintained to support the different collaborative and cognitive functions that they supported. We approached the problem by augmenting paper with computing power, rather than replacing paper.

SYSTEM DESIGN
We developed a paper interface on top of an existing groupware system for case management in a distributed healthcare team. The system is used to track a set of patients in a workflow. For example, it is possible to create forms for each patient that states what drugs that have been given to them. Furthermore, forms can be placed in folders and sticker-notes can be glued onto documents.

In LINDA2, clinicians use digital pens, i.e., the Anoto system [5], to record medical data directly into the computer domain by writing on special paper forms. The digital pen scans paper printed with a unique pattern to capture the pen stokes.

Figure 1: LINDA2 provides a combined paper and walk-up display interface that let clinicians stay with their paper-based practices to record data to the medical computer system.

Furthermore, the digital pen provides a Bluetooth transceiver that sends the pen strokes to the computer domain.

In principle, the LINDA2 system enables users to have both virtual and physical representations of the same document and folder. Figure 1 illustrates the use of LINDA2.

Our initial experiments with the digital pen revealed a shortcoming of the paper-based approach. We found that the digital pen provided limited user feedback. It was difficult to know if the pen and the software applications captured and interpreted correctly what was written. The prototype pen provided only limited feedback through a system of light emitting diodes and pen-vibrations which were useful mainly to signal that the pen strokes was well-received by the computer. Furthermore, normal paper is naturally limited in providing feedback and memory aid, such as pop-up menus with predefined values from the underlying software applications.

To overcome the feedback problems, we experimented with two distributed user-interface approaches. These provided feedback by means of *auditory* and *visual* cues. A *distributed 'menu'* was developed in which feedback was channelled through external devices such as earphones and walk-up displays.

Both the auditory and the visual systems were written as thin clients. These clients were responsible for sending feedback to the appropriate devices, processing pen strokes and to access the underlying system. A character recognition engine and a set of string matching algorithms (i.e., Soundex and Bestmatch) interpreted the pen strokes.

In the auditory approach, a Bluetooth-enabled mobile headset provides the feedback. A special paper form was created with areas where users could tick with the digital pen to interact with the system.

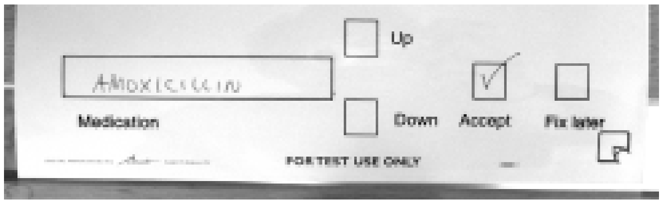

Figure 2: The paper form that constitutes part of the distributed user-interface.

For example, a clinician could write a medication on the form and the system will send its interpretation of the text in the earphone by means of an audio message. To confirm that it is interpreted correctly users mark the Accept box and the system verifies once again the selection in the earphone. To 'scroll' in a list of items, users tick the up and down areas of the form. We also added an option for cases when the system could not interpret the text appropriately (i.e., a Fix later box). Figure 2 shows the paper interface.

In the visual cue approach, feedback is provided by a walk-up display. Users fill in a value on the paper form and directly see how the system interpreted the text on the walk-up display. To confirm the system's suggestion, users simply tick the Accept box, as in the auditory approach, and the display shows an accept message.

We found that the auditory solution was cumbersome to use in our prototype due to bad response time. Additionally, auditory feedback could be difficult to use in a noisy clinical work environment. The walk-up display scenario is more preferable from an interaction standpoint. However, this solution is less mobile than its auditory counterpart.

CONCLUSION

This paper presented an approach to clinical computing that provided a paper interface to a case management system. It let clinicians stay with paper-based practices and allow them to physically structure their workplace to deal with cognitive and collaborative issues. Further work involves the full implementation of the prototype systems to evaluate the usefulness of the distributed user-interface approach.

REFERENCES

1. Luff P, Heath C, and Greatbatch D. 1992. Tasks-in-interaction: paper and screen based documentation in collaborative activity, in *Proceedings of CSCW '92,* New York: ACM Press, 163-170.
2. Kirsh D. 1995. The intelligent use of space*. Artificial Intelligence*, 72, 1-52.
3. Norman D. *Things that make us smart: defending human attributes in the age of the machine.* Addison-Wesley, Reading MA, 1993.
4. Bång M and Timpka T. Cognitive Tools in Medical Teamwork: The Spatial Arrangement of Patient Records. *Methods of Information in Medicine* (In press).
5. Anoto AB, 2002.

Analysis of Optimized Micro-Generator Architectures for Self-Powered Ubiquitous Computers

P D Mitcheson, T C Green, E M Yeatman and A S Holmes

Department of Electrical and Electronic Engineering
Imperial College of Science Technology and Medicine
Exhibition Road, London, SW7 2BT
pmitcheson@iee.org, t.green@ic.ac.uk, e.yeatman@ic.ac.uk, a.holmes@ic.ac.uk

ABSTRACT

It is advantageous for many types of ubiquitous computing artifacts to be capable of extracting energy from their environment, making them self-powered and self-sustaining. This paper presents a comparison of the two recently proposed micro-power generator architectures; the velocity damped resonant generator and the coulomb damped resonant generator. Also presented and analyzed is a new device, the parametric generator.

The analysis has shown that the parametric generator is likely to be useful when the input vibration amplitude is an order of magnitude greater than the dimensions of the micro-generator. It has also shown that for resonant generators, the efficiency of the technology used to realize the energy conversion is likely to be of greater importance in determining the output power than the architecture of generator used. Equations have been developed to enable the designer of a micro-generator to choose an optimal architecture.

Keywords

Vibration-to-electrical energy conversion, Self Powered Systems, Micro-power generators, Inertial-generators, Microelectromechanical devices.

INTRODUCTION

In order to successfully realize fully ubiquitous computing, it is necessary that each computing node is capable of running without user intervention. For computing nodes that are required to be physically small and mobile, a major obstacle to achieving this goal is the requirement of a continuous power source. Although there is continuing research into micro-batteries, this type of power source will only ever contain finite energy. Current thin-film micro-batteries can only achieve a capacity (per unit area of silicon) of 65µAh/cm^2 at around 4V [3]. A more attractive solution is for the node to scavenge energy from its environment, forming a self-powered system. This energy could be in the form of solar energy [2], thermal gradients or some form of movement [1], [4]. This paper presents a comparison of the main types of micro-generator architectures based on motion. These motion devices convert ambient mechanical vibration into electrical energy for use by ultra-low-power electronics, such as a computing artifact.

INERTIAL GENERATORS

In order to convert an input mechanical motion into electrical energy, some form of transducer is required. For the purposes of modeling the mechanics of the generator, the transducer can be realized as a damper with a suitable characteristic. The energy dissipated in the damper is the energy extracted from the mechanical system and converted into electrical energy. A typical inertial generator is shown in figure 1.

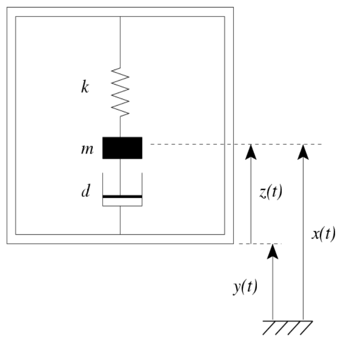

Figure 1: A Generic Resonant Inertial Generator

Types of Damping

In mechanical systems, damping is often modeled as one of the following types:

- Velocity damping
- Coulomb damping
- Hysteretic damping

A typical example of velocity damping is air resistance. An example of coulomb damping is dry friction of a body sliding along a surface, and hysteretic damping is normally used to model energy dissipated by the structural deformation of materials.

Any micro generator is likely to have little effect on the moving body that is providing the power. Thus, for the given constraints on generator mass and volume, it is important to choose the type of damping which allows most mechanical power to be coupled into the generator, dissipated in the damper and thus converted into electrical energy. It should be noted that current materials may or may not allow the chosen damping scheme to be sensibly mapped onto a particular technology.

Implementing the Damper

The following phenomena can be used to convert kinetic energy into electrical energy: Electrostatic field; electromagnetic field; piezoelectric effect.

An electrostatic generator can be realized with a moving plate capacitor. A parallel plate capacitor operated with a constant charge and a comb capacitor operated in constant voltage are both realizations of coulomb dampers. A velocity damper can be realized with a moving magnet whose flux is linked with a stationary coil. Piezo devices are best modeled as hysteretic dampers.

The Parametric Generator

In order to maximize the work done against an electric field, and thus maximize the electrical energy generated, the force-distance product should be maximized. The point of maximum force occurs when the acceleration of the input motion is at its maximum. Consequently, with the parametric generator, the moving plate of the capacitor snaps from minimum separation to maximum separation when the input acceleration is greatest. The relative plate motion (Z(t)) is shown in figure 2.

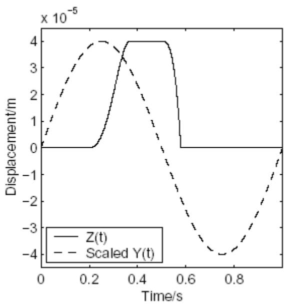

Figure 2: Optimal Parametric Motion

DISCUSSION OF RESULTS

Electro-mechanical generators have many independent variables associated with them. This means that assessing and fairly comparing the performance metrics of different generators is difficult. Analytic expressions for generator performance have been developed and verified with time domain simulation. As an example, the analytic expression for the optimal power generation by parametric generator is:

$$P = 2\beta \frac{\omega^3 m Y_0 Z_0}{\pi}$$

Where m is the proof mass, ω and Y_o are the input frequency and amplitude, β is the fraction of maximum acceleration at which generation starts and ω_c is the input frequency normalized to the resonant frequency. It has been found that when the force is optimized in order to maximize the energy dissipated in the damper, it is possible to plot the performance characteristics of all three generators on normalized axes. Figure 3 shows the optimal generator configuration for a set of normalized operating conditions. The types are, from lightest to darkest: Parametric Generator, Coulomb Damped Generator and Velocity Damped Generator. As can be seen, the parametric generator is optimal when the source of motion has an order of magnitude greater than the dimensions of the device. If a resonant device is used, there is little to differentiate the coulomb and velocity damped cases, and so a choice should be made depending on the efficiency of the implementation of each.

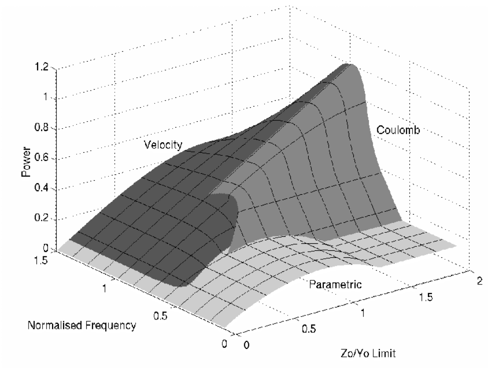

Figure 3: Optimal Generator Configuration

A parametric generator has been fabricated on a quartz wafer, to minimize parasitic capacitance, and is in the process of being tested. In order to realize a successful generator, the micro-power power-electronics and control and synchronization of the generator are important. There are trade-offs with the level of control used to ensure optimal power generation for a given input, and the power consumed by the control circuitry itself.

ACKNOWLEDGMENTS

This project is funded by the EPSRC and the DC Initiative (EU Framework IV) under the ORESTEIA project.

REFERENCES

1. Amirtharajah, R. and Chandrakasan, A.P., Self-Powered Signal Processing Using Vibration-Based Power Generation. *Solid State Circuits, IEEE Journal of*, Volume 33, Issue 5, May 1998, pages 687-695.
2. Lee, J.B., Chen, Z., Allen, M.G. and Rohatgi, A., A High Voltage Solar Cell Array As An Electrostatic MEMS Power Supply. *Micro Electro Mechanical Systems, 1994, MEMS '94, Proceedings, IEEE Workshop on*, pages 331-336.
3. West, W.C., Whitacre, J.F., Brandon, E.J. and Ratnakumar, B.V., Lithium Micro-Battery Development at the Jet Propulsion Laboratory. *IEEE Aerospace and Electronics Systems Magazine*, Volume 16, Issue 8 , Aug. 2001, pages 31-33.
4. Williams, C.B. and Yates, R.B., Analysis of a Micro-Electric Generator for Microsystems. *Solid-State Sensors and Actuators, 1995 and Eurosensors IX. Transducers '95, The 8th International Conference on*, Volume 1, pages 369-372.

"BM3G": Innovative Business Models for 3G Mobile Applications and Services

Avi Wasser
University of Haifa, Mount Carmel, Israel
Graduate School of Business
Center for the Study of Organizations
awasser@research.haifa.ac.il

Maya Lincoln
The Technion, Israel Institute of Technology
Faculty of Industrial Engineering
Center for Dynamic Enterprise Modeling
mayal@technion.ac.il

ABSTRACT

The study's main objective is to research, structure and plot Innovative **B**usiness **M**odels for **3G** Mobile applications and services. The generated business models will suggest an optimized business framework between the participants in the Mobile Internet value chain, including Mobile network operators as well as application providers, content providers, enterprise and corporate users, private users, handset manufacturers and infrastructure suppliers.

Keywords

Ubiquitous Computing, Business Models, Billing, Mobile Internet, UMTS.

INTRODUCTION

BM3G suggests an optimized business framework that facilitates the business relationships between participants in the Mobile Internet value chain, including Mobile network operators as well as application providers, content providers, enterprise and corporate users, private users, handset manufacturers and network and infrastructure suppliers. Each of BM3G's business models supports different business environments, taking into account a set of possible business parameters describing the different types of engagements and other related environmental variables and constraints.

Mobile data takeoff concerns

One of the main reasons for the failure of 2G mobile data technology was the lack of adequate and standardized business models (Zoller et al., 2001). Players in the 2G Mobile Internet arena developed and implemented different business models, without considering any collaboration with adjacent value chain entities. Let us consider for example the range of business models implemented by different Mobile network operators for their end users: each operator suggested a different business model to its customers, e.g., fixed price, metered connection, packet charging, edge charging, Paris metro charging, expected capacity charging, etc (Cushnie, 2000) resulting in a confusing, fuzzy and difficult-to-compare range of service packages. This concern gets even more complicated when Mobile data roaming issues are involved: in this case different business models are implemented by the user's home and roaming Mobile operators, with regard to the same given service. A similarly confusing business environment also subsisted between other participants in the 2G wireless arena. This absence of adequate and standardized business models was therefore reflected in significant financial loses for wireless carriers and other industry actors such as content and services providers. Users were confused by the lack of satisfactory business models, and reacted with suspicion to the Mobile Internet suites that were publicly offered. It could be possible to assume that an appropriate business model is a compulsory condition for a successful market penetration of any commercial ubiquitous application.

BM3G solutions

The new generic business models aimed to be developed in this research will involve Mobile operators, applications / service / content providers as well as handset manufacturers and network / infrastructure suppliers as a research group. The arrowed lines in Figure 1 (based on Cushnie, 2000) demonstrate the relationships between the involved entities. The lines represent the exchange of billing information, charging requests and transfer of funds.

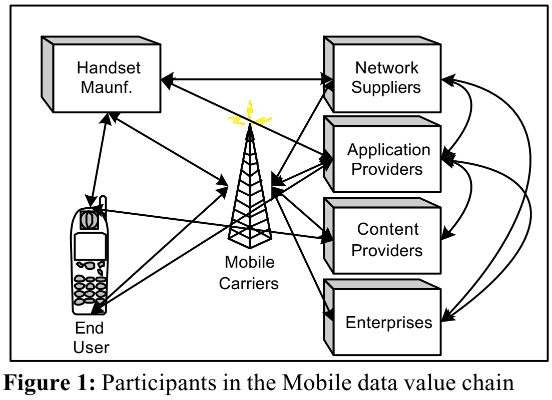

Figure 1: Participants in the Mobile data value chain

For suppliers, implementation of generic models can boost the adoption of Mobile data services, reduce customer churn, and generate new revenue streams. The customers on the other hand, will benefit from simple, consistent and reasonable pricing schemes that will help them overcome the mishaps of the previous generations' Mobile data offerings. The new generic business methodologies will make it feasible to plan and adopt lucrative business opportunities for the UMTS environment without changing the involved technologies (infrastructure, gateways, portal types, and existing services). The study offers the business basis towards a successful commercial launch of any Mobile service/application in UMTS.

METHOD

The research consists of three research stages that are designed to facilitate the achievement of the final result: a complete set of generic business models compatible for the 3G environment, ready for implementation and adoption by the commercial participants in the 3G environment.

First Stage

Process and analyze data regarding existing business models that are currently implemented in the 2.5G, and review future 3G business models. The results are based on a field research of Mobile operators, infrastructure suppliers, ISPs, and content/applications providers.

Second Stage

Generate and consolidate new generic business models between Mobile network operators and the other participants in the 3G Mobile value chain.

Third Stage

Test and examine whether the generic business models are applicable by simulating them using actual business scenarios and business cases.

RESULTS

The field research was aimed to research existing/planned business models of participants in the Mobile data industry. It included processing and analyzing data that was gathered from leading European participants in the Mobile Internet value chain, consisting of Mobile network operators, Mobile virtual network operators, application providers, content providers, enterprises, handset manufacturers and infrastructure suppliers. The study identified the following "meta-models" and a variety of their variations:

Metered charging: charging the subscriber for the connection to the service provider on a fixed basis and then charging for metered usage of the service as a function of use time or number of sessions.

Fixed price charging: assumes a lump sum payment for a predefined set of services.

Volume Charging and *Packet charging:* packet charging is specific to packet switching (Keshev, 1997). This model requires the implementation of packet counting systems in the network and the deployment of billing systems that can process packet data. A variation to this model is *content (*or *value) based charging* in which content type is a parameter that has an effect on the pricing.

Expected capacity charging: allows service providers to determine the bandwidth that a subscriber receives under congested conditions. The subscriber pays an agreed price for that level of service.

Edge Pricing: this model is based on the concept that networks cross-charge each other for the usage at the network 'edges' (Shenkar et al., 1997). The gathered billing information can be used for the deployment of metered, fixed, or expected capacity charging methods.

Metro (or QoS) Charging: assumes that subscribers will be assigned with a preferred "travel class" that indicates an associated cost for different network traffic (Odlyzko, 1999).

DISCUSSION AND FUTURE RESEARCH

The gathered information and knowledge have contributed to the global understanding of the existing and predicted Mobile Internet business models by focusing on their different charging methods and cost drivers, as well as on their advantages and disadvantages. This accumulated knowledge base laid the foundation for further developments and improvements. The comparative tests have shown a significant difference between the business models that were implemented in different business environments along the value chain. It was also shown that different business models were implemented in different parts of Europe. Among existing 2.5G models, the research identified fixed (flat rate) charging, metered connection charging, and Metro (QoS) charging with numerous variations and combinations. A significant difference was also viewed in the planned 3G business models, where business models mostly involved packet charging, edge charging, expected capacity charging, and value (content) based charging. It should be noted that the majority of the models that were developed by BM3G's participants were dealing only with the relationships between the end user (subscriber) and the Mobile network operator. BM3G's future research will include the generation of comprehensive generic business models, and an actual simulation based on real life business cases.

REFERENCES

1. Cushnie J., Hutchison D. (2000): "Charging and Billing Models for Future Mobile Internet Services", Distributed Multimedia Group, Lancaster University.
2. Keshav S. (1997): "An Engineering Approach to Computer Networking", Addison-Wesley.
3. Odlyzko A. (1999): "Paris Metro Pricing: The Minimalist Differentiated Services Solution", AT&T Laboratories Research, April.
4. Shenker S., Clark D., Estrin D., Herzog S. (1996): "Pricing in Computer Networks: Reshaping the Research Agenda", ACM Computer Communication Review, Vol. 26, pp 19-43.
5. Zoller E., Matthews J., Van Housen L., O'loughlin M., Davison J. (2001): "Wireless internet business models: global perspective, regional focus", an Ovum report, August.

Context, CC/PP, and P3P

Patrik Osbakk
Computing Laboratory
University of Kent
Canterbury, Kent
CT2 7NF, UK
+44 1227 823824
pjo2@ukc.ac.uk

Nick Ryan
Computing Laboratory
University of Kent
Canterbury, Kent
CT2 7NF, UK
+44 1227 827699
N.S.Ryan@ukc.ac.uk

ABSTRACT

In this extended abstract we describe how CC/PP, a framework for describing user preferences and device capabilities, can be extended to support general contextual information. Whilst there has been some concern for privacy issues in the intended uses of CC/PP, this more general use increases the need for protecting sensitive contextual information. We describe how an additional 'privacy profile' can be used together with P3P, to determine what information is revealed

Keywords

Context, Privacy, CC/PP, P3P

INTRODUCTION: CC/PP and Context

Composite Capabilities/Preference Profiles (CC/PP) [1] provides a way for user agents, typically browsers, to specify metadata about device capabilities and user preferences. Using this information, services and content might be adapted to the client's capabilities and needs with much greater flexibility than can be achieved using, for example, the information in existing HTTP headers.

The information contained in a CC/PP profile can be considered as contextual information in the sense that it describes the environment in which the device operates and the user desires to operate. As such, it represents just a small part of the information domain of context-aware systems [2]. By extending CC/PP beyond its original limited base it can be used to specify different types of context, such as user identity and contact details, location and current activity, as well as information about their environment, such as noise levels or weather conditions. Indeed, one of the very few published examples of a CC/PP profile that recognises a use beyond device description includes location and weather conditions in a "UserSituation" component [3]. Such information can be used to improve the personalisation process, both for content adaptation and as part of a user's interaction with a context-aware service. It can also provide a standardised platform-independent structure that can be used to communicate context information in a ubiquitous computing environment.

Extending CC/PP to this role is straightforward. New components and vocabularies may be defined for any purpose. The context information that could be included in a CC/PP profile will be limited only by available vocabularies and the encoding format. A standardised vocabulary and encoding format would give the greatest interoperability between systems, but private vocabularies may also be defined to suit application specific needs.

CC/PP and Privacy

For some users, revealing even the basic device information contained in a typical CC/PP profile may be a cause for concern. When the profile contains more general contextual information, privacy becomes a major issue. To use a secure communication channel is not enough as the profile content might be used without the users knowledge and for purposes that they would not welcome. It is therefore important to provide a way for users to protect their privacy.

Several possible approaches to protecting a user's privacy might be taken when working with CC/PP. The whole profile might be protected and disclosed only to trusted parties. An initial 'minimal profile' might first be sent as part of a process to determine whether a service can be trusted [7]. Alternatively, the selection of which parts of the profile are disclosed could be made dependent on knowledge about the service. This latter approach has been investigated in more detail because of its potential flexibility.

To be able to make automatic decisions regarding which parts of the profile to disclose or withhold, information about the user's privacy preferences needs to be available. The idea of a classification and clearance scheme has been used in this work to structure the privacy preferences. Each part of the profile is given a classification level that indicates the sensitivity of the information. For experimental purposes, a level from 0 (public) to 5 (private) has been used. Sites are then assigned a level of clearance depending on how trusted they are and what profile information they should be able to access. This scheme is somewhat limited and we anticipate using a more formal and flexible scheme, e.g. using a Role Based Access Control (RBAC) model [4], in the future. Below, we describe how the clearance level assignment can be automated using P3P.

Just as context information can be represented in a CC/PP profile by specifying a context vocabulary so also can privacy information with a privacy vocabulary. The design of the vocabulary would be dependent on what privacy scheme is used and how it is implemented. But, because the privacy information is local to the client device and is never disclosed, the vocabulary and its implementation may vary from system to system.

Profile Filtering

A single CC/PP profile is used to combine device capabilities, user preferences, context, and privacy information. This combined profile is for local use only and needs to be passed through a filter to produce the revealed CC/PP profile, or profile differences, to be sent with a request. The filtering process enforces the privacy requirements by creating a profile that contains only the information that the user desires to disclose to the recipient. If a site has not been assigned a clearance level only public level information is disclosed.

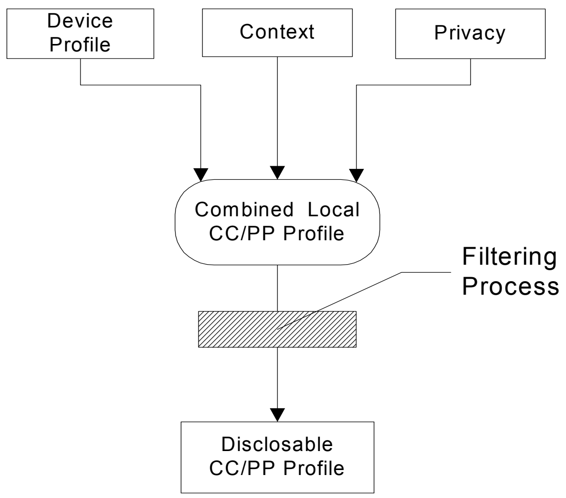

P3P

The Platform for Privacy Preferences (P3P) [5] is a privacy-enhancing technology. It allows a privacy policy to be described in a machine-readable standardised format. The P3P policy is intended to be used to make informed decisions about the interaction between a user and a remote service. But P3P can also be used to protect a CC/PP profile [6] [7].

The way P3P has been used to protect the users privacy in this investigation has been to use it to establish the level of clearance to assign to a site. This is important as not all recipients of a CC/PP profile are likely to be known in advance and so clearance levels cannot be pre-assigned. If a RBAC model was used P3P could be used to establish what role to assign instead. To establish the level of clearance, a site's P3P policy is retrieved and compared with defined rule sets.

There can be one rule set defined for each clearance level, except for the public level 0. In the rule sets a user can specify what a P3P policy must and must not declare for a clearance level to be attained. If a rule set is not defined for a clearance level it is assumed that it cannot be attained by evaluating a P3P policy. The comparison in this investigation has been done with a modified version of the JRC P3P Appel Evaluator [8]. The comparison starts with the rule set that grants the highest clearance level. If the evaluation results in a positive result the site is assigned the current clearance level temporarily. If the result is negative the comparison continues with the next lower clearance level rule set. The comparison continues until a clearance level has been assigned or no more rule sets are available, resulting in a clearance level of 0 being assigned.

Conclusion

We have described how CC/PP can be used to communicate context, how the users' privacy can be protected and how this protection can be automated using P3P. The use of P3P and CC/PP profile filtering has been described in a typical user agent/web server scenario. Other experiments have shown that this approach is also applicable where the profile can be queried by a remote system that can supply a P3P policy, either in a peer-to-peer environment or through an intermediate context service. It is hoped that technologies like CC/PP and P3P may help to make context-aware personalisation services more readily available.

REFERENCES

1. Composite Capability/Preference Profiles (CC/PP): Structure and Vocabularies [W3C Working Draft 15 March 2001], World Wide Web Consortium (W3C).
2. B. Schilit, N. Adams and R. Want, 'Context-Aware Computing Applications', IEEE Workshop on Mobile Computing Systems and Applications, 85-90, 1994.
3. CC/PP Implementors Guide: Harmonization with Existing Vocabularies and Content Transformation Heuristics [W3C Note 20 December 2001], World Wide Web Consortium (W3C).
4. Sandhu, R. S., Coyne E. J., et al. Role-Based Access Control Models. *IEEE Computer* 29,2 (1996), 38-47.
5. The Platform for Privacy Preferences 1.0 (P3P1.0) Specification [W3C Recommendation 16 April 2002], World Wide Web Consortium (W3C).
6. Nilsson, M., Lindskog, H., Fischer-Hübner, S. Privacy Enhancements in the Mobile Internet. *Proceedings of the IFIP WG 9.6/11.7 working conference on Security and Control of IT in Society*. (Bratislava, June 2001).
7. CC/PP Implementors Guide: Privacy and Protocols [W3C Working Draft 20 December 2001], World Wide Web Consortium (W3C).
8. Hogben, G. JRC P3P Appel Evaluator Java Module. Available at http://p3p.jrc.it/.

Controlling Your Environment Through Sketching

Edward Lank, Richard Augustine, Shi Chen, Heather Kane, Sandy Ngo and Joy Tutt
San Francisco State University
Computer Science Department
1600 Holloway Ave.
San Francisco, CA 94122 USA
lank@cs.sfsu.edu

ABSTRACT
This paper describes a storyboarding study on using unconstrained sketches by untrained users to control an interactive room. One interesting result which we observe is a similarity between a number of the sketches. This result argues for the possibility of designing a sketch based perceptive user interface which can interpret sketches drawn by a new, untrained user of the interactive room.

Keywords
Smart room, Sketch-based interfaces, Perceptual user interfaces.

INTRODUCTION
Sketch-based interfaces have been shown to be effective in early-stage design tasks [3]. We propose using sketch-based interfaces to control an interactive room. There are characteristics of sketching that are beneficial as a means of expressing concepts in a computer interface. Sketching is a natural means of communicating, commonly used in discovery tasks. As such, people are comfortable sketching. Unlike voice interfaces, sketching is private. In our interactive room, we try to preserve some aspect of personal or private space. Finally sketching is flexible. Because there is a many to one mapping of sketches to action, users can sketch the same command in a variety of ways.

Sketching to control interactive environments such as smart rooms also makes intuitive sense. These rooms often contain a number of pen interface devices such as electronic whiteboards and touch sensitive tables [1, 4]. As well, devices which users bring into the environment, for example PDAs and tablet computers, also have pen interfaces.

The ability to recognize arbitrary sketches (those sketches which are not created using a pre-defined notation) is an active area of research. Saund et al. and Sezgin et al. describe systems which attempt to perform low-level sketch recognition for the purpose of editing [2] and beautification [3]. Our work seeks to extend this to higher-level interpretation of arbitrary sketches. In order to accomplish higher-level interpretation, the sketch must be interpreted in the context of an application domain, which, in this case, is the control of an interactive room.

INTERFACING WITH AN INTERACTIVE ROOM
The SFSU interactive room includes two electronic whiteboards (touch sensitive plasma displays), a table with individual interactive displays, two web cameras, and two microphones. A sketch application, a web browser, a PowerPoint viewer, and a text editor are being developed. The room's infrastructure is managed by JINI running on a wireless network. Due to JINI's ability to dynamically incorporate devices, PDAs and laptop computers can be added to the room.

A number of interactive workspaces have been designed, including the iRoom of Fox et al [1], and the iLand system of Streitz et al. [4]. Authors of these systems note that an important concern in designing interactive workspaces is the user interface. Users should be able to effortlessly see how to interact with the interface, i.e. the user interface should be transparent.

To allow users to orient themselves in the user interface, we have designed an interface which mimics the layout of the physical room (see figure on next page). At the left, a simplified picture of the physical room is displayed. On the right hand side, we have incorporated three panes. The **Applications** pane contains software applications available in the room, the **File Server** pane stores application data, and the **Portable Devices** pane contains the devices which people have carried into the environment. The purpose of this interface is to permit users to freely sketch commands in an interactive environment, so the use of three panes on the right hand side was an arbitrary decision based on our view of the difference between software applications, application data, and physical devices.

EXPLORATORY STUDY
We conducted an initial exploratory study using storyboarding. The goal of this study was to determine whether a sketch-based interface was usable, and to determine commonality between sketches of different

users.

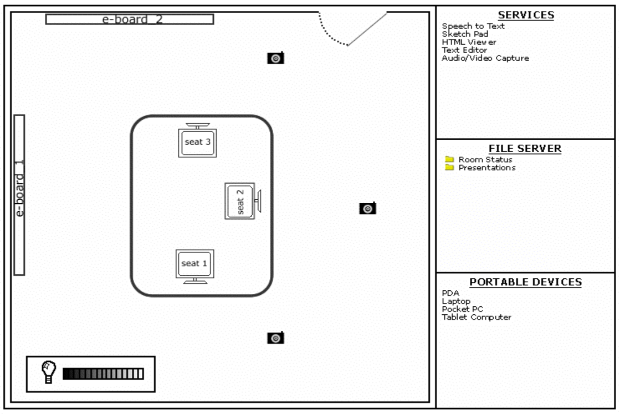

Our initial study involved 28 subjects. The subjects included employees in an industrial research lab, sales reps, a multimedia designer, mathematics teachers, and students. Each subject drew sketches of four different actions, drawn from a list of twelve actions, on a paper mock-up of the interface. Possible actions included connecting various devices together, starting applications on devices, and moving applications from one device to another. Participants knew, from the consent form, that the interface was of an interactive conference room, but were not in the interactive room during the study. Thirteen subjects were read a brief introduction to the user interface to orient them in the interface and were told that text was not permitted; the remaining fifteen had none.

Of our 28 subjects, 22 (79%) used only connectors (lines and arrows) and containers (circles and rectangles) to specify their task. Three subjects used a combination of connectors/containers and symbols or icons. The final three subjects used icons to describe their task. The use of connectors and containers is significant, allowing easy identification of the components involved in a command.

Education level had no discernable effect on users' drawings, but we did note a possible correlation between confidence with the experiment and the use of connector/container diagrams vs. icons. Without exception, those participants who used icons expressed a concern as to whether or not they were drawing the "correct" diagram. They were eager to avoid misrecognition. While a sample of three is too small to draw strong conclusions, there is evidence that confident users draw connector/container diagrams, and less confident participants tend to draw more complex artifacts.

We had another participant recognize the sketches and categorize them based on what actions they appeared to represent. The recognition rate of the participant examining the diagrams was 67%. This recognition rate includes obviously poor data where participants could not orient themselves in the interface and/or chose not to draw any diagram at all. Additionally, the verbal introduction seemed to be of no benefit. The recognition rate was actually higher – though not by a statistically significant margin – for users with no verbal introduction to the interface, 71% to 62%. Of the 15 subjects with no verbal orientation, 13 drew connector/container diagrams, and two drew connector/container diagrams with additional icons. Of the 13 subjects with a verbal orientation to the interface, nine drew connector container diagrams, one drew connector/container with icons, and two drew icons.

CONCLUSION

In this paper, we describe the design and some early implementation of a sketch-based interface for controlling an interactive room. The results of an exploratory study argue that users draw recognizable diagrams even when provided with no guidance as to the diagrams that are expected.

There are obvious limits to unconstrained sketching. We are just beginning to examine the limitations and benefits of sketching for control.

In the future, we plan to conduct a more ambitious wizard of oz study. We are in the process of implementing the sketch recognition engine and integrating the components of our interactive room.

ACKNOWLEDGMENTS

The authors would like to acknowledge the support of San Francisco State University for initial project funding. We thank Eric Cheng for designing icons in our user interface, and Raj Gopal and Katie Everitt for valuable discussions. Finally, we would like to thank the participants in our study.

REFERENCES

1. A. Fox, B. Johanson, P. Hanrahan, and T. Winograd, "Integrating Information Appliances into an Interactive Space", IEEE Computer Graphics and Applications 20:3 (May/June, 2000), 54-65.
2. E. Saund, J. Mahoney, D. Fleet, D. Larner, E. Lank, "Perceptual Organization as a Foundation for Intelligent Sketch Editing", in Proceedings of the AAAI 2002 Spring Symposium (Sketch Understanding), Stanford, March 25 - 27, 2002.
3. T. Sezgin, T. Stahovich, and R. Davis, "Sketch Based Interfaces: Early Processing for Sketch Understanding", in Proceedings of the Workshop on Perceptive User Interfaces, Orlando, Florida, November 15 – 16, 2001.
4. N.A. Streitz, J. Geibler, T. Holmer, S. Konomi, C. Mller-Tomfelde, W. Reischl, P. Rexroth, P. Seitz, and R. Steinmetz. "i-LAND: An Interactive Landscape for Creativity and Innovation". In Proceedings of CHI '99, 120-127.

Directed Information Harvesting

Geoff Lyon, John Grundbäck, Cyril Brignone, Salil Pradhan
Hewlett-Packard Laboratories
1501, Page Mill Road
Palo Alto, CA 94304
{gml, grundbac, cyrilb, salil}@hpl.hp.com

ABSTRACT
Directed information harvesting is a method for probing a highly mobile ad-hoc network for data. This method is designed to overcome certain shortcomings of normal data diffusion, in particular how data responses are routed back to the requesting node when the network topology is sparsely populated, dynamic and fragmented.

Keywords
Sensor networks, Ad Hoc, Context aware, Mobility

INTRODUCTION
At Hewlett Packard Laboratories, we are developing context aware appliances [1] with the ability to sense the users location, surrounding environment and are able to provide appropriate service links to local entities. These devices are worn by the users and create a dynamic ad-hoc network, containing a majority of people carrying nodes [2]. Data queries and responses propagate from person to person throughout the network via the casual proximity of pairs or groups of device carrying users. Such networks may also add to the quality of the data being exchanged as the users unconsciously create the contexts for data exchange and aggregation. We call this style of people centric networking a *gossip* network.

At an abstract level, a *gossip* network has similar characteristics to an ad-hoc sensor network. However many of the existing ad-hoc sensor network routing solutions are inappropriate as *gossip* networks, with much lower nodal densities, communicate via sporadic connectivity. Also, the high nodal mobility in *gossip* networks makes the creation of persistent multi-hop routing connections extremely difficult. Therefore we cannot be assured of consistent data paths for data query and response routing. To overcome the problems of network routing under such conditions we propose a technique called *directed information harvesting.*

AD HOC SENSOR NETWORKS
In a traditional ad-hoc sensor network, wireless battery powered nodes sense their physical environment and communicate with their neighbors. Nodes relay received packets to other nodes in the network, enabling multi-hop communications. Routing optimizations [3,4] are usually effective in limiting data flooding and system energy consumption. When one considers data propagation in sensor networks the term diffusion is often used. In a densely populated and stable network, data diffusion is linear and resembles ripples on a pond. This radial data flooding can be enhanced by directed diffusion [5]. This creates more direct routing for response messages returning back towards the querying node. Such techniques require that the network retains or re-establishes a direct return path for any network query. This requirement is usually satisfied in densely populated networks. However as the network becomes more sparsely populated and the individual nodes become more dynamic, the return path is more difficult to establish. Gossip networks are characterized as possessing a highly fragmented topology where the links between adjacent nodes may only occur at sporadic intervals or perhaps only once, for example when two strangers pass by each other on a street. Under such network conditions, data flooding is still possible, but finding a return path for network queries becomes extremely difficult.

DIRECTED INFORMATION HARVESTING
While a gossip type network may contain a large number of mobile nodes, their localized density and connectivity distributions will be quite low and subjected to continuous dynamic change. Under such a fragmented network topology, alternate techniques for message routing need to be considered. Often a node responding to a query does not have routing visibility back to the requestor, and may never regain this within it's local vicinity. Therefore an alternate return route must be established. In addition, the query messages may propagate to more than one matching node and so multiple response messages maybe generated, each requiring the establishment of a return path to the requestor. To achieve this, we propose a technique called *directed information harvesting.* This expands on the existing methods for directed data diffusion by introducing the notion of a more easily accessible repository to which response messages are routed, in the absence of a return path to the requesting node.

In *directed information harvesting* messages propagate through the network by normal diffusion with the addition of temporal storage on the forwarding nodes; we wish the query to reach as many other nodes as possible, including nodes that drift in and out of the network over time, not just those present when the message was created. Each node receiving the query will forward the message and, if a data match is found, will append a response to this message, thus continuing to propagate both the query any responses through the network. We wish to direct any responses back to the query originator, however this is difficult as the network topology is likely to have experienced considerable change since the query's generation. Instead, responses are routed towards an intermediate repository, which resides on the Internet. The location of this intermediary is specified as part of the query message, allowing any node with Internet connectivity to route response messages using IP. This creates a requirement for at least one, and more practically, additional nodes to have Internet access. In reality we envision gossip nodes to be embedded into user's cell phones or personal digital assistants, where the chance coupling or grouping of these devices shapes the network. Response messages are forwarded to the repository when these devices open a connection to the Internet or an Internet proxy, for example when synchronizing with an Internet connected workstation or during a cellular call. The repository harvests all the query responses and may aggregate the resulting data for the query-generating node. Once this node connects to the Internet, it can access the query results.

DISCUSSION AND FUTURE WORK

We have shown that directed information harvesting is a valuable addition to the existing techniques for data message routing in ad-hoc networks. In particular people-carrier networks of personal devices can benefit from this technique. However such devices are traditionally single user centric, individually operated and only store highly personalized data. The use of these devices to form a *gossip* network requires a more communal approach to device connectivity. Initially, users are likely to be skeptical of allowing their devices to store and route information for others, in an unconscious manner, and averse to too much of their devices battery life being consumed for this purpose. These concerns highlight the need for further work on device privacy, security and utilization to ensure the acceptance of *gossip* style networking solutions.

REFERENCES

1. Context Aware appliances at HP Laboratories. http://www.hpl.hp.com/research/cp/cmsl/research/as/hpswatch_research.htm
2. Glance, Natalie and Snowdon, Dave. "Pollen: Virtual networks that use people as carriers", in Proceedings of the International Symposium on Handheld and Ubiquitous Computing, 1999.
3. IP Flooding in Adhoc Mobile Networks. http://www.ietf.org/internet-drafts/draft-ietf-manet-bcast-00.txt
4. Mobile Adhoc Networks. http://www.ietf.org/html.charters/manet-charter.html
5. Directed diffusion: A scalable and robust communication paradigm for sensor networks Intanagonwiwat, Chalermek, Ramesh Govindan and Deborah Estrin *In Proceedings of the Sixth Annual International Conference on Mobile Computing and Networking (MobiCOM '00), August 2000, Boston, Massachusetts.*

Fingermouse: A Wearable Hand Tracking System

P. de la Hamette, P. Lukowicz, G. Tröster
Wearable Computing Laboratory, ETH Zürich
pdelaham@ife.ee.ethz.ch, lukowicz@ife.ee.ethz.ch, troester@ife.ee.ethz.ch

T. Svoboda
Computer Vision Laboratory, ETH Zürich
svoboda@vision.ee.ethz.ch

ABSTRACT
In this paper, we describe a system that implements limited hand tracking and gesture recognition capabilities in a small low power module. A first application is a mouse replacement, worn on the body, that captures the hand movement and allows the user to control a mouse pointer. The hardware is based on 2 image sensors, a distance sensor and a low power DSP. The software consists of different tracking algorithms supported by stereo vision based background clutter removal. Our system proves that computer vision can be useful even for mobile systems with strictly limited computational and power resources.

Keywords
Stereo vision, real-time hand finger/tracking, wearable

INTRODUCTION
In many situations, in particular when referring to real world objects, gesture and hand tracking are an attractive input mode for wearable computers. Unfortunately, gesture recognition and hand tracking in a mobile environment are tough challenges. In particular varying lightning conditions and a moving, cluttered background make a reliable recognition difficult.
Research on mobile gesture recognition has so far focused on algorithmic improvements paying little attention to the computation resources needed for their implementation. As a consequence most recognition systems (e.g. [3,4]) have relied on high end, power hungry systems or even stationary workstations connected to the wearable over a wireless network. This greatly reduces the usefulness of such systems as an input device for wearable computer.
Our approach addresses the computational resources problem in three ways. First, we restrict the scope of the recognition problem by requiring the hand to be in a certain fixed distance from the camera and to start tracking in a predefined position. Secondly we use a distance sensor together with stereoscopic vision to get rid of the cluttered background. Finally we have implemented a special purpose hardware module that allows us to execute the necessary algorithms with a fraction of the power consumption of a general-purpose machine.

As a result we have a system that is *small*, *low power* but nevertheless able to acquire images and compute them independently, so that it acts as an *autonomous* peripheral to a wearable computer.

THE HARDWARE:
AN AUTONOMOUS, LOW POWER VISION PLATFORM

The hardware consists of the sensors, the computing unit and the interfaces, as shown in Figure 1.

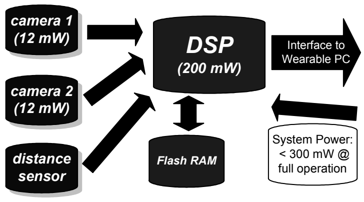

Figure 1: System architecture

The Sensors
The *Fingermouse* includes two low power CMOS cameras that capture black and white images with resolution of 256 by 256 pixels and 10-bit depth. The optics are small, light lenses with a very wide viewing angle of approximately 100 degrees, so that the hand can be close to the sensors.
Furthermore, a third sensor measures the distance of the hand from the *Fingermouse*. This can be used in the stereo algorithms or as a means to find out if a hand is present at all.

Figure 2: *Fingermouse*, top, cameras, dist. sensor

The Computing Unit
The main computing unit has been designed to execute computer vision algorithms efficiently. We chose the *Texas Instruments TMS320VC33* low power DSP, running at 75MHz. It features a floating point unit and enough internal RAM to store an image. Both are needed for most computer vision applications. An additional flash RAM chip stores the boot code and could also be used to store images or any other data. The complete power supply for the *Fingermouse* is included in the system as well, so that it can run from a simple battery. Interfaces are RS-232, JTAG and a 32-bit bus.

Figure 3: *Fingermouse*, bottom, DSP, Flash RAM

THE SOFTWARE: REAL-TIME TRACKING ON A LOW-POWER DSP

An important constraint in our system is that the background moves, since the device is worn on the body.
In this special case, we want to determine the position of a hand/finger in the captured image. This position is to be determined absolutely (e.g. the coordinates of the fingertip) *or* relatively, which would mean only relative movements would be tracked. A standard PC mouse also tracks the relative movement only.
Some algorithms are also able to determine the orientation (angle) and the scale (relative size) of the recognized hand. We will now briefly describe two of the algorithms that we have implemented.

Tracking the Contour

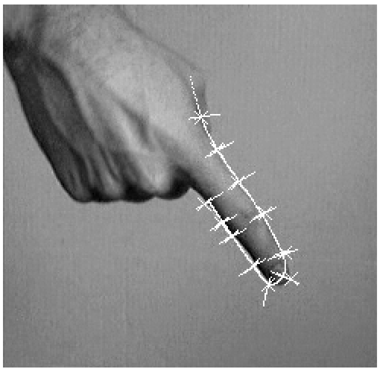

Figure 4: Hand, predicted contour and measured points (x) used for update

The algorithm used here [1], knows the shape and initial position of the hand. Starting from that, a prediction is made for the position, orientation and scale of the hand in the next captured picture. A new measurement in the predicted area then has to update the actual parameters. These allow again the prediction for the next step.
The measurement is performed only at a few control points along the shape of the hand. This has the advantage that the algorithm needs less computation resources. Kalman Filtering [1] is used to make the tracking robust against false measurements (that arise), for example, from a cluttered background.

Stereo Vision: Seeing the Foreground

Since the *Fingermouse* features two identical sensors in a stereo configuration, the so-called *disparity* may be used for the hand segmentation [2]. The optics of the sensors are designed so that a far background is projected identically on both image sensors. The hand is much closer to the cameras and projects differently (Figure 5a, 5b). Thus, the background vanishes in the difference image (Figure 6a). Simple thresholding and morphology operations are used to remove eventual noise (Figure 6b). Having a segmented image, a number of parameters may be computed like centroid, extremal boundary, the most significant orientation and others.

Figure 5a and 5b: Captured image by left and right camera

Figure 6a and 6b: Difference image, idem with threshold

DISCUSSION AND FUTURE WORK

We have shown a system that implements simple but useful computer vision algorithms in a compact low power module. Our primary application is a gesture-based mouse like input device for a wearable computer with limited power and computational resources. However our hardware platform is suitable for a variety of vision and signal processing algorithms found in wearable applications, which it can execute at a fraction of the cost (in terms of power consumption) of a conventional computer system.
In the next stages of the project the performance of the system as a wearable mouse will be evaluated quantitatively and the implementation of more complex gesture analysis will be investigated.
At a later point, an ASIC solution could reduce the system's power consumption by an order of magnitude.

ACKNOWLEDGMENTS

The *Fingermouse* was co-designed by Pascal Flammant, as a diploma thesis at the Swiss Federal Institute of Technology.

REFERENCES

1. A. Blake and M. Isard, *Active Contours*. Springer, 1998
2. R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2000
3. T. Starner, J. Weaver and A. Pentland. Real-Time American Sign Language Recognition using Desk and Wearable Computer Based Video, *IEEE PAMI Vol. 20, No. 12*, pp 1371-1375
4. T. A. Mysliwiec, FingerMouse: A Freehand Computer Pointing Interface. *VISLab Technical Report*, University of Chicago, 1994

FLAME: An Open Application Framework for Location-Aware Systems

George Coulouris, Hani Naguib, Kam Sanmugalingam

Laboratory for Communications Engineering
Department of Engineering, Cambridge University,
Cambridge, UK

{gfc22,han21,ks286}@cam.ac.uk

http://www-lce.eng.cam.ac.uk/qosdream/

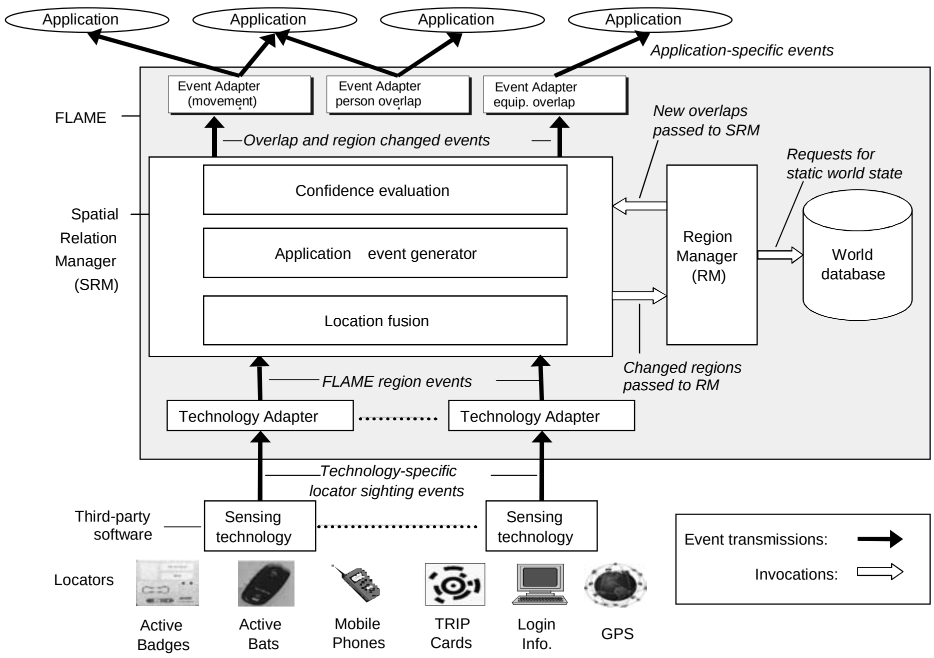

Figure 1- Location Management System

Abstract

QoSDream [1] is a research platform and framework for the construction and management of context-aware and multimedia applications. QoSDream simplifies the task of building applications by integrating the underlying device technologies and raising the level at which they are handled in applications. The platform is highly configurable allowing researchers to experiment with different location technologies and algorithms used to handle the information being generated by these technologies.

Keywords

Location-aware, middleware, sensor fusion, geometric modeling.

Introduction

The aim of the QoSDream framework is to provide a configurable, generic software framework that reduces the effort required to develop location-aware applications to a minimum.

FLAME is an implementation of QoSDream, implemented in Java. Its APIs are available as Java or Corba interfaces and it employs the Corba's Notification service for event handling. It interfaces to a database containing information about the physical world through the Java Data Objects (JDO) standard, allowing the use of a wide range of database systems.

Conceptual model

Applications are presented with events that provide information relating to a simple model of the physical world. This model contains the following abstractions:

Locators: Objects that can be sensed. (Active Badges/Bats, GPS devices); the locators are sensed, typically at frequent intervals, and each sensing generates a low-level, technology-dependent event.

Locatables: Objects whose location we wish to track. (people, equipment, …). Locatables are usually associated with one or more locators. This association allows the location of locatables to be inferred from the location of their locators. Applications subscribe to event channels that supply information about changes in the state of locatables and or locators.

Regions: Static regions (buildings, rooms, areas in rooms); dynamic regions associated with locatables (people's locations, their visual and auditory fields, …) The use of (polygonal) regions as the only geometric construct simplifies the architectural model and enables the inherent imprecision of each location technology to be precisely modelled (by associating an appropriately-sized region with each locatable). It also assists in the filtering and reduction of the event streams passed to applications.

Design concepts

- Raw location events filtered and abstracted to produce smaller streams of application-relevant update events.
- State of the physical and logical world models updated rapidly to reflect new location information as it arrives.
- Location and other state information stored persistently in a consistent, queryable form in order to allow a wide range of application services to be supported.
- Support for multiple sensor technologies;
- A simple spatial model for the representation of locatable entities;
- Processing of heterogeneous sensor data and its assimilation into a spatial model in real time;
- A simple event architecture for the presentation of location information to applications;
- Extensibility via pluggable algorithms and standard interfaces wherever possible.

Architecture

Figure 1 shows the main components of QoSDream in a typical application environment, the flow of events and the stages of transformation of the wide streams of technology-specific locator-generated events such as *Badge: ID 123 sighted in Room: ID 44*) into much narrower streams of application-relevant events such as *Dr. Smith is near Computer XYZ.*

Database: provides initial state (static regions, identities and other attributes of locators and locatables). Also holds a synopsis of the real-time location information.

Region manager (RM): Large numbers of static and dynamic regions are generated in a typical application. The region manager stores and retrieves regions according to their identities and their physical locations. The choice of data structure and the algorithms used for accessing regions have a critical impact on system performance. The current implementation uses a quadtree-organised data structure to enable the rapid detection of overlapping regions.

Spatial relation manager (SRM): generates application-related events in order to satisfy currently active subscriptions. The SRM employs several procedures:

Location fusion: integrates the regions associated with the (possibly several) locators associated with each locatable, using their intersection if there is one, else it retains their disjoint regions – representing an ambiguous location.

Application event generator: generates events giving the regions and relations between regions required by applications. For example, *Dr. Smith entered Room A22*.

Confidence evaluator: No sensor data is 100% reliable. The confidence evaluator applies heuristics to associate a confidence level with the regions passed to applications.

Event adapters: An example of such an adapter is the 'Person Movement Adapter' which listens on the SRM's Region Overlap Events and generates a Person Movement Event whenever it notices that a person has moved from one static region to another.

Preliminary evaluation and conclusions

As an indication of the impact the region manager can have on performance, table 1 shows the number of region sightings that can be handled by the SRM with each of the two region managers. These measurements were obtained by simulating sightings using actual sensor data that had been logged previously. The world model used for this represented an actual building with 52 rooms and 34 people. It is important to note that these results are only useful as a measure of the impact a region manager can have on the overall performance of FLAME. The number of sightings per second that can be achieved by FLAME depends on various other aspects including: the size of the world model, the number of static regions in the world model.the number and type of location sensing technologies.

Region Manager	Sightings/Second
QuadTree-based Region Manager	255
Set-based Region Manager	121

Table 1. Region Manager performance impact

Performance exceeding 1000 sightings per second has been achieved with a version of FLAME that used a less sophisticated database model than JDO; we anticipate that similar or better performance can be achieved by optimizations that have not yet been implemented following the adoption of JDO.

Recent efforts have been focused on achieving simplicity, modularity and clarity of interfaces. FLAME has emerged from several iterations of design and evaluation of earlier versions and is now available from us.

Reference:

1. Naguib, H., Coulouris, G. Location Information Management, in *Proceedings of UBICOMP 2001*, Atlanta , USA, Oct 2001, Springer

Game Engines for Use in Context Aware Research

Perry McDowell, Andrzej Kapolka, Michael Capps, Mike Zyda
MOVES Institute, Naval Postgraduate School
833 Dyer Road, Rm Sp-517
Monterey, CA 93943-5118 USA
+1 831 656 7591
{mcdowell, akapolk, capps, zyda} @nps.navy.mil

ABSTRACT
One of the biggest difficulties to overcome in creating and testing context aware applications is the interface with the real world. This includes both inputting data from the real world into a symbolic format and outputting the data to the user in a useful fashion. In this poster, we describe how we used a commercial game engine to overcome these difficulties and take the place of the real world.

Keywords
Context-aware, game engines, augmented cognition

INTRODUCTION
Ubiquitous computing is at a similar point to that which home computing was twenty-five years ago. Currently, while the hardware is cheap enough to permit market saturation, the applications have not yet reached the future that Weiser envisioned nearly a decade ago[1]. The main problem is that interfaces are lacking, and the only people using such systems are small groups of enthusiasts who are willing to endure a cumbersome interface. For ubiquitous computing to gain widespread use, the equivalent of the GUI must be created so novice users can use it easily.

We feel that creating context-aware applications are this equivalent of the desktop GUI for ubiquitous computing. By correctly determining the context of a situation, the computer can provide much better and more personalized assistance to the user. Building a system where the device can intelligently use all that information without the user having to input it will make it available to novice users.

Simple Context Aware Architecture
As part of our research, we have determined that, in order to be effective, context-aware ubiquitous applications must be able to perform three major actions, which are illustrated in Figure 1:

1. on the left hand side, the application converts the real world into a symbolic representation;
2. in the center, the application uses that representation to determine the context of a situation and a course of action for the user;
3. on the right side, the application informs the user of the recommended course of action and/or supply him with information he requires.

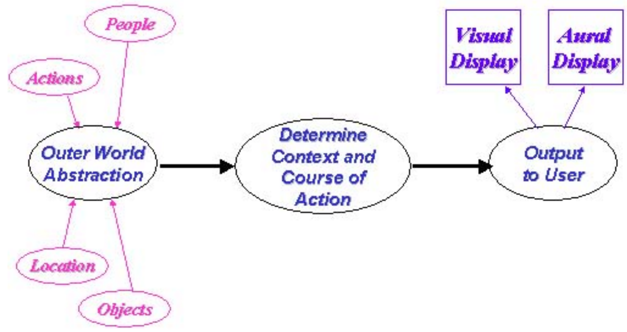

Figure 1 - Diagram of simple context-aware application

Notice that, of the three actions in Figure 1, only the center one is actually concerned with determining context. The left and right portions are concerned with issues secondary to the core issue of context. Therefore, all the time spent on those issues is time not spent on determining context.

Difficult Ancillary Problems
The left side of Figure 1 is a difficult problem because truly context-aware systems must use a significant number of real world attributes, some obvious such as location, time, user's schedule, or computing resources available, and many not so obvious, such as the physiological status of the user or environmental data. The difficulty lies in gleaning this information from the real world. The sensors involved in determining this information can be extremely difficult to create and use, often requiring frequent and/or extensive calibration.

The right side of Figure 1, outputting information so that the user can exploit it, is equally troublesome. While the data can be output using a wide range of techniques, most of these techniques involve emerging technologies. These include text to speech, augmented reality, and mobile/ wearable computers. Like all emerging technology, each of these can be difficult to use, and again much of the researcher's time is spent fixing secondary equipment.

Researchers who wish to limit their research strictly to context find themselves spending a great deal of their time on building, maintaining and using what should be only supplemental equipment to their actual investigation into context-aware computing. Additionally, much of this equipment is quite expensive, which forms another barrier to investigation.

Our Solution – A Commercial Game Engine

Dey approached the difficulty of dealing with several different types of these inputs and outputs by creating the Context Toolkit[2]. The Context Toolkit had widgets which allowed designers to abstract the method of determining information important for determining context. Our method to perform context-aware research without having to build the entire required infrastructure to support them was to use a high end game engine to simulate the real world for our context-aware application. The game engine performs many tasks important, yet secondary, to the actual building of the game itself. A game engine proved an exceptional way to simulate both the left and right sides of Figure 1, since its myriad of functions serves the same purposes in the game world.

We decided to use the UNREAL Game Engine, by Epic Games, because of its superb rendering, sophisticated artificial intelligence, advanced physics, and ease of use. This allowed us to test our context-aware applications while removing all the distractions of building systems which had to operate in the real world.

Many other researchers are using game engines to perform research in a wide variety of fields. Starner *et al* used them to create Augmented Reality applications[3], while Manninen also used the UNREAL engine to power a game which demonstrated contextual issues of a mobile application.[4] Our work is very similar to Bylund and Espinoza's, who used the Quake III engine to drive their context aware GeoNotes system via Dey's Context Toolkit.[5, 6]

Proof of Concept

In order to test our hypothesis as to the suitability of a game engine as a vehicle for building context-aware applications, we built a simple simulation of a Special Forces mission which was demonstrated at the Augmented Cognition Workshop in December 2001. While the mission simulated is significantly simpler than an actual mission would be, and the context-awareness of the simulation does not approach the final expectations, it serves as a proof of concept. Additionally, it effectively demonstrates several of the features of our application.

In this simple scenario, as the soldier carries out his mission, the context-sensing device helps him perform tasks by giving him directional guidance in the form of augmented reality arrows overlaid on the screen, pointing out important information that is either hidden or is difficult to notice, displaying the device's "degree of certainty" as to the identification of an object, and automatically recording important information. This allows the user to apply all brainpower to the task at hand.

Results

We were able to build an application with rudimentary context sensing simulating a soldier on a mission in six weeks using two programmers. The situations set up in the simulation adequately tested the context-sensing ability of the application. The information used to determine the context of the soldier's situation included his location and that of his enemies, ammunition status, mission status, and orders from higher authority.

We found that using a game engine has several of the advantages we had anticipated. Because the game engine keeps track of all information about every entity in the world, the application only has to query the engine to get any required information. Additionally, there were no delays due to malfunctioning equipment, which would have certainly been the case in a real world scenario.

However, there were some difficulties encountered while building the system. We discovered it was necessary to "dumb down" the information provided by the engine. Since the engine can provide perfect information about any object in the world, we found it necessary to limit the scope and accuracy of the information the engine delivered to more accurately reflect the information a soldier in the field would have. For example, the game engine provided the location of a hidden enemy, but that information was not provided to the context module.

Additionally, using a game engine is akin to learning a new programming language, and there was a learning curve involved for the programmers.

Finally, game engines are continually in flux, with the company often providing updates to their code. These updates are not always fully backwards compliant, and it is possible than an update will break existing code. When this happens, the programmers are required to debug the code, and many times these errors are not readily apparent.

REFERENCES

1. Weiser, M., Some Computer Science Issues in Ubiquitous Computing, *CACM, Vol 36*, Jul 1993.
2. Dey, A., *Providing Architectural Support for Building Context-Aware Applications*, PhD Dissertation, Georgia Institute of Technology, December 2000.
3. Starner, T. *et al*, Towards Augmented Reality Gaming, *Proceedings of IMAGINA 2000*, Monaco.
4. Manninen, T., Towards Communicative, Collaborative and Construtive Multi-Player Games, In *Proceedings of Computer Games and Digital Cultures Conference.* Mayra, F. (ed.), June 7-8, Tampere, Finland, 2002.
5. Espinoza, F., *et al*, GeoNotes: Social and Navigational Aspects of Location-Based Information Systems, *Proceedings of Ubicomp 2001*, Atlanta, Ga., 2001.
6. Bylund, M., & Espinoza, F., Testing and Demonstrating Context-Aware Services with Quake III Arena, *CACM, Vol. 45*, Jan 2002.

Intelligent Ubiquitous Computing to Support Alzheimer's Patients: Enabling the Cognitively Disabled

Donald J. Patterson, Oren Etzioni, Dieter Fox, and Henry Kautz
Dept. Computer Science & Engineering
University of Washington, Seattle, WA 98195
http://cs.washington.edu/homes/kautz/ac/
kautz@cs.washington.edu

ABSTRACT

Assisted Cognition systems provide active cognitive aids for people with reduced memory and problem-solving abilities due to Alzheimer's Disease or other disorders. Two concrete examples of the Assisted Cognition systems we are developing are an ACTIVITY COMPASS that helps reduce spatial disorientation both inside and outside the home, and an ADL PROMPTER that helps patients carry out multi-step everyday tasks.

Keywords

medical applications, context-aware computing, aging in place, artificial intelligence, proactive environments

INTRODUCTION

To date computer systems designed to help people suffering from cognitive disabilities due to aging, disease, or accidents have been rare and of limited scope. Recently, however, researchers in ubiquitous computing and artificial intelligence have come together to envision systems that can act as proactive partners in increasing the independence and security of people who have problems of memory, planning, and carrying out tasks of everyday life[7, 1, 5, 4]. A major motivation for developing intelligent caretaking systems is that the growing elderly population in many nations is leading to a crisis of demographics that will bankrupt current human and medical resources. The most common disease, Alzheimer's, currently affects four million Americans; by 2050, the number is expected to rise to 15 million, out of a world total of 80 million.

Both the physical and social environments greatly affect a cognitively impaired person's level of functioning. For example, an Alzheimer's patient may come to rely on the memory, remindings, guidance, and prompts of a caregiver, such as a spouse or other family member [6, 8]. Thus the social environment can provide *active interventions* that extend the patient's ability to handle the challenges of everyday life. But there is a limit to any person's capability to provide such help: physical and emotional 'burnout" of caregivers, often with serious health consequences, is a common phenomena [2].

The goal of Assisted Cognition is to develop computer systems that can provide such *active* assistance to an Alzheimer's patient. In brief, Assisted Cognition systems (i) sense aspects of an individual's location and environment, both outdoors and at home, relying on a wide range of sensors such as Global Positioning Systems (GPS), active badges, motion detectors, and other ubiquitous computing infrastructure; (ii) learn to interpret patterns of everyday behavior, and recognize signs of distress, disorientation, or confusion, using techniques from state estimation, plan recognition, and machine learning; and (iii) offer help to patients through various kinds of interventions, and alert human caregivers in case of danger.

The Assisted Cognition Project is an interdisciplinary effort between the University of Washington's Dept. of Computer Science Engineering, the Alzheimer's Disease Research Center, and the UW Schools of Nursing and Medicine. Industrial partners include Intel Research and Elite Care, whose Oatfield Estates are a living prototype of a ubiquitous-computing enabled residential home for the elderly in Portland, Oregon. The home senses and records nearly all activity going on in or around its campus, including the movements of all residents and staff, operation of all lights and appliances, movement of all doors, *etc.* This enormous and detailed real-world sensor stream provides one of the primary data sources for our project.

THE ADL PROMPTER

A common problem in the early to middle stages of Alzheimer's is a difficulty in carrying out complex tasks, while the ability to perform simple actions is relatively unimpaired [6]. For example, an Alzheimer's patient may be able to perform the individual steps in dressing him or herself, but be unable to recall the proper sequence of actions, *e.g.* that socks go on before shoes.

The ADL PROMPTER is a system that helps guide an impaired individual through the steps of such an "activitiy of daily living" (ADL). Input to the system comes from a sensor network embedded in the home environment. Data from the network (see Fig. 1), together with a rich probabilistic model of how activities are typically carried out, is integrated to create a model that predicts *what the patient is trying to do*.

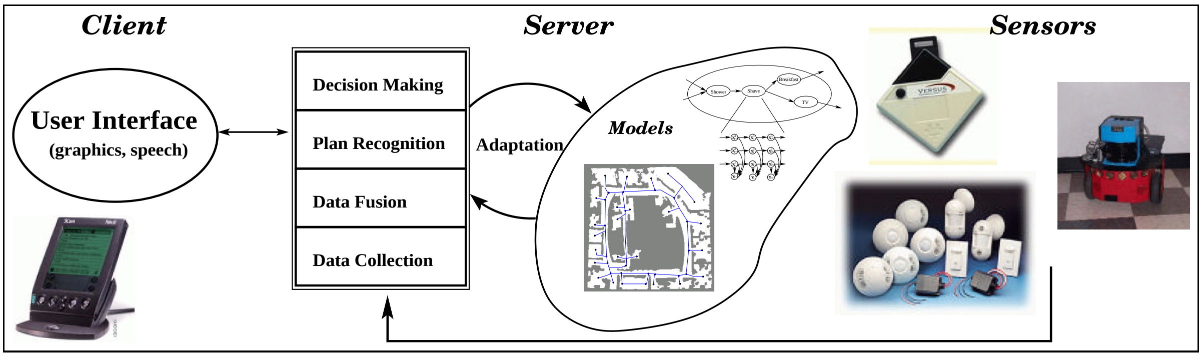

Figure 1: Architecture of Assisted Cognition systems. The server module consists of layers of increasing levels of abstraction. Probabilistic models of sensors and the environment are applied to process noisy sensor data. User activities are tracked using hierarchical Bayesian models. Probabilistic descriptions of these activities help to determine the user's plans and goals, which enable the system to choose which interactions to trigger.

The sensors may include ones for sound, motion, position of objects, movement of doors, operation of appliances, and so on. For example, the system might note that it is morning, and that the patient entered the bathroom and turned on the sink. Some time passes, and the patient remains motionless. The system predicts that the "morning toothbrushing and bathing" activity has begun but become stalled. Finally, the system decides to intervene by verbally prompting the patient to pick up the toothbrush. Note that a prompt is not given unless it is deemed necessary, and prompts are not preprogrammed for certain times of the time.

THE ACTIVITY COMPASS

The ACTIVITY COMPASS is a tool for helping a patient move independently (and safely) throughout his or her community. The compass is based on a client/server architecture where the client handles interaction with the user, and the server stores sensor readings, constructed models, and background information. Our current ACTIVITY COMPASS client is a GPS and wireless-enabled Palm PDA.

The Assisted Cognition differs from an ordinary GPS guidance device in several crucial ways: (1) Instead of requiring the user to manually enter a destination, the system attempts to predict the user's destination based on learned patterns of behavior. (2) The Assisted Cognition can *proactively* engage the user: for example, if it infers that the user is wandering and may be lost, it may suggest that the patient head home by a audible prompt and a simple graphic display (*e.g.*, an arrow pointing in an appropriate direction). (3) The Assisted Cognition can link current data about the user's movements with external environmental information that is relevent to the user's goals.

An example of such external information are bus routes and real-time bus locations (which are available in the Seattle area). The following scenario illustrates how such information could be used: *1. Don is walking toward a bus stop; 2. The system notes that at this time of day Don frequently catches the bus home; 3. Real-time bus information shows that bus leaves in 5 minute, and next one is in an hour; 4. The system predicts Don will miss bus at his current rate of walking; 5. The system decides to intervene by alerting Don to walk faster.*

TECHNICAL COMPONENTS

Both the ADL PROMPTER and ACTIVITY COMPASS are based on a layered architecture (see Fig. 1) linking sensor data to simple behaviors, behaviors to plans and goals, and predictions of the success or failure of those plans to potential interventions. Each layer takes in noisy and uncertain information, abstracts and fuses the data to reduce (but not always eliminate) uncertainty, and passes this higher-level view of the world to the next layer. Feedback from the effects of invention feed back down through the layers, in order to improve the accuracy and effectiveness of the underlying models. More detailed technical discussion appears in [3].

ACKNOWLEDGEMENTS

The Assisted Cognition project is supported by grants from the Intel Corporation and the Intelligent Information Systems Institute (ISII).

REFERENCES

1. G. Baltus, D. Fox, F. Gemperle, J. Goetz, T. Hirsh, D. Magaritis, M. Montemerlo, J. Pineau, N. Roy, J. Schulte, and S. Thrun. Towards personal service robots for the elderly. In *Proc. of the Workshop on Interactive Robotics and Entertainment (WIRE-2000)*, 2000.
2. N. R. Hooyman and H. A. Kiyak. *Social Gerontology: a Multidisciplinary Perspective (6th Edition).* Allyn and Bacon, 2001.
3. H. Kautz, L. Arnstein, G. Borriello, O. Etzioni, and D. Fox. An overview of the assisted cognition project. In *Working Notes of the AAAI-2002 Workshop on Automation as Caregiver: The Role of Intelligent Technology in Elder Care*, 2002.
4. C. E. McCarthy and M. E. Pollack. A plan-based personalized cognitive orthotic. In *Proceedings of the 6th International Conference on AI Planning and Scheduling*, 2002.
5. A. Pentland. Machine understanding of human action. In *Proceedings 7th International Forum on Frontier of Telecom Technology*, Tokyo, Japan, 1995.
6. B. Reisberg, S. H. Ferris, J. J. Leon, and T. Crook. The global deterioration scale for the assessment of primary degenerative dementia. *American Journal of Psychiatry*, 1982.
7. J. M. Sanders. Sensing the subtleties of everyday life. *Research Horizons*, Winter 2000.
8. D. Shenk. *The Forgetting: Alzherimer's: Portrait of an Epidemic.* Doubleday, New York, NY, 2001.

iStuff: A Scalable Architecture for Lightweight, Wireless Devices for Ubicomp User Interfaces

Meredith Ringel, Joshua Tyler, Maureen Stone, Rafael Ballagas, Jan Borchers
Stanford University Department of Computer Science
{merrie, jtyler}@cs.stanford.edu, stone@stonesc.com,{ballagas, borchers}@stanford.edu

ABSTRACT

iStuff ("interactive stuff") is an architecture for a toolkit of lightweight, wireless, platform-independent, physical user interface components. We present examples of several iStuff devices we have developed, and discuss the software infrastructure that supports them and allows for easy configurability and extensibility.

Keywords

User interface, tangible interface, physical interface, tuple space, interactive workspaces, wireless devices.

INTRODUCTION

iStuff is a toolbox of wireless, platform-independent, physical user interface components designed to leverage the tuple-space-based software infrastructure of Stanford's iRoom, a technology-augmented room used as a testbed for ubiquitous computing and user interface research[4]. Application users can easily and dynamically configure iStuff physical interface components to flexibly set up sensors and actuators in a ubicomp environment without having to solder components, run wires, or write device drivers. As a prototyping toolkit, iStuff aims to facilitate research at the crossroads of ubiquitous computing and HCI.

In the past, Ishii's Tangible Bits project [3] introduced the notion of bridging the world between "bits and atoms" in user interfaces. More recently, Greenberg's Phidgets [2] provide physical widgets, designed for rapid development of physical interfaces that expand Desktop GUIs. In contrast, our approach assumes an entire interactive room as its environment. Consequently, our devices must be dynamically retargetable to different applications and platforms. Additionally, devices are lightweight because they can leverage the existing interactive room infrastructure.

iSTUFF ARCHITECTURE

Several criteria guided our design of the iStuff architecture. We wanted our devices to have completely autonomous packaging and battery-powered operation, the ability to communicate wirelessly with existing applications in the iRoom, and simple, affordable circuitry. We found that we were able to use several different hardware technologies to create iStuff devices that fulfilled these criteria; so far we have working iStuff made from custom-built RFID components, from X10 components, and from standard FM radios and transmitters. From these platforms we have built several different types of devices (see Figure 1) including buttons, sliders, LED's, buzzers, and speakers. Schematics for some of these devices can be found at http://www.stanford.edu/~borchers/istuff/. Additionally, we have integrated commercial products as iStuff devices, such as MIDI controllers for 2D input, and portable microphones with voice recognition software for speech input.

Ultimately, the specific communications medium (RFID, X10, Bluetooth, 802.11b, etc.) employed is irrelevant, because of the iRoom's underlying software infrastructure. This infrastructure is based on the Event Heap—an event-based communication mechanism that connects all of the platforms, applications, and devices in the room. To add a new physical hardware iStuff platform (such as Bluetooth), one merely needs to write a glue layer that translates hardware input into an event or, conversely, translates events to hardware output. Once the device driver is written for a platform, each individual device can be uniquely mapped and configured without modifying the driver.

Input devices (e.g., buttons, sliders) transmit device specific data (e.g., device ID, position) via their wireless protocol to a receiver which is connected to a proxy computer via a parallel, serial, or USB port. This data is then packaged into an event and placed on the Event Heap using a hardware glue layer. Applications or devices can be controlled by iStuff by adding a listener notification method for specific events, which can be done with a few lines of Java code. Output devices (e.g., LEDs, buzzers, speakers) work in a complementary manner, with software listeners converting events with device-specific data to hardware output.

Device events can be translated to different application events via a patch panel application that has user configurable, dynamic mappings. This allows the user to pick up an iStuff device, go to a configuration webpage and change the mapping of the device to control a running application on any machine in the room in less than 30 seconds.

We are currently expanding the iStuff framework by including new types of discrete and continuous input and output devices, and by experimenting with new technologies for wireless transmission such as Bluetooth.

DISCUSSION

By embodying most of the "smarts" of iStuff in a computer proxy which sends and receives IDs and handles the posting and retrieving of event tuples, we were able to make the physical devices themselves very lightweight (Figure 2). They merely need the capability to send or receive their device-specific data.

Figure 1.Various iStuff. The iDog sends a button press event when turned over. The iPen is an augmented SMARTBoard pen, where the embedded button operates as a "right click" to the Windows OS. The X10 buttons are standard X10 hardware. All other devices work through a simple RF receiver that plugs into the USB port of the Proxy PC.

Utilizing the tuple-based event model of the Event Heap has allowed us to create configurable devices which are platform-independent. This architecture provides great flexibility and rapid prototyping of input devices [1]. It has allowed us to turn several commercially available items like X10 controllers into iStuff, in addition to our own custom-built hardware. The "patch panel" software increases the flexibility of our toolkit by allowing for dynamically configuring mappings between events and devices, thus permitting exploration of the ramifications of having input devices that are not tied to a specific machine or display.

iStuff has proven its worth in our lab by allowing us to quickly create experimental interfaces. For instance, we have integrated iButtons into meeting capture software as a way to provide customized annotations, and we have made a videogame where iSliders can substitute for mouse input (refer to our poster for more details on usage scenarios).

As technology continues to move beyond the desktop and toward ubiquitous computing environments, we predict that device architectures like iStuff will become increasingly commonplace. By leveraging the infrastructure in its environment, iStuff enables rapid prototyping and configuration of new and creative user interfaces for ubicomp settings.

ACKNOWLEDGMENTS

We would like to thank Michael Champlin, Joyce Ho, Robert Brydon, and Adam Rothschild for their contributions, and the NSF for graduate student fellowship support.

REFERENCES

1. Borchers, J., Ringel, M., Tyler, J., and Fox, A. Stanford Interactive Workspaces: A Framework for Physical and Graphical User Interface Prototyping. *IEEE Wireless Communications*. Special Issue on Smart Homes, 2002 (in press).
2. Greenberg, S. and Fitchett, C. Phidgets: Easy Development of Physical Interfaces Through Physical Widgets. *Proceedings of UIST 2001*, 209-218.
3. Ishii, H. and Ullmer, B. Tangible Bits: Towards Seamless Interfaces Between People, Bits and Atoms. *Proceedings of CHI 1997*, 234-241.
4. Johanson, B., Fox A., and Winograd, T. The Interactive Workspaces Project: Experiences with Ubiquitous Computing Rooms. *IEEE Pervasive Computing Magazine,* 1(2), April-June 2002.

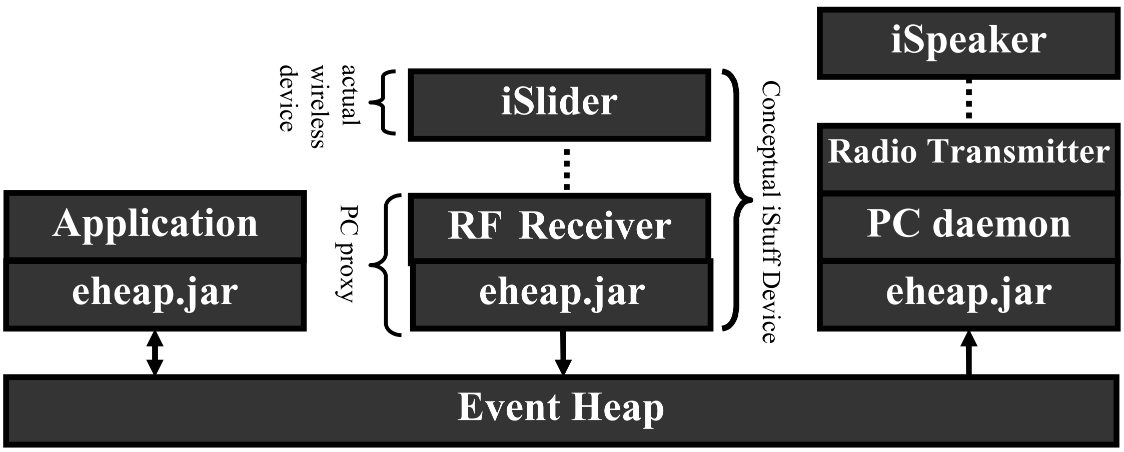

Figure 2. The iStuff architecture – much of the functionality is handled by the computer proxy, allowing the actual physical device to be quite simple and lightweight. The Event Heap mediates communication between disparate devices and applications – in the above diagram, the Application can receive events from the iSlider and send events to the iSpeaker via the Event Heap; the slider and speaker could easily be replaced by other iStuff input and output devices.

Miniaturised Modular Wireless Sensor Networks.

John Barton, Kieran Delaney,
Cian O' Mathuna,
National Microelectronic Research Centre
Lee Maltings, Prospect Row, Cork, Ireland
+353 021 4904088
{john.barton, kieran.delaney, cian.omathuna}@nmrc.ie

Joseph A. Paradiso
MIT Media Laboratory
20 Ames St. E15-351
Cambridge, MA 02139
joep@media.mit.edu

ABSTRACT
This paper focuses on the development of miniaturised modular wireless sensor networks that can be used to realise distributed autonomous sensors for future ad-hoc networks in the ambient systems and intelligent environments arena. Such modular, mobile networks are key enabling technologies in the field of ubiquitous computing.
The modules are fabricated in a 3-D stackable form with a novel modular PCB design which can be mounted on artefacts or on parts of the body, can measure acceleration, rotation, shock, elevation etc. and have a low-power RF channel-shared link to a base station (for sports, exercise, entertainment, health). The modular nature of the design allows for extra panels to be developed and added easily.

Keywords
Modularity, wireless sensor networks, 3-D packaging.

INTRODUCTION
Major research efforts are currently targeting the "disappearance" of the computer into the fabric of our environment. In the future, the spaces we live in will be populated by many thousands of objects (often described as "artefacts") with the ability to sense and actuate in their environment, to perform localised computation, and to communicate, even collaborate with each other. Artefacts are playing a large role in research towards intelligent systems and ubiquitous computing. There are two prime drivers: the smaller these objects are the more effective they will be in providing opportunities for integrating the physical and digital worlds, and the greater the number of objects within these systems/networks the more valuable the networks are. The main properties required to maximise the capabilities of such networks are that it should have, granularity (i.e. high resolution), reconfigurability modularisation and mobility. The system level implementation will be realised through concurrent hardware and software engineering; innovation in software should be matched by invention in hardware. It is important that novel hardware technology platforms are used for object and system development, incorporating 3-D stacking, mu lti-chip and micro-sensor integration, thin and flexible substrates, active polymeric materials, smart materials, and ultimately micro-nano-systems. To do this, new form factors for hardware need to be investigated, optimizing performance. In this light, the key initial considerations are interconnection and modularity of the hardware.

PROBLEM STATEMENT
Recent developments in wireless and micro-sensor [1,2] technologies have provided foundation platforms for considering the development of effective modular systems (see figure 1). They offer the prospect of flexibility in use, and network scalability. Currently, most sensor networks are strongly integrated into the assembly process of their target systems (E.g. the automobile, production line equipment, aircraft, etc). Thus, they carry a high infrastructural overhead. Emerging autonomous formats include wireless units designed to collect data and transmit to central (or distributed) hosts. Interesting examples include passive/active tags, inertial measurement units (IMU), the 1cm^2 wireless integrated micro-sensors at UCLA [3], and the "Smart Dust" project [4,5] at the University of Berkeley.

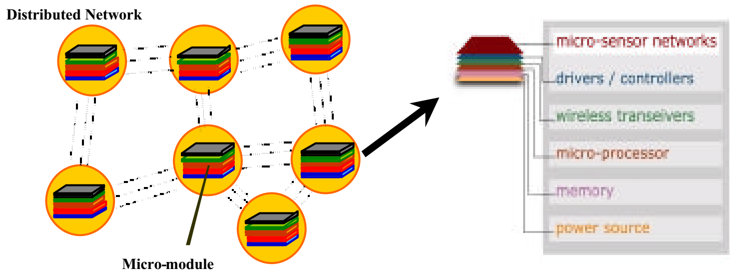

Figure 1: System Format with Modular Nodes.

The availability of modular systems will provide a valuable research tool for ambient systems and ubiquitous computing. Thus, a modular approach initially for a wearable sensor network, including functionality as an inertial measurement unit has been adopted by the NMRC. The module comprises an ensemble of 16 sensors, electronic interfaces, and a wireless transmitter manufactured using a combination of current surface mount techniques and multichip packaging (MCP). The sensors included accelerometers, gyroscope, compass, pressure sensors, bend sensors, electric field sensor and sonar pinger The module includes; an integrated PIC micro-controller with A/D converter, separate 256K EEPROM memory for local data storage and a 433MHz RF Transceiver with 20kbit/s data rate within a multi-chip module (MCM). Current prototypes consist of miniature sensor packages that can be worn on limbs and torso or mounted within artefacts.

BACKGROUND
The miniaturised wireless sensor networks presented here are the evolution of a project collaboration between NMRC and MIT Media Lab with the aim of miniaturisation and ruggedisation of the MIT Media Lab

“expressive footwear”[6]. Prof. Joe Paradiso and his team at the MIT Media Lab have prototyped a sensor component network of 16 sensors that are integrated with a wireless transmitter and mounted onto a shoe (expressive footwear) [7,8]. The sensors were located either in the insole or on a PCB attached to the shoe. A dancer, equipped with a pair of this footwear, has every movement monitored by an off-stage computer. Although currently used to explore applications in interactive dance, this system also has applications in areas like sports training and medicine, interactive exercise, and podiatry diagnosis and therapy.

MODULE EVOLUTION

The re-design of the module took place in 2 phases: the development and ruggedisation of the wearable sensor platform and the re-design of the circuits for the miniaturised modularised form factor.

The ruggedisation was necessary to improve the module reliability. This was done in two steps. Firstly, the wearable sensor portion of the circuit was re-designed and fabricated on Copper-clad flex. Sensors incorporated in the flex circuit included bend sensors, dynamic pressure sensors and force sensing resistors. Secondly, these circuits were laminated between protective plastic sheets and reliability tested through a series of non-standard reliability tests with a high yield.

During the design phase of the module, a building block technique for the autonomous wireless sensor network was developed. The formal re-design of the module PCB was completed with the aim of miniaturising and modularising the circuit to allow it to be unobtrusively worn anywhere on the body. The final design was realised as a 90mm x 30 mm two layer PCB which could be mounted as is or separated into three 30mm x 30mm panels which can be positioned on any portion of the body. Figure 2(a) shows the module before segmentation while 2(b) shows a separated module.

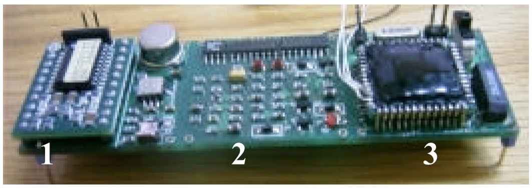

(a)

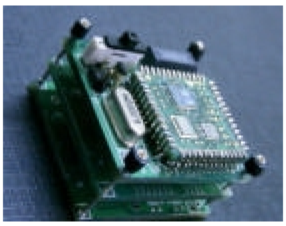

(b)

Figure 2: Autonomous Sensor Network Node with Modular Design Format (a) before, and (b) after segmentation.

The 3 panels visible are 1) the inertial measurement panel 2) the force sensors interface panel and 3) the wireless transceiver panel. The modular format of the PCB allows for extra panels to be designed and manufactured as and when required. Though stacked in 3-D in Figure 2(b), miniature flex cable connectors on each panel allow the modules to be connected in a variety of different ways and in an unobtrusive manner. Preliminary versions of this module have been utilised in projects including a wearable network, interactive glove and a localisation system.

CONCLUSION AND FUTURE WORK

This paper has presented the work done in evolving a module for a miniaturised wireless sensor network. The implementation of this form factor is useful in numerous applications, including for sports, exercise, entertainment, and health; in addition, the imaginative use of flex circuitry may provide for further form factors to be evaluated (for example, connected panels could be wrapped around the wrist). The current size is too large to expand the application potential of the form beyond niche level; the realistic number of stackable panels is currently four. To expand the viability of the format, it is appropriate to look at the potential for further miniaturisation of this module with an initial target volume of 1 cm^3.

The continuation of the project will focus on challenges, which are key to bringing the intelligent environments concept to reality: 1) The further development of high density, (3-dimensional) packaging technology platforms to enable increased miniaturisation of micro-sensor modules incorporating application-specific sensors, data acquisition, signal processing and wireless communication 2) Applications in a mobile/wearable domain with emphasis on context, and design 3) The realisation of a distributed system of micro-sensor modules. There will be an increased focus upon design and development of appropriate form factors for mobile/wearable applications and realisation of the target dimensions.

ACKNOWLEDGEMENTS

This project was funded by the Irish Government Higher Education Authority (HEA).

REFERENCES

[1] Estrin et al, “Next Century Challenges: Scalable Coordination in Sensor Networks”, Proc. Of ACM MOBICOM, 1999.
[2] S. Meguerdichian et al, “Coverage Problems in Wireless Ad-Hoc Sensor Networks”, Proc. Of IEEE INFOCOM, 2001.
[3] G. Pottie, W. Kaiser, "Wireless integrated network sensors," Communications of the ACM, vol. 43, pp. 51--58, May 2000
[4] J. M. Kahn, R. H. Katz and K. S. J. Pister “Mobile Networking for Smart Dust”, ACM/IEEE Intl. Conf. on Mobile Computing and Networking (MobiCom 99), Seattle, WA, August 17-19, 1999.
[5] Warneke, B, Last, M, Leibowitz, B, Pister , K S J, “Smart Dust: Communicating with a Cubic-Millimeter Computer”, IEEE Computer Society, vol. 34 no. 1, p. 43-51, January 2001.
[6] J. Paradiso, K. Hsiao, A. Benbasat, Z. Teegarden, “Design and Implementation of Expressive Footwear,” IBM Systems Journal, Volume 39, Nos. 3 & 4, October 2000, pp. 511-529.
[7] http://www.media.mit.edu/resenv/
[8] J. Paradiso, K. Hsiao and E. Hu, "Interactive Music for Instrumented Dancing Shoes", Proc. of the 1999 International Computer Music Conference, October 1999, pp. 453-456.

Mobile Bristol: A New Sense Of Place

Morris Williams
University of the West of England
Faculty of Computing and Mathematical Sciences
Morris.Williams@uwe.ac.uk

Constance Fleuriot*
University of Bristol
Mobile and Wearables Research Group
fleuriot@cs.bris.ac.uk

Keri Facer
Futurelab, Bristol UK
keri.facer@nestafuturelab.org

Josephine Reid, Richard Hull
HPlabs Bristol, UK
Josephine_Reid@hp.com,
Richard_Hull@hp.com

Owain Jones
University of Bristol
Geographical Sciences
owain.jones@which.net

ABSTRACT

This poster describes a two day workshop conducted with children aged 11-12 in which they designed and created soundscapes. This was a first stage in examining the potential of wearable computing devices to extend children's mobility and play opportunities within the urban environment. Such opportunities are a crucial part of child development but are often curtailed by real and perceived threats. The creation of the soundscapes was followed by a brainstorm session in which the children suggest directions for application and technology development that would augment their experience of outdoor environments.

Keywords

Wearables, soundscapes, augmented reality, children, mobile technologies

INTRODUCTION

This poster describes a two day workshop held during June 2002 at Hewlett Packard Labs, Bristol. The workshop was the first step for a project called 'A New Sense of Place', which brings together researchers from a variety of disciplines to explore the impact of pervasive computing and mobile technologies on children's play, learning, mobility and interaction with the outdoor environment. We are interested in asking if there are ways in which these technologies can facilitate children's fuller, freer, safer yet adventurous use of the urban environment. The objectives of the workshop were to explore whether young people would engage with the technology and wearable devices that we are developing for a wider project, Mobile Bristol[1] .

RATIONALE

In the UK (as elsewhere) concerns for the conditions of modern childhood are particularly articulated through visions of the child in problematic relations with urban environments [3]. The urban is the site of issues of considerable concern in relation to childhood, such as traffic, rises in levels of asthma, fears for the safety of children, particularly 'stranger danger', and fear of other crime against, or even by children. These perceived dangers have often "driven children from the street into their bedrooms" [6]. Children however need to play and to use the outdoor environment as a key part of their development [4], exploring, developing routes and short-cuts, and negotiating 'ranges' with parents and carers. De Certeau [1] argues that "the childhood experience that determines spatial practice later develops its effects, proliferates, floods private and public spaces, undoes their readable surfaces and creates within the planned city a "metaphorical" or mobile city...a great city". Those opportunities for walking have been eroded; A New Sense of Place is then very much concerned with urban childhood and the implications wearable technologies may have in urban settings. As Ward [7] said, "I want a city where children live in the same world as I do".

RESEARCH QUESTIONS

Our research questions relate to the longer term nature of our work and are:

- How can we develop technologies to empower children and young people to explore their environments, and to play safely without adult control and intervention?
- Can wearable devices overcome parental concerns about child safety and mobility?
- How does a real space become an important place through familiarity of use or the feelings associated with it?
- What landmarks will people choose to augment, with what will they augment them and what routes will they create between those landmarks?
- For whom will they design the personal landscapes they are creating and how will they be used?

THE TECHNOLOGY

The current technology is comprised of a bag containing a small portable computer (HP Jornada) interfaced to an ultrasonic location system, headphones, batteries and networking hardware (to access a wireless 802.11).

We have used the term 'soundscape' to describe the work that is created; the overlaying of a physical landscape with digital audio [2]. The client device detects its location within the physical space using an ultrasonic positioning system [5]. A directory on the server contains the information which defines each 'aura'. This consists of an unique

[1] (see http://www.hpl.hp.com/hosted/mbristol)

identifier, a location, a radius of applicability and the url of the audio file associated with that aura. Each soundscape is built from a number of such auras. When the user enters an aura the client device uses the url, to play the associated audio file to the user.

THE WORKSHOP

Workshop Objectives

The inter-disciplinary research group came together after a Mobile Bristol conference in January 2002. The objectives of the workshop were:

- to introduce the idea of wearables and demonstrate the technology to the children
- to teach the children how to use the technology and enable them to build their own soundscapes
- to allow the children (and others) to experience their own and their colleagues soundscapes
- to brainstorm about the future use of the technology

Workshop Content

Ten 11-12 year old school students (6 girls and 4 boys) from John Cabot City Technology College were selected by the school to participate in the two-day workshop. The children were introduced to the technology, the idea of the soundscapes and the process of production.

Working in pairs with an adult enabler, the children were encouraged to plan their soundscapes working on whiteboards, paper, or straight onto the system. They soon engaged with the process and clearly enjoyed working with the headsets and handheld computers. The adult enablers took a hands-off approach as the children worked on their soundscapes. The five pairs easily completed and fine-tuned their designs, which they had the opportunity 'to sleep on' between day one and two.

On the second day we held a brainstorm to elicit the children’s thoughts on the technology, if it was of any value to them and whether it might be developed to have an impact on their lives. From a general perspective, the response of the children (and the many teachers who visited) was overwhelmingly positive.

Thy reflected on the ways in which the technology could be used to improve their local areas - providing information reports to police, giving instructions on how to behave in those sites, playing music and sound-based games that would enhance the experiences in those areas.

The children commented on the potential of the technology to keep their parents informed of their locations - if only to ensure that they would come to collect them. They also suggested applications of the technology for people with a range of disabilities. The children thought it could provide a useful additional 'layer' of experience on entering unfamiliar environments. Importantly, they seemed to see the potential for people to use this technology to produce soundscapes for each other - very different from a broadcast model.

PRELIMINARY FINDINGS

- Children are valuable, adaptive and creative users in the participative design of ubiquitous computing experiences and the devices that might enable them.
- Children are able to articulate potential wider uses after experimental use in a confined environment. These preliminary responses indicate a positive role for wearables in extending children's mobility.
- Gender differences in the potential usage of such devices emerged - girls appear to be driven by the social potential of such devices e.g proposing personalised soundscapes for friends, boys appear to focus on pragmatic uses e.g. communicating class information in school.

WHERE NEXT?

The research group aims to work with young people as creative collaborators. The workshop was an exciting first step in the development of the project, bringing young people and these emerging technologies together, and we aim to build on this. Future workshops are planned that will take place in the outdoor space, developing the project outside the protected environment of HP labs. Full scale trials are envisaged for the technology in public spaces around Bristol and a dedicated workshop for young people in one of Bristol's large parks is being discussed. A major research project focused specifically on young people, education and mobile technologies is currently under consideration by the ESRC. Future projects also involve developing the soundscape creation tools, and identifying the capabilities needed for the software and hardware.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the staff and pupils of John Cabot Technology College.

REFERENCES

1. Certeau, M. de The Practice of Everyday Life. University of California Press, Berkeley 1984
2. Hull, R., Reid, J.,and Kidd, A. Delivering Compelling Experiences through Wearable Computing. Submitted to: Pervasive Computing
3. Jones, O. Melting Geography: Purity, Disorder, Childhood, Space in Holloway, S. and Valentine, G. (eds.) Children's Geography : Living, Playing, Learning, 29-47. Routledge, London, 2000.
4. Moore, R.C. Childhood's Domain: Play and Place in Child Development. Croom Helm, Beckenham, 1986.
5. Randell C. , H. Muller. Low Cost Indoor Positioning System. in UbiComp 2001 Atlanta: ACM.
6. Summers, D. A modern child's home appears to be its refuge. Financial Times. London 25 April 1995.
7. Ward, C. The Child in the City. Architechural Press, London, 1978

funded by Equator see www.equator.ac.uk

Optimized Maintenance and Inspection of Distributed Processes Based on Local Tracking and Messaging

M. Polaschegg, M. Schmid, Ch. Steger, R. Weiss
Institute for Technical Informatics
Graz University of Technology
{polaschegg, schmid, steger, weiss}@iti.tugraz.at

ABSTRACT

Inspection and maintenance in large industrial plants with a large number of distributed processes are essential tasks to guarantee trouble-free production performance. However, finding an optimal schedule for these maintenance jobs is very tedious. Minimizing the time for job execution can save costs and can increase the efficiency of a plant enormously. We propose an environment for the optimization of job performance with object tracking and location based messaging. Allocation of jobs to human service agents and local scheduling are automatically generated using the combinatorial optimization algorithm Simulated Annealing. A messaging service automatically guides each service agent to the jobs to be performed and to the necessary tools.

Keywords

Bluetooth, RF tags, Object tracking, Location based messaging, Optimization, PDA, Simulated Annealing

INTRODUCTION

The efficiency of tuned production processes in industrial plants mainly depends on the performance of their inspection and maintenance jobs during life cycle. Service jobs always have to be executed during runtime, providently or because of actual failures of components. Additionally, revisions, were plants are offline, allow to perform extended maintenance and inspection, i.e., in power plants.

However, finding an optimal partitioning and scheduling for these service jobs is very tedious. For a human service agent the problems are (i) to find an optimal schedule for a given number of jobs, (ii) to find the right location of the process to be inspected, and (iii) to find the location of the appropriate tools, which are often shared between agents. Jobs may also have dependencies, i.e., a process may only be inspected after fulfillment of other jobs, or new jobs have to be dynamically scheduled during inspection because of actual failures.

Therefore, we propose an optimization system, which is based on our LOMOT, a **Lo**cation based **M**essaging and **O**bject **T**racking environment. Our system tries to optimize the schedule of service jobs with the aim to minimize the time for the fulfillment of all service activities.

SYSTEM OVERVIEW

Figure 1 denotes a basic example for distributed processes with human service agents and tools in a given surrounding.

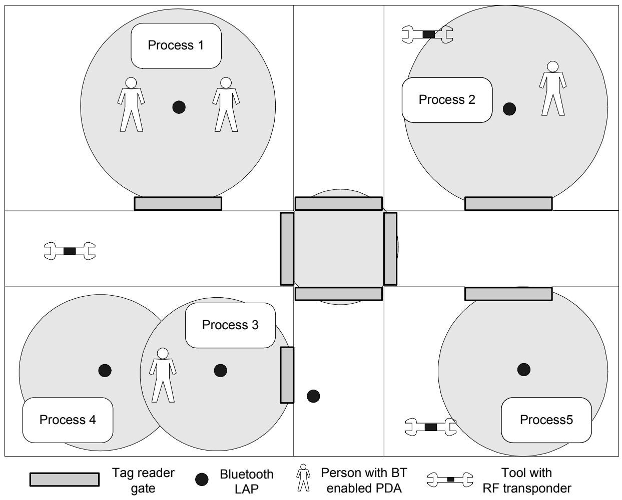

Fig. 1: Example of distributed processes with the installation of LOMOT devices for object tracking and messaging.

LOMOT is intended to track objects and to support location based messaging of objects within a predefined surrounding. Basically, LOMOT takes two types of objects into account. First, there are passive objects which can only be tracked. The tracking mechanism is based on RF tags. Appropriate tag reader gates with a width of about 2 meters detect the passage of such tags, which are attached to the objects (tools). The granularity of this object localization depends on the maximum distance between the reader gates. Second, active objects are persons (human service agents), who wear Bluetooth (BT) enabled personal digital assistants (PDAs). Tracking is performed by detecting their entry and leaving of reception areas of one or more BT access point (AP). This mechanism provides an accuracy of about 10 meters. During BT access, location based information is automatically transmitted from a message server to the PDA.

Structure of LOMOT

Figure 2 denotes the client-server structure of our LOMOT system.

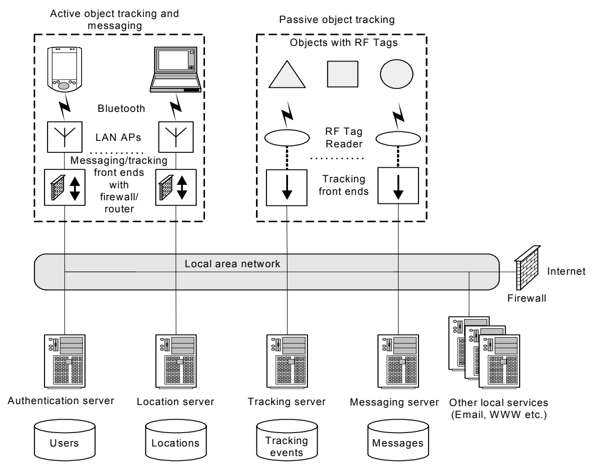

Fig. 2: Overview of our location based messaging and object tracking system LOMOT.

A main goal of LOMOT is to provide a secure environment. BT APs and tag readers are connected by front ends to the LAN. The front ends connecting the BT APs to the LAN contain router, firewall and a local DHCP server. This mechanism avoids direct connections of PDAs to the LAN and possible attacks to the infrastructure of LOMOT. All active objects are identified by an authentication server with their BT address, a user ID and a PIN code. Transmissions of data between BT APs and PDAs are encrypted, too. The tracking server annotates each pass of an object through a tag reader gate and each access of a PDA to a BT AP. The location server contains positions of all tag readers, BT APs and the topology and boundaries of the surrounding. This server also provides the service to locate all persons/objects by combining the location information with tracking events originating from the object to be searched. The messaging server stores location and user dependent information to be transmitted to the PDAs.

Optimization

Allocation and scheduling of jobs is optimized using Simulates Annealing (SA) [4]. This central optimizer collects information about each job, its predicted execution time, its priority and the dependencies of jobs among themselves. First, an initial allocation of jobs to human service agents is performed and an initial schedule is calculated. Necessary tools are allocated to agents and jobs, too. Finally, the optimization result is transmitted to the PDAs and all agents can start with their job execution. Nevertheless, some fuzziness in job execution has always to be taken into account. Duration of jobs can differ from predicted values due to the changing performance of human agents. Human agents also can ignore instructions from the optimizer by mistake. Our system solves these problems by transmitting the current state of job execution to the central optimizer, which immediately calculates a new schedule for each agent. To avoid conflicts when accessing tools a semaphore mechanism has also been introduced.

RELATED WORK

Similar problems for job scheduling can be found in literature. W. Shen [1] proposes an agent-based approach. Resources like machines, operators, robots etc. can be treated as intelligent agents connected by a LAN. Such a system supports distributed local scheduling instead of a centralized scheduling mechanism.

Bluetooth and RF tags for location aware computing are topics of interest of the following two projects. R. Kraemer [2] demonstrates a successful system for location based messaging with BT. The implemented fairguide for the CeBIT 2001 consisted of 130 LAN APs in a hall of 25000m^2 with 200 BT enabled PDAs. Moschgath et al. [3] introduce a Sm@artLibrary to demonstrate localization and location management technologies, mainly based on RF tags. Tag readers are integrated in shelves, door frames and freely available web pads to localize and identify books containing RF tags.

CONCLUSION

We have presented an environment for object tracking and messaging based on BT and RF tag technology. The goal of this system is to optimize inspection and maintenance in large industrial plants with a large number of distributed processes. A distributed optimization scheme using SA provides near-optimal schedules with the ability to handle dynamic changes in the schedule.

ACKNOWLEDGMENTS

We would like to acknowledge the support from Texas Instruments.

REFERENCES

1. Shen, W. Distributed Manufacturing Scheduling Using Intelligent Agents, *IEEE Intelligent Systems*, pp 88-94, 2002.
2. Kraemer, R. Bluetooth Based Wireless Internet Applications for Indoor Hotspots: Experience of a Successful Experiment during CeBIT 2001. *Proceedings of IEEE LCN* 2001 (Tampa FL, Nov. 2001), 518-524.
3. Moschgath, M.-L. Hahner, J. Reinema, R. Sm@rtLibrary – An Infrastructure for Ubiquitous Technologies and Applications. *Conference on Distributed Computing Systems* (Mesa AZ, April 2001), 208-213.
4. Kirkpatrick S. Gelatt Jr. C. D. Vecchi M.P. Optimization by simulated annealing. *Science,* 220(4598):671-680, 1

Parsimony & Transparency in Ubiquitous Interface Design

Christopher R. Wren
Cambridge Research Laboratory
Mitsubishi Electric Research Laboratories
201 Broadway - 8th floor
Cambridge MA USA 02139
wren@merl.com

Carson J. Reynolds
Media Laboratory
Massachusetts Institute of Technology
20 Ames Street
Cambridge MA USA 02139
carsonr@media.mit.edu

ABSTRACT
Minimalism in ubiquitous interface design allows computational augmentations to seamlessly coexist with existing artifacts and the constellations of task behaviors surrounding them. Specifically, parsimony and transparency contribute to improved learnability and user acceptance of novel interfaces.

Keywords
parsimony, transparency, minimalism, design, tangible, ubiquitous computing, learnability

DESIGN OVERVIEW
We present a design philosophy for ubiquitous computing centered on parsimony and transparency. By transparently integrating aspects of the digital world into real artifacts, we strive to provide ubiquitous interfaces to computation that do not obscure or destroy the highly refined interaction modalities of the host artifact in the physical world. We believe that carefully crafted coexistence of the physical and the digital, based on minimalism, leads to more learnable interfaces.

We also present a system that demonstrates this design philosophy: an augmented go board. The game of go is a demanding test case because it is surrounded by a set of behaviors and aesthetic concerns that have been refined over the course of thousands of years. Our system provides a flexible, modeless augmentation of go by adhering to minimalism.

Minimalism doesn't necessarily imply limited functionality. Transparent design means minimizing cognitive demands on the user by limiting the changes to the pre-existing constellation of behaviors surrounding the artifact being augmented. For example, the augmented go system transparently adds game recording and an automatic move clock to traditional face-to-face play with no change to the traditional experience.

Furthermore, parsimony means minimizing the introduction of interface elements and inappropriate metaphors that could lead to clutter. The traditional activities of solitary review, study, and problem solving are enhanced by the addition of minimal visual augmentations that are appropriate to the game context and therefore preserve the game aesthetics. The traditional experience is actually improved because the user is free not only from distractions in the interface, but also from the usual distractions of notes, reference books, and newspapers.

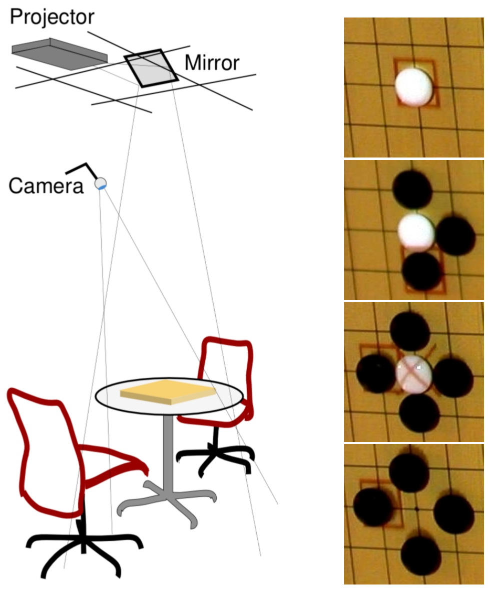

Figure 1: Left: The system. Right: A sequence of moves shows the visual annotation. The red box indicates the last move. The red × indicates a white stone that should be removed.

An overview of the physical configuration of the system, as implemented, is shown on the left side of Figure 1. The right side of Figure 1 shows examples of the projected patterns used to augment the board.

IMPLEMENTATION
The system we implemented is governed by the design philosophies of parsimony and transparency articulated above. The system itself consists of a light-table comprised of a video camera and projector situated above a go board on normal table. The system projects visual annotations that form a superset of the traditional board functionality: a game clock, a remote or artificial opponent's moves, and a minimal interface for exploring game variations. The vision system explicitly supports our design philosophy by accommodating the traditional style of game play. Unlike some other tangible user interface (TUI) light tables, like Underkoffer and Ishii's URP[3], our system relies solely on unmodified, traditional artifacts. Moreover the system is adaptive to various light-

ing and geometrical arrangements of the go board. What this means is that users interact with the light table-enhanced go board as they traditionally would, but are also provided useful augmentations. Consequently, players already familiar with go can learn to use the interface quickly because we've minimized the behavioral adjustments needed to use our interface.

EXPERIMENTAL VALIDATION

We verified our assertions about learnability by conducting a user study that we will briefly detail here. The experiment had two conditions: a well designed graphical user interface (GUI), and our augmented physical board. The experiment compared these conditions within subject (over 10 subjects). The independent variable was the choice of condition. The dependent variable was time on task. The task for both conditions was to play out a game and explore a variation of that game. The ordering of which condition subjects first encountered was randomized. We infer that our system is more learnable because the time on task for users with no experience with the game was less with our system, than with the GUI. Given the null hypothesis that the mean times are the same for the two conditions, but with unequal variances, the two-tailed probability that this data represents the null hypothesis is $\boldsymbol{p} = 3.3817 \times 10^{-6}$. Our system was also subjectively reported by subjects as preferable.

Table 1 shows the data from the experiment. The prefix on the subject number indicates which condition was experienced first: 0 for TUI, 1 for GUI. The † denotes that the participant failed to complete the task before the time limit (n.b. this only happened in the GUI condition), and μ and σ give the mean and 95% confidence interval for the data. Questions 1 and 2 asked the subject the difficulty of using our system and the GUI, respectively (with 1 indicating easy). Question 4 asked the subjects their preferred system (with 1 indicating a preference for our system). Answers were given on a five point scale. The remaining questions established the subjects' familiarity with go and with computers in general. Our thesis is that transparent, parsimonious TUIs are more learnable, and we have shown this to be true in a limited domain of comparison (with respect to a representative GUI).

Sub	TUI	GUI		Q1	Q2	Q4
0-1	4.58	7.97				
0-2	7.62	10.13		1	2	1
0-4	5.27	8.65		1	1	5
0-6	4.20	11.45	†	3	5	1
0-7	4.77	10.08		2	3	2
1-0	3.78	11.45	†	1	4	1
1-1	4.35	5.83		1	3	2
1-2	4.93	11.45	†			
1-5	4.28	11.45	†	1	3	2
1-6	4.28	11.45	†	1	4	1
μ	4.81	9.99		1.38	3.13	1.88
$\sigma\pm$	0.34	0.61		0.25	0.41	0.45

Table 1: Time on task data (in minutes) and select answers from the questionnaire.

RELATED WORK

One aspect of this work is the constructive coexistence of the physical and the virtual. In this respect this work is similar to the work of Wellner on the DigitalDesk[4] and is informed by Ishii's pioneering efforts in tangible user interfaces[3, 2].

Another aspect of this work is the desire for transparency and minimal cognitive demands on the user. In this respect it is inspired by work on sympathetic interfaces[1] and supported by the prior literature on perceptual interfaces[5].

CONTRIBUTIONS

This work focuses on a design principle for augmenting the traditional, physical tasks that consume a large part of everyday life. This is in contrast with much of the tangible interface literature that focuses instead on the useful, but different task of giving graspable manifestation to digital information [2].

This focus on existing artifacts demands tranparency and parsimony. Transparency seeks to minimize the impact on the constellation of existing human behavior that surrounds the host artifact. Parsimony supports transparency by minimizing the clutter that distracts attention, and avoiding inappropriate metaphors that add cognitive load. Parsimony is also important for preserving the aethetics of the host artifact. This minimization, avoidance of complexity, and preservations of traditional aesthetics all contribute to the increased learnability of our interface.

The go board is a particularly challenging artifact because of its highly refined aesthetic and behavioral constellation. We have presented an augmented go board that sucessfully showcases a minimalist ubiquitous computing design approach based on parsimony and transparency.

REFERENCES

1. M. Johnson, A. Wilson, B. Blumberg, C. Kline, and A. Bobick. Sympathethic interfaces: using a plush toy to direct synthetic characters. In *Proceedings of CHI*. ACM Press, 1999.

2. B. Ullmer and H. Ishii. Emerging frameworks for tangible user interfaces. *IBM Systems Journal*, 39(3&4):915–931, 2000.

3. J. Underkoffler and H. Ishii. Urp: A luminous-tangible workbench for urban planning and design. In *CHI*, pages 386–393, 1999.

4. P. Wellner. The digitaldesk calculator: Tangible manipulation on a desk top display. In *Proc. ACM SIGGRAPH Symposium on User Interface Software and Technology*, pages 107–115., 1991.

5. C. Wren, F. Sparacino, A. Azarbayejani, T. Darrell, T. Starner, Kotani A, C. Chao, M. Hlavac, K. Russell, and Pentland A. Perceptive spaces for peformance and entertainment. *Applied Artificial Intelligence*, 11(4):267–284, June 1997.

Principles of Context Inferences

Eleftheria Katsiri
Laboratory for Communications Engineering
University of Cambridge
http://www-lce.eng.cam.ac.uk
ek236@eng.cam.ac.uk

ABSTRACT

This paper discusses the principles of making inferences about a user's context from location time-series. These principles have been implemented and embedded in a statistical model, used in a Sentient [1] Notification Service. This model makes estimates about the probability of a situation occurring in the future, based on users' locations and past behavior and using methods such as Bayesian inference [2] and Hidden Markov models [3]. It is also used for modeling user activities, such as making coffee. This information is then used by a Context-Aware Notification Service [9], to notify all registered users who are interested in the particular activity.

Keywords

Context-Aware, notification service, activity modeling, Bayesian inference, Hidden Markov models

INTRODUCTION

Imagine the following scenario: at 8.15 am, on a Thursday morning, in an IT company, a visitor wishes to meet the Director of Finance. The director's secretary, queries the system for the probability that the Director will be in his office shortly. Receiving an encouraging answer, she asks the visitor to wait and types in her PDA *"notify me as soon as the Director of Finance appears anywhere on the premises"* Meanwhile, the Director has just arrived and enters the building by the garage. Immediately a notification is sent to his secretary of his arrival. While he is still in the elevator he receives notification that *"It is 98% certain that someone is making coffee"*. He can't resist stopping by the coffee-room for cup of coffee on his way to his office.

Activity modeling and likelihood estimations

Current location platforms [4], [7] and context-aware systems [5], [6] do not provide ways of inferring user-activity over quantitative data. E.g. a notification such as *"notify me when someone is making coffee"* is impossible with state of the art systems. Similarly, contemporary systems do not provide any likelihood estimation of a situation of interest that allows decision making. E.g. requests similar *to "I want to be registered for notification about whether Person X is in his office, only if there is a strong possibility that Person X will come in the office today*" are unfeasible.

Statistical models

An analysis of sightings in the Laboratory for Communications Engineering (LCE) [8] over a period of 72 hours has produced a set of predictions of people's movements. In this experiment, the data were produced by the Active Bat system [4]. The following figure depicts a prediction of where Person M is more likely to be "seen". Room 11 is Person M's office and Room 10 is the office he was moving into at the time of the experiment. The next section discusses the methodology of this experiment.

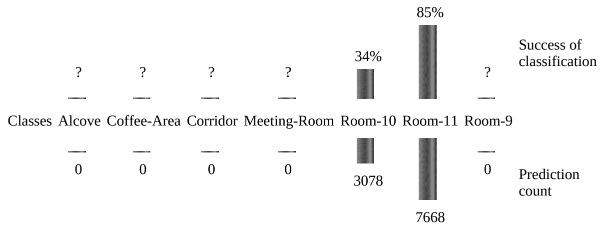

Figure 1: Probability of Person M's sightings during the day

Bayesian Predictive Discriminant analysis

Discriminant analysis is a tool that aims to construct a model for predicting the value of one discrete variable using other variables. The predictions use infinitively many models weighted by their probabilities to do the prediction. The following figure depicts a prediction of the probability of people's sightings in the coffee area during morning hours.

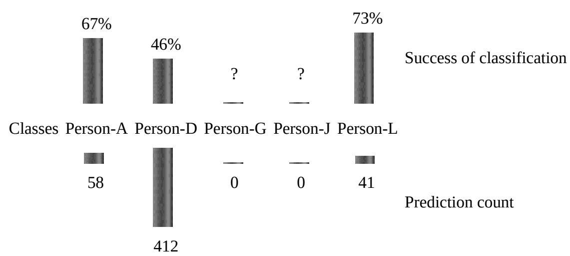

Figure 2: Probability of morning sightings in the coffee-area

In this case, while doing the cross validation, it was predicted 58 times that data item should belong to the class "Person-A" and 67.2% of these classifications were correct. So it is estimated that if the system predicts previously unseen data item to belong to the class "Person-A", there is

67.2% chance that this prediction is right. The reliability of this estimate can be rated by stating the fact that the estimate is based on classifying 58 items (11% of the sample) as members of the class "Person-A". Below each estimate (fig2) a bar chart indicates the percentage of the sample size used to calculate this estimate.

Preliminary Evaluation

We asked the people who were picked up by the system if they are frequently in the coffee area in the morning and both Person A and Person L confirmed to their regularly making coffee there in the morning. Person D, confirmed that he makes tea regularly, and therefore is also often in the coffee area in the morning. This observation led us towards looking into the following problem. Can we detect a user making coffee just by his movements? The following section explains how we solved this problem successfully.

Activity modeling using a Hidden Markov Model : Detecting the coffee-making process

We have designed a coffee-making detector, using a Hidden Markov Model with four states. We noticed that the coffee-making process in our lab requires a person to first approach the coffee machine and remove the jar. Next, one must go to the sink, empty the jar and fill it with fresh water from the tap. Then, one has to re-approach the coffee machine fill it with the water, and ground and press the start button. We furthermore assumed that if the user completes the coffee machine-sink-coffee machine cycle without leaving the combined area of the coffee machine and the sink, then there is very strong possibility that he is making coffee. The Hidden Markov model for modeling the coffee making process, is depicted in the following figure.

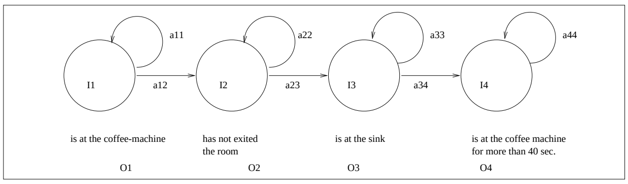

Figure 3: A Hidden Markov Model that detects coffee-making

EVALUATION

The outcome of the above experiment has been 4 inferences and corresponding notifications all of which were correct. Three people (Person L, Person A and Person P) have been correctly detected to be making coffee in a total 4 times within 72 hours. Two of these people (Person L and Person A), are regular coffee-makers as can be inferred by the predictability of their movements in the coffee area in the morning, when the coffee is usually prepared. (fig2) Person D, although is frequently seen in the coffee area in the morning, does not drink coffee, and therefore never makes any. Person P also admitted to her having prepared coffee, during the duration of the experiment.

ACKNOWLEDGEMENTS

I am grateful to Professor Andy Hopper for his precious guidance and support throughout this work. I would like to thank the University of Cambridge for supporting the funding of my PhD.

REFERENCES

1. Andy Hopper. *The Royal Society Clifford Paterson Lecture: Sentient Computing*,
2. http://www.cs.Helsinki.FI/research/cosco/Projects/NONE/SW/. BAYDA 1.31 – Bayesian Predictive Discriminant Analysis.
3. Rakesh Dugad. A Tutorial on Hidden Markov Models. *Technical Report*, Signal Processing and Neural Networks Laboratory, Department of Electrical Engineering, Indian Institute of Technology, Bombay, 1996.
4. Andy Harter, Andy Hopper, Pete Steggles, Andy Ward and Paul Webster. The Anatomy of a Context-Aware Application. *In Proceedings of the fifth annual ACM/IEEE* international conference on Mobile Computing and Networking August 15-19, Seattle, 1999
5. Location-Aware Information Delivery with ComMotion, Springer-Verlag, 2000
6. Dey K. Anind and Gregory D. Abowd. CyberMinder: A Context-Aware System for Supporting Reminders. *In CHI ,2000*
7. http://www-lce.eng.cam.ac.uk/qosdream/. QoSDREAM: Quality of Service for Distributed Reconfigurable Adaptive Multimedia
8. http://ww-lce.eng.cam.ac.uk Laboratory for Communications Engineering
9. D. Ipina and E. Katsiri. A Rule-Matching Service for Simpler Development of Reactive Applications. *In Middleware2001*. http://computer.org/dsonline/0107/features/lop0107.htm

Plan-Aware Behavioral Modeling

Matthew Lease
Intel Research Lab @ Seattle
mlease@intel-research.net

ABSTRACT
One exciting promise of ubiquitous computing is that increased fidelity in natural world sensing will enable computers to better serve people via responsive or anticipatory behavior. We believe better behavioral modeling of users can play an important role in helping to make this possible. We propose the technique of plan-aware behavioral modeling to reason about deviations between a user's intended plan of behavior and behavior actually observed. We have been investigating this general technique in the context of supporting a specific activity: musical rehearsal and performance. Preliminary evaluation suggests the technique merits further investigation.

Keywords
Behavior modeling, Bayesian network, proactive

INTRODUCTION

The advent of ubiquitous computing has introduced exciting new opportunities for computers to both become significantly more aware of their physical environments and to actualize tangible behavior within these environments. A common hope is that these new capabilities will enable computers to better serve people via responsive or anticipatory behavior. In order to achieve this goal, better behavioral modeling of users is necessary. Fortunately, continuing improvements in sensing technologies provide increasing opportunities for computers to monitor user behavior more accurately and in new ways. By analyzing the data collected from this monitoring, detailed behavioral models may be constructed. For example, when a user performs a certain action, a computer could consult that user's behavioral model to reason about what caused the user behave in that way (i.e. diagnosis). Such a model could also be leveraged to help predict a user's future actions, enabling the computer to provide proactive assistance appropriate to the expected behavior. Finally, such models can be studied to improve our own understandings of human behavior.

We have been exploring behavioral modeling for situations where a computer knows a user is trying to perform some given task and can monitor the user's progress towards completing the task. More specifically, the user has a plan of successive steps he intends to perform but may deviate from it due to limitations in execution accuracy. In order to restrict our focus to reasoning about plan-execution inaccuracies, we assume any deviations are accidental and rely on the user to clearly communicate intentional changes to the plan. We believe there are many interesting scenarios in which these conditions are met. Our goal is to enable ubiquitous assistance in such scenarios in order to improve user accuracy in executing a desired plan.

To provide a real-world example, consider in-vehicle route planning assistance [1]. In this scenario, a driver asks his vehicle's telematics system for driving directions to a given destination address. The system determines the vehicle's current location via GPS and plots an appropriate route, automatically updating the directions whenever the driver makes a wrong turn, etc. While such automated route re-planning is obviously a valuable feature, the driver would probably prefer to have not made the wrong turn at all. By real-time monitoring of such mistakes and recognizing patterns of their occurrence, we believe it might be possible to predict an upcoming wrong turn and proactively tailor interaction with the driver, providing sufficient assistance to prevent the mistake's occurrence.

SUPPORTING MUSICAL PERFORMANCE

We have been investigating the application of plan-aware behavioral modeling to providing ubiquitous assistance for musical rehearsal and performance. We chose this application because it provided a real, existing activity to support and a novel environment for evaluating ubiquitous computing techniques. In this application, the user is a musician and the plan of behavior is represented by a composer's score, which indicates a sequence of notes to be performed. While others have studied musician mistake patterns relevant to cognitive planning [2], our interest is in exploring those mistake patterns which allow an opportunity for providing automated assistance. Other work has shown one can estimate in real-time a musician's position within a given score [3], though the estimate produced will contain a degree of uncertainty which must be accounted for when using the estimate. We propose augmenting existing estimation techniques with plan-aware behavioral modeling in order to predict whether an upcoming note will be performed correctly and to automate diagnosing the cause of detected mistakes. In the first case, if the model suggests the musician will have difficulty performing a given note, we would like to offer proactive assistance to prevent the mistake's occurrence. The success of a solution to this problem may be evaluated by comparing a musician's performance accuracy with and without the assistance enabled. In the second case, we would like to offer responsive assistance to prevent the mistake from reoccurring. The success of a solution to this problem may be evaluated similarly.

OUR APPROACH

We have focused our initial analysis on solo piano performances. Such performance data may be readily acquired due to the commercial availability of pianos supporting real-time output of MIDI events describing performer-piano interaction. Given a tolerance for additional sensing uncertainty, existing techniques may also be employed to extract similar data in real-time from ensembles of arbitrary instruments [3].

We will be using a Bayesian network [4] to exploit knowledge of the plan in addition to musical context and historical observations of behavior. Bayesian Networks make it practical to perform inference under uncertainty by expressing independence assumptions between parameters. As the musician plays, any mistakes made and the musical context of those mistakes will be used to update the network's prior and conditional probabilities. We expect these probabilities and the structure of the network can incrementally gain complexity as automated learning and domain expertise iteratively influence and inform one another. Model accuracy will be evaluated using standard model validation techniques.

In order to provide the musician with ubiquitous assistance, we will ask him to read the score from an electronic display. As he performs, we will automatically generate digital annotations consistent with how the musician would manually annotate a paper copy of the score during rehearsal (Figure 1). These annotations will be created both proactively, to prevent mistake occurrence, and retroactively, to prevent mistake re-occurrence.

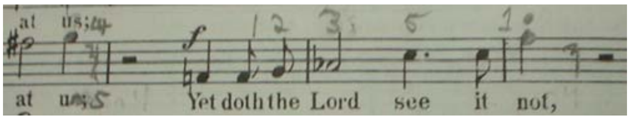

Figure 1: Annotations denoting articulation & chord structure used in a performance of Mendelssohn's Elijah

PRELIMINARY RESULTS

Our results thus far consist of analyzing an existing data set we were able to obtain from another researcher [5]. Previous analysis of this data set was restricted to evaluating methods for estimating score position. In contrast, our interest was to study musician behavior and to begin developing tools for analyzing the data in this light. The data set consists of both MIDI data and a subjective coding of mistakes as identified by a domain expert. We began our analysis by exploring the many existing tools which support interactive MIDI playback and visualization. After surveying the visualization options available, we found piano roll notation (see Figure 2) to be the most valuable in examining precise details of note onsets and durations. However, we also found it necessary to complement our use of piano roll notation with a format more closely resembling the familiar, traditional Western staff notation in which the original score was written. Unfortunately, traditional staff notation presents an abstract

Figure 2: A musical excerpt shown in piano roll notation

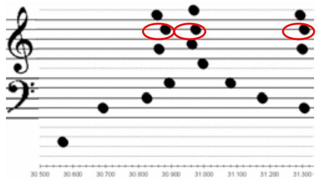

Figure 3: Modified staff notation shows precise timings in a familiar format

notion of time which can only coarsely describe the actual timing of performance events. After unsuccessfully trying to find an existing tool which addressed this mismatch, we adapted a spreadsheet program to create our own modified version of staff notation (Figure 3). As with the piano roll notation viewer, we were able to zoom in on regions of interest to better visualize precise timing of performance events. Use of these complementary approaches has begun to reveal some interesting mistake patterns for analysis. For example, one performer tended to be late with his middle finger when playing chords with his right hand (Figure 3).

FUTURE WORK

In addition to the applications mentioned thus far, we would also like to investigate how plan-aware behavioral monitoring could enhance ubiquitous computing support for several other scenarios: distance and self-education, fabrication and maintenance, laboratory work, and assisting the elderly and disabled.

ACKNOWLEDGMENTS

We would like to thank Ron Patterson and Henry Kautz for their assistance with this work.

REFERENCES

1. M. Ness and M. Herbert. A prototype low cost in-vehicle navigation system. Vehicle Navigation and Information Systems Conference, 1993, 56-59.
2. C. Drake and C. Palmer. Skill acquisition in music performance. Cognition 72 (2000), 1-33.
3. L. Grubb et al. Automated Accompaniment of Musical Ensembles. *Proc. of 12th National AAAI (1994)*, 94-99.
4. E. Charniak. Bayesian networks without tears. *AI Magazine 12 (1991)*, 4, 50-63.
5. H. Heijink et al. Data processing in music performance research: Using structural information to improve score-performance matching. *Behavioral Research Methods, Instruments & Computers 32 (2000),* 4, 546-554.

SiSSy – Smart-its child Surveillance System

Henrik Jernström
PLAY Studio, Interactive Institute
Hugo Grauers gata 3
41133 Göteborg, Sweden
henrik.jernstrom@interactiveinstitute.se

ABSTRACT

In this paper I describe what functions are necessary to support people responsible for children by using computer-augmented artifacts. The suggested solution, *SiSSy*, is based on using input from *Smart-It* devices to prevent children from finding themselves in danger by alerting parents of the situation. To collect information of how the system should be designed I gathered a group of possible future users to participate in a brainstorming workshop. I report the results from the workshop and how these have influenced the design of SiSSy.

Keywords

Smart-Its, Ubiquitous Computing, Context Awareness

INTRODUCTION

The vision of Ubiquitous Computing has the ambition to relieve us from the demanding and intruding properties of the computer, get it out of sight - out in the periphery of our lives, while still being able to reap the benefits of computational services. Computers should conform to humans and humans' needs, rather than the other way around. One solution for this is to attach computers to objects in the real world and connect them to each other to make them support people's requirements whenever and wherever they need it [2]. Ubiquitous Computing now seems within our grasp by the increasing growth of computer power, smaller and more accurate sensors, greater communication bandwidth and less power consuming devices. One research project that tries to realize some of the possibilities of Ubiquitous Computing is Smart-Its.

The Smart-Its project

The Smart-Its project [3] is collaboration between TecO Karlsruhe, ETH Zürich, VTT Finland, University of Lancaster, PLAY Interactive Institute and FAL Viktoria Institute and is funded by the European Union's Disappearing Computer Initiative. The goal is to make inexpensive computational devices for post-hoc computational augmentation of everyday things. Smart-Its are devices that perceive their environment with various sensors, communicate with peers, and have adaptable and modifiable behaviors. The vision of these devices is for them to be as cheap, non-intrusive and common as modern radio-tags so that they can be used to build and test ubiquitous computing scenarios by augmenting a great number of everyday objects and places with Smart-Its.

THE SiSSy

Children disappear from their parents in shopping malls, crowded places and sometimes children even get lost from daycare. It is impossible for parents and daycare personnel to always have full control over what children are up to, where they are going, and if there is more than one child - who is doing what. The Smart-its child Surveillance System, SiSSy, is an approach to tag children and parents with Smart-Its devices which can sense the environment and determine whether a situation is dangerous or if the child is engaged in something hazardous. The idea is to make SiSSy easy to use without requiring any specific infrastructure such as large transmitting equipment or satellites communication. The system is by its nature primarily suited for children between ages two and five and will, implemented and ready, be tested on children of that age.

From the parents point of view

To ground my research and find the relevant functionality for SiSSy I chose to engage possible future users by arranging a brainstorming workshop. Five parents took part in the workshop, which lasted for three hours including introduction and scenario presentations. The comic strip-like scenarios were used to illustrate the intention of SiSSy and give the participants a notion of the Smart-Its capabilities. To make the discussion easier and give it a more concrete form a map over a playground was used with small toy figures exemplifying children and parents. The session was recorded with a video camera.

Figure 1: Parents discussing different scenarios using a map and toy figures

The participants' opinions are summarized as follows:

- It is important that SiSSy warns *before* anything has happened. It has to sense sudden and rapid movement away from other devices, which should trigger an alarm.
- SiSSy must be totally reliable to have parents' complete trust. However, they also recognized that if this is well accomplished, there might be a risk of parents becoming inattentive to what their children are doing.
- SiSSy should be able to sense moisture to prevent accidents in water, altitude to detect dangers of falling, and car exhausts and traffic noise to make it possible to decide whether the child is near traffic or not.
- The behavior of SiSSy should be easy to modify. Different sensors can be used for different purposes in different situation and SiSSy should support this. It is more likely that SiSSy will be used if it works in more than one place.
- Smart-Its could also be used for tagging the surroundings and not just for tagging children and adults. The benefits would be to create secure and non-secure areas when a family is temporarily at a location e.g. at a playground, on a picnic, etc.
- There were also arguments for making SiSSy absolute position aware, motivated by that if you can't find a missing child there is really no point with using SiSSy.
- Wristwatch-like device for the parents with a vibrator alarm with preferably text messages that calls for the parent's attention [1] and informs about the situation, who, were and what.
- The child device should be attached properly on the child's belt and must be silent. This to prevent any risk that SiSSy could cause disturbance in the psychological development for a child from being under surveillance.

Figure 2: Smart-Its assembled in a cassette box.

Smart-Its description

The Smart-Its prototype used for SiSSy is designed and manufactured by TecO, University of Karlsruhe. It consists of two main boards: an RF-board equipped with a 868.35 MHz radio transceiver and a sensor board that can detect sound (using a microphone), touch pressure, light level, 2-dimensional acceleration and temperature. An I2C bus connects the boards together and both boards are equipped with a microchip PIC microcontroller running at 20MHz. The system allows for ad-hoc networks to be formed spontaneously, and has support for sensor data processing and context recognition.

Implementation

My work will progress in order to satisfy the user requirements as far as it is technically possible using the current Smart-Its platform (making absolute positioning unfeasible). The physical characteristics of SiSSy devices are shown in my design prototype (figure 2) with the intention to attach the box to children's belts. Parents' devices are connected to a wristwatch with a cable, used for attracting attention when an alarm is triggered (cf. [1]), using vibration and colored LEDs. In addition, several stand-alone devices have been constructed. SiSSy's behavior will adapt to the situation and decrease or increase the broadcast radius whether the situation is dangerous or not (as one of the requirements for initiating an alarm is the loss of contact between the devices). The system will also be able to route information through intermediate devices to make the functionality more flexible. The first SiSSy ready for use will be evaluated by the same parents who took part in the above described workshop. The test will be carried out on a playground similar to what the workshop map represented.

DISCUSSION

The perfect SiSSy application has to be completely reliable for it to help parents and personnel working with children. Complex contexts should be detected and understood by the combination of different sensors, their input and appropriate algorithms and the system should determine whether it is dangerous for children or not (falling, traffic, crying etc). Because of the inherent complexity of the problem, it is most likely that SiSSy will primarily be an experiment to explore the Smart-Its concept and bring up important questions for the use of Smart-Its in a wider sense. The goal is to demonstrate a complex application based on ad-hoc networking and real-time information from distributed and mobile devices.

ACKNOWLEDGMENTS

The Smart-Its project is funded by the Commission of the European Union as part of the research initiative "The Disappearing Computer" (contract IST-2000-25428).

REFERENCES

1. Hansson, R., Ljungstrand, P. and Redström, J. Subtle and Public Notification Cues for Mobile Devices. In: *Proceedings of UbiComp 2001*. Springer.
2. Weiser, M. "The world is not a desktop". *Interactions*; January 1994; pp. 7-8. ACM Press.
3. http://www.smart-its.org/, (14 September 2002)

SoundPryer: Joint Music Listening on the Road

Fredrik Axelsson, Mattias Östergren
Interactive Institute Mobility Studio
Karlavägen 108, 104 50 Stockholm
+46 (0)8 783 24 80
{ fredrik.axelsson, mattias.ostergren }@interactiveinstitute.se

ABSTRACT
The roads become increasingly congested. People in cars need more entertainment. Eavesdropping the music played in other vehicles is an interesting alternative. We introduce the SoundPryer prototype, which is a mobile music-sharing tool. It is an application for handheld devices equipped with wireless LAN adapters and thus provides sharing in a wide variety of mobile settings. The application accommodates situations ranging from walking to driving a vehicle, but specifically adapted for the latter. Meetings in traffic are highly contingent and therefore SoundPryer is designed as a peer-to-peer application over ad hoc networks. Sharing is accomplished by streaming MP3 files.

Keywords
Music sharing, Mobile Computing, Ad hoc Networking

INTRODUCTION
We believe jointly listening to music shared between cars on the road would be an entertaining addition to the experience of car driving. Sharing music is an important activity in the everyday handling with music. It is a major factor when establishing the identity of a person or a group. It also increases a person's musical appetite [1, 2]. Recently, sharing of digitally encoded music, such as MP3 files, has gained a lot of attention. Through the use of Internet peer-to-peer applications, such as Napster and Gnutella, a large body of people gained access to each other's files. When spending time in a vehicle, listening to music is a popular entertainment. In a recent report [1] it is found that, in a group of music enthusiasts, on average, 82 % of the time in a car, is spent listening to music. However, in a car the only possibility to share music is to play it loud to other passengers. In traffic there is plenty of people around, and it is likely that they listen to music, but there are no means to share music experiences. We believe a tool to enable sharing of digitally encoded music among people in vehicles would be interesting for two reasons. First it will expand the benefits of sharing to mobile situations. Second, it will make road use more enjoyable. A user will hear new music and associate its source with vehicles in the immediate surrounding. Imagine he or she is passed by a pick-up truck and thinking: "Sigh. Another country fan." When hearing the music being played in that vehicle, the user is surprised: "That driver is listening to Disco!" Napster and Gnutella are instead designed to share files among users connected to the Internet and doing it anonymously. These applications are focused on searching and retrieving files and the music experience of playing a downloaded file, is unbounded in time and place. SoundPryer aims at sharing music experiences when co-located in traffic. It tackles the challenge to facilitate sharing in order to provide joint music experiences, during brief and opportunistic meetings, where Internet access is not available. We are investigating handheld computers equipped with wireless LAN (WLAN) interface cards as the platform for sharing music in traffic. We are currently collecting requirements and experimenting with various designs in order to ultimately craft and evaluate a prototype application for the PocketPC operating system. It will use IEEE 802.11 compatible WLAN cards. Sharing is accomplished by streaming MP3 files over ad hoc connections.

RELATED WORK
Music sharing has received some attention by the research community, since it yields some collaborative properties. The principal systems suggestions are: Music Books, Music Buddy [2] and ShoutCar [3]. Music Books, explores the fact that music encoded on tapes, CDs etc. possesses tangible properties. By coupling a small book with music resources on the Internet, tangibility and the flexibility of digitally encoded music is combined. Sharing is accomplished by handing someone one of these books. Music Buddy aims at stimulating sharing and the socializing that occurs around it. The user uploads a listing of his or her music collection to a web server. He or she is then clustered together to other users according to musical taste. By browsing this cluster or clusters related to it, the user learns about new music and new acquaintances. ShoutCar lets the user listen to samples of music while driving. The samples are stored in servers and they are accessed through a web interface. First, a user assembles a selection of samples and later, in the car, they may be played.

DESIGN REQUIREMENTS
We believe that handheld devices equipped with WLAN cards will enjoy widespread usage in the future. Such devices will become personal accessories and be present in many different situations in life. Today many networked applications for these devices assume some properties of the mobile situation in order to work, such as connectivity with an office LAN for Internet access. However, we believe applications should accommodate use in a much wider spectrum of situations. In our case mobile music

sharing will be available irrespective of the mobile setting at hand; may it be walking by a road, or driving in traffic.

Traffic is a highly mobile situation. High speed makes meetings among people very contingent and a sharing application must be able to act promptly on them. Such application requires three properties of the network. Besides a high and predictable bandwidth, the networking must also be of low delay and the range of the wireless network must be at least 100 meters. This is an ideal situation for mobile ad hoc connections over WLAN in contrast to current and future mobile telephony and Bluetooth networks. The transient nature of ad hoc networks implies that distributed applications cannot rely on infrastructure, such as routers or servers; however, a peer-to-peer software architecture is fitting.

The functionality of SoundPryer is dependant of the number of co-located peers. When co-located, all receivers are dependant of the number users actually actively providing music. In order to maximize the possibility that there is something to listen to, SoundPryer requires the user to assemble a play-list before he or she may activate the playback mechanism. Each item in this list is then immediately and always available for other peers, through streaming, whenever they are selected for local playback.

There must be a balance in playing for others and listening to what others play. The playback mechanism in SoundPryer is provided through two modes: *manual* and *automatic*. The manual mode is useful when the user wants to control the music selection. It is ideal for situations when being interested to listen to local files. Still, it supports toggling to a remote source. This provides quickly checking out what someone else in the vicinity is listening to. In the automatic mode SoundPryer provides an experience for the user. This mode plays a mix of local and remotely sourced files however it always strives at playing music from remote sources. We suggest a rule to prevent a situation where, for instance, two peers are in automatic mode and repeatedly attempt to select each other. A user may only manually select and the application may only automatically select, a peer that plays a local file. Now, if all peers are in automatic mode, all users but one will be able to select a remote peer.

When in a sharing situation it is essential to be aware from whom you are receiving music. The challenge is to obtain clear and concise coupling between real-world properties relevant to tell who you are and how they are reflected by the application. In SoundPryer, a user is represented by three properties: a nickname, a silhouette of a vehicle, and the color of it. When choosing a silhouette, a user will select among a predefined set of stylized shapes. The choices correspond to, for instance, a convertible, a pick-up truck, a motorcycle, a bike etc. There is also a humanoid figure, to accommodate usage when not driving.

Supporting activities related to music-sharing experiences is important. The SoundPryer logs the date, time, silhouette, color, nickname, the title and artist of all the remote streams received and played. A user may also add such entry to a "favorites" database. Having these logs is helpful, when later trying to find a track at a retailer.

In many situations the user of SoundPryer will be occupied with other cognitive demanding tasks, such as driving in intense traffic. The GUI is designed to mostly be in the fringe of the user's attention and not interfere with the situation at hand. To effectively display the status at a glance, it uses large controls of high-contrast colors. Making use of the touch sensitive screen and large interface controls will accommodate easy finger-based interaction.

THE SOUNDPRYER PROTOTYPE

Internally the playback mechanism is represented by four states. The states represent the four combinations of mode and the source of music. The states are: *manual/local*, *manual/remote*, *automatic/local* and *automatic/remote*. In manual/local, SoundPryer plays and broadcasts a file in the play-list. In this state, the user may manually select a peer present in the list of available remote peers. Then the playback switches to the manual/remote state. In manual/remote, the application receives and plays the broadcast stream from a remote peer. When connectivity is lost, or deliberately terminated, the application resumes to manual/local state. The automatic/local state is similar to the manual/local state. However, when a remote peer appears, the playback and broadcasting of a local file is stopped. The source of playback is then automatically changed from local to remote and the application enters the automatic/remote state. In the automatic/remote state, when connectivity is lost, a new remote peer is selected automatically. As long as there are at least two peers present, after a fixed interval, SoundPryer terminates the current connection and randomly selects a new peer. If no peers are present, operation resumes to the automatic/local state.

ACKNOWLEDGMENTS

We would like to thank our colleagues for valuable comments on our work. SITI funded this project.

REFERENCES

1. Brown, B. Geelhoed, E., and Sellen A. The Use of Conventional New Music Media: Implications for Future Technologies. In *Proc. of INTERACT'01*, pp. 67-75, Tokyo, Japan.
2. Brown, B. Geelhoed, E., and Sellen A. Music Sharing as a Computer Supported Collaborative Application, In *Proc. of ECSCW'01*, pp. 179-198, Bonn, Germany.
3. Åkesson, K-P., and Nilsson, A., Designing Leisure Applications. To appear in *Journal of Personal and Ubiquitous Computing*, vol. 6, 2002, Springer-Verlag.

SpeakerPhone: a platform for dynamic human-navigable soundscapes

Jonah Brucker-Cohen and Stefan Agamanolis
Human Connectedness group
Media Lab Europe
Sugar House Lane, Bellevue, Dublin 8, Ireland
{jonah, stefan}@mle.ie

ABSTRACT
SpeakerPhone is a high-density array of speakers for presence applications. Scenarios for its use include networked soundscapes, data visualization and spatialization, interactive location-based narratives, and customizable sound landscapes. SpeakerPhone is an enabling technology that focuses on the opportunity to provide rich immersive experiences through highly tailored spatial sound environments. It differs from past research in its modular structure for realizing a diverse range of applications and its emphasis on revealing the physical pathways of data through spatialized audio. The project attempts to contribute to an overall infrastructure for human-navigable technology-mediated spaces.

Keywords
Spatialized audio, data visualization, location-based storytelling, responsive audio, telepresence

INTRODUCTION
SpeakerPhone is a framework for creating hybrid spaces combining audio spatialization, data visualization, telepresence, narrative environments, and audience participation. Unlike traditional audio spatialization techniques in which listeners must be passive, SpeakerPhone enables the creation of multi-layered, architecturally integrated, physically-navigable audio soundscapes. These dynamic soundscapes can be controlled telematically or through direct sensor feedback. SpeakerPhone also enables modes of data visualization for information moving across networks as a way of increasing awareness of these concealed information pathways.

RELATED WORK
SpeakerPhone's focus on ubiquitous spatialized audio, data visualization, telematic systems, and narrative spaces uncovers a diverse amount of related work. Projects range from Perry Cook's *NBody* musical instrument modeling project [1], which recreated instrumental sound through specifically arranged speaker arrays, to Interval Research's *ThunderWire* [2], an audio-only media space for networked audio transmissions between remote locations. *VisualRoute* [3] and the Unix *traceroute* utility illustrate methods of accessing and visualizing the physical path of information across international networks. Finally, the interactive Argentina-based performance *De La Guarda* [5] demonstrates the power of a mobile audience mediated by technology.

Benefits of SpeakerPhone
SpeakerPhone's approach differentiates itself from past research by both using sound as a medium to reveal the path of data in a physical space as well as by focusing on scenarios for the capture, transmission, and precise layering of spatial audio information from one architectural environment into another.

TECHNOLOGY
The SpeakerPhone prototype consists of a high-density array of 24 miniature speakers installed in a room at our laboratory [Fig. 1]. These speakers can be controlled via a computer interface to a microcontroller that instantaneously switches between relays that address different speakers in the array [Fig. 2]. The SpeakerPhone architecture provides the ability to create customized effects, such as making a sound travel around the room in various directions and rates of speed, or having a sound follow someone moving through the room.

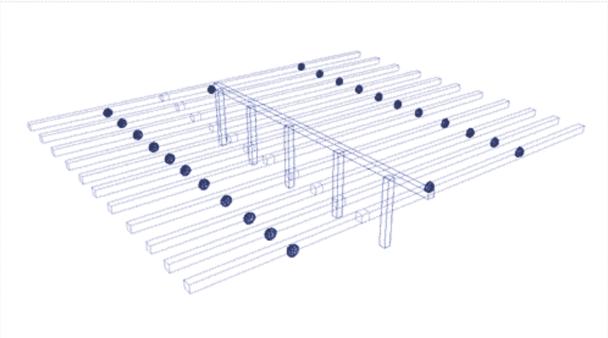

Figure 1. 3D model of the speaker arrangement in a room at our laboratory

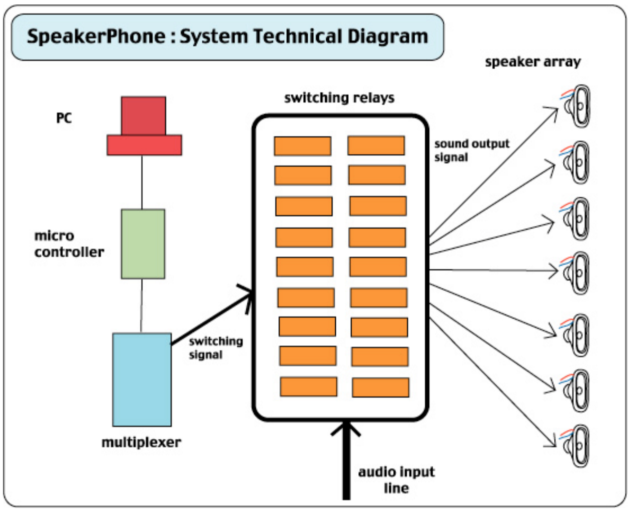

Figure 2. SpeakerPhone technical system

SCENARIOS

SpeakerPhone provides an inexpensive modular platform for the realization of a diverse range of potential applications such as networked soundscapes, data spatialization, interactive narrative environments, and customizable sound landscapes. Below are a few specific scenarios in which the system might be employed.

1. Data Spatialization: Speakerphone uses sound as a means of exposing data moving through physical spaces. SpeakerPhone attempts to free information from the hidden pathways of wires. The movement of sound becomes an audible illustration of the information overload infiltrating our daily connected lives.

2. Networked Audio Mapping: When various necessities dictate our being separated from our friends, family, and colleagues, the need to maintain some kind of contact becomes more urgent. By creating a continuous two-way ambient auditory link between a pair of similar spaces through accurate sound layering, we can create a hybrid "connected" space. Such a scenario would require mirror arrays of microphones and sophisticated echo-cancellation technology.

3. Narrative Audio Targeting: Escaping the passive audience model, SpeakerPhone enables dynamic narrative sound applications and allows for a mobile relationship between audience and content in a story, presentation, or performance. The audience can help drive the narrative because their investment in the narrative becomes both physical and mental.

4. Smart Speakers: Like the Audio Spotlight [4], SpeakerPhone can focus audio on specific locations in a space and transform the audio landscape based on sensor feedback of various kinds. With further enhancement to the technology, individual speakers in the array could be made to sense their surroundings and create dynamic soundscapes based on their proximity to each other or other objects.

5. Telematic Audio Control: Remote networked control of audio placement would allow for collaborative audio environments created across distance by multiple performers or participants.

6. Pathways of Data: SpeakerPhone's node-based architecture enables users of the system not only to dictate the final location of audio but also the path it travels to a specific destination. In this way, the system suggests the ability to customize the routes taken by other kinds of data in communications networks.

FUTURE RESEARCH

Future research on SpeakerPhone will include building a new type of speaker architecture that provides independent computational capability within each speaker node, and forming a self-organizing ad-hoc networking protocol for controlling playback and movement of audio from node to node. Other improvements include equipping each speaker with its own audio processing circuitry and adding a wireless transceiver for collecting new audio data or broadcasting the audio source or stored sound in each speaker to other nodes. This would simplify deployment of the system in new environments and enable it to be controlled via standard wireless networking protocols.

SUMMARY

SpeakerPhone enables the creation of dynamic, physically-navigable audio spaces in which listeners do not have to wear headphones or remain still in one location. The scenarios presented illustrate a variety of application possibilities that relate to interactive narrative, data visualization, spatial audio mapping, collaborative live performance, and ambient communication. The project emphasizes the integration of audio with architectural environments both as an enabling technology for presence applications and as a way of enhancing our understanding of how we interact with digital information in the physical spaces we inhabit.

ACKNOWLEDGMENTS

This research has been supported by sponsors and partners of Media Lab Europe.

REFERENCES

1. Cook, P. R., Essl, G., Tzanetakis, G., Trueman D. N >> 2: Multi-speaker Display Systems for Virtual Reality and Spatial Audio Projection, ICAD '98, Glasgow, UK.
2. Hindus, D., Ackerman, M., Mainwaring, S., Starr, B. Thunderwire: A Field Study of an Audio-Only Media Space, in *Proc. CSCW '96*, ACM Press, 238-247.
3. Visual Route, http://www.visualroute.com (last visited 12 June 2002).
4. Pompei, F. J. The Use of Airborne Ultrasonics for Generating Audible Sound Beams, in *Journal of the AES 47*, 9 (September 1999), 726-731.
5. De La Guarda, http://www.delaguarda.com (last visited 12 June 2002).

Teaching a practical Ubicomp course with Smart-Its

Michael Beigl, Philip Robinson,Tobias Zimmer, Christian Decker
TecO, University of Karlsruhe
Karlsruhe, Germany
Email {michael, philip, zimmer, cdecker}@teco.edu

Abstract
We propose the Smart-Its platform as a platform for teaching Ubicomp in practical student courses. We present four experiments that were undertaken during summer term 2002 along with the outcome. Applications, context, sensor and Wireless ad-hoc communication were the key issues in these experiments. We also introduced the use of the Smart-Its hardware and development platform as tools for such teaching activities.

Keywords
Teaching, smart devices, context, sensors, communication.

INTRODUCTION
Teaching Ubiquitous computing topics in a practical student course requires a platform that is easily used yet provides a realistic environment for students to solve real-world problems, while experimenting. This paper therefore presents Smart-Its as a platform for conducting such experiments in student courses. Smart-Its [1] are small devices (1.5x4 cm) containing sensing (various), computing (5 MIPS) and communication (via RF) capabilities. These Smart-Its are intended to be attached to everyday objects so as to enhance the functionality of the object, deliver information (collected through sensors) to other devices and to produce acoustic or visual output. The Smart-Its platform also includes a software library and development support that is simple to program when implementing Ubiquitous Computing scenarios. The platform is powerful enough to allow programmers to implement algorithms and rather complex context detection. Smart-Its provide access to functionality like sensors and communication through libraries and operating system functions. For our practical course we identified four main topics where Smart-Its are helpful. Smart-Its are used in the experiments as **Smart-Its nodes** for sensing data in the environment, for computing context out of these sensor data, for communicating these data to other nodes, and for building and testing applications. For the experiments, Smart-Its nodes had to be attached to everyday objects by students and then had to be programmed. The paper briefly describes the content of the course held first in summer term 2002, experience collected and also the development environment and technical parameters.

Related Work
Up to now there are only a few platforms available that can serve as a basis for a practical course. One of them is the Berkley Mots [2] system and the TinyOS. In contrast to the Mots Smart-Its provide some better performance values (e.g. faster communication) and also a development environment including wireless programming with systems in place over the RF channel. Smart-its also support analysis of context and sensor data via backend systems. First ideas for tools for building ubiquitous computing systems were presented in the first workshop on Ubicomp Tools at Ubicomp 2001 [3]. However, only a minority is available for use now, e.g. the CoolBase or the Phidget system. The drawback of these systems is their reliance on a wired or simple IR-base communication architecture, which restricts their application range.

Teaching a practical course in Ubicomp
In an example practical course we identified four main topics to be taught; the premier topic was based on building and testing applications in the area of ubiquitous computing, while the remaining three topics demonstrated Ubicomp-enabling technologies. The second dealt with context and sensor technology, the third with context communication between different nodes in a network and the fourth with basic communication and energy issues. The students conducted general experiments for topics 2-4, followed by a more detailed experiment based on topic 1. The motivation to use the Smart-Its platform and the expected outcome of these topics are described in the paragraphs below:

Ad-hoc application experiments

This part of the course was dedicated to collecting practical experiences with applying technology into the everyday environment. In the experiment students were faced with a complex and distributed application, where multiple nodes worked together to fulfill a task. The major outcomes were that students learnt to develop distributed applications, how to address security and privacy issues and where to implement Smart-Its nodes. The student also had to solve practical problems associated with attaching Smart-its nodes to objects and collecting information. A definite advantage of using Smart-Its is the short implementation time for small and medium-sized application settings. Such settings require nodes that are not fixed to wired sensor structures (with long set-up times) and are not reliant on a complicated infrastructure. Furthermore, as Smart-Its computation is processed onboard, applications and experiments are not confined to the lab, plus they can be implemented within a restricted time frame, having no previous knowledge of the technology.

Context and sensors

In this experiment students learnt about the available sensors, which sensors or combination of sensors can be used to retrieve a context and what algorithms are appropriate for use in small sensor nodes. They also became familiar with the characteristics of sensors, e.g. their required sampling rate, the energy consumption and their error rate. The Smart-Its sensor board provides a general set of sensors that can be used for sensor-based applications and for experiments on deriving context from sensor data. Such experiments always require information retrieval yet, due to the data transmission deficiencies of wireless networks, they often require processing of data at the sensors' location. Smart-its support this by providing a complete computing and sensor infrastructure integrated into the boards.

Abstract context aware communication

The outcome of this experiment was a general understanding of how to express context information for use in other applications, how to filter and combine them into new (richer) types of contexts and how to communicate them between applications. Smart-Its support this through a context protocol layer and an application programming interface (API). This layer is intended to be extended by new types of contexts that may result from previous work on contexts retrieval with Smart-Its.

Network and energy issues

As most of the energy in a node is spent on communication, another consideration in the experiment was the implementation of energy saving and energy aware protocol extensions. A major outcome of the experiment was the understanding of communication issues for ubiquitous computing networks, their characteristics and solution strategies. Smart-Its support access to different layers of communication; In Smart-Its the network layer can be accessed through a simple API providing all functionality to implement various types of algorithms including energy aware protocol enhancements.

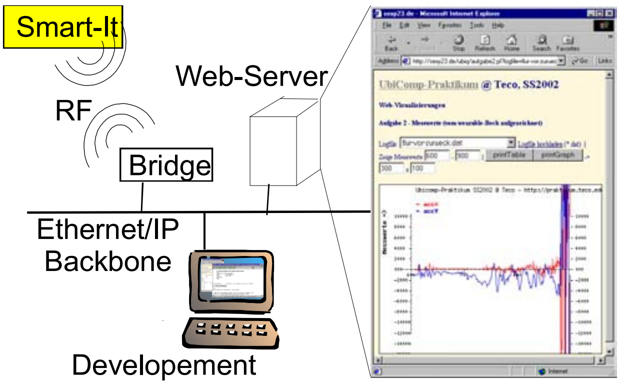

Smart-Its Development and Technology

The Smart-Its hardware consists of a generic RF communication board connected over I²C to a sensor board (total size: 4x1.5x1 cm), on which application development is done. Wireless broadcast communication facilitates interact with other Smart-Its in their vicinity (at 125kbit/s), and optionally with servers over a RF-Ethernet/IP bridge. The backend is used for analyzing sensor data through a tool running on a Web-Server (figure 1) while the experimentation process is running. The programming environment (also connected via Internet) consists of the Smart-Its Software Development Toolkit (SDK), compiler and program download utility. The latter one allows you to download compiled programs directly onto a Smart-Its from the development PC.

Smart-Its provide sensors for collecting light, pressure, temperature, acceleration and audio information. The freely programmable microprocessor on the Smart-Its provides about 12kbyte space for programs and about 8kbyte for data. Sensor values can be accessed by simple functions using the SDK. To share information with other Smart-Its, programs can communicate and express sensor values and contexts via APIs. A subscription mechanism of the communication subsystem allows one to watch for selected or all information on the communication channel.

Experiences with Experiments and outlook

During the practical course 4 experiments were carried out. In "Night&Day" students automatically adjusted the time-of-day setting of the Smart-It by observing the environment through the sensors and deriving the context. In "F1 Skates" Smart-Its were attached to inline skates and a kickboard; context had to be retrieved and communicated to a second Smart-It attached to the body of the user for output status (Figure 2). In a third experiment a protocol enhancement for energy aware communication of context had to be implemented. The fourth and last experiment, spaces of trust, was about using context collected in a room to generate areas of trust and keys for exchanging documents among those present.

For many of the technologies in Ubicomp students had no previous knowledge. Although we thought this would be a major problem, we found out it was not. E.g. the programming language for Smart-Its is plain C. None of the students were familiar with C, yet all were able to run programs successfully after 2 weeks. The provided example code and code-parts were enough to have a quick start into the environment. Also, although there was no previous knowledge of dealing with hardware, the students managed to acquire this skill, e.g. built cables for connecting sensors and batteries to Smart-Its, soldered additional piezos for acoustic output etc. Instead, the major problems appeared when implementing Smart-Its in the physical world. Cable deterioration and inappropriate housings without damping were found as the major sources of error during the run of the course.

Figure 2. Practical Course Experiment

Before the next course held in 2003 we will especially improve these parts of the system. We will also look into using Smart-Its as a platform for other Ubicomp topics like network routing and HCI studies.

ACKNOWLEDGMENTS

The Smart-Its project is funded by the Commission of the European Union as part of the research initiative "The Disappearing Computer" (contract IST-2000-25428). General information about the Smart-Its platform is available at http://smart-its.teco.edu.

REFERENCES

[1] 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, Ubicomp, pp. 116-122, 2001

[2] J. Hill, R. Szewczyk, A. Woo, S. Hollar, D.E.. Culler, K. S. J. Pister. System Architecture Directions for Networked Sensors in Architectural Support for Programming Languages and Operating Systems, pp 93-104, 2000

[3] Ubitools workshop <http://choices.cs.uiuc.edu/UbiTools01/cfp.html>

The Personality Analyzer

Carola Gustavsson, Alireza Haghighi, Christina Wisser
HCI/Interaction Design program, IT-university of Göteborg
Box 8718, SE-40275 Göteborg, Sweden
carola@hemsjo.com, it1haal@ituniv.se, christina@meditik.se

ABSTRACT
In this paper, we describe the Personality Analyzer. It is an artifact that analyses the users personality. The motivation for the project was to enter deeply in the aesthetics of interaction design with a physical artifact as an outcome. Focus for the project has been on aesthetics and expressions rather than functionality. The primary target group was girls from 8 years and up. During development we used, brainstorming, mind maps, sketches and interviews and we also made mock-ups for initial testing. The result was an artifact that analyses the users personality. The most unique aspect of the artifact is the input device. A user describes her personality to the system by selecting clothes used to dress a paper doll. All technology is hidden from the user.

Keywords
Ubiquitous computing, physical input devices, RFID-tags, aesthetics, tangible interface

INTRODUCTION
In the Interaction Design master's program at the IT university of Göteborg, we had an assignment focusing on expressions and aesthetics in information technology rather than on functionality [1]. One of the results of this course was an exhibition at Universeum, Sweden's national science discovery center [4]. We decided to create something physical, that a lot of people could interact with. The target group for Universeum is teenagers and particularly girls, and that became the target group for this project as well. Different kinds of tests are popular amongst the target group, and that was a major reason for concept we chose. Another goal was to make something that appealed to the visitors and arouse their curiosity. One of the main goals from the aesthetics point of view was to make the technology transparent or hidden for the user. All this taken together resulted in a personality analyzer idea.

Idea generation
To develop our ideas, we used brainstorming, mind maps and sketches in our early works as well as exploring aesthetics on different websites. A small cardboard model of the personality analyzer was made and tried out with test users to evaluate the concept. Interviews were carried out to collect data about personalities and trends. We also performed a survey on various ways to get input and output to the artifact.

Figure 1. The paper doll wardrobe with clothes and the dressing mirror.

THE PERSONALITY ANALYSER
The project was inspired by other studies in the area of interaction with physical objects [2]. We were also inspired by how other researchers had used RFID tags to augment everyday objects (e.g. [5]). The main inspiration however came from paper dolls, popular among young girls. These dolls can be dressed with various paper 'clothes', by placing them on top of the paper doll and gently folding the clothes around the doll (cf. [3]).

The Personality Analyzer consists of three interconnected parts – the wardrobe with clothes, the mirror with the paper doll and the drawers below the dressing mirror, containing pieces of paper describing a certain personality type.

A large female doll figure was made of plywood, ca 40 cm high. Suitable clothes were made of textiles glued onto plywood. RFID-tags were hidden inside the clothes, and an RFID reader with three multiplexed antenna coils were hidden behind the large mirror (Fig. 1, right). The doll was attached directly onto the mirror. The clothes were attached to the wardrobe wall using Velcro, and similar Velcro pieces were placed on the doll, to allow for the clothes to be fastened to the doll. Below the mirror are four drawers containing pieces of paper. The drawers cannot be opened manually since they have no knobs. The artifact was designed to attract the visitors. This was done using bright colors and a wave-like shape.

Figure 2. The rear side of a clothes piece showing the RFID tag and Velcro.

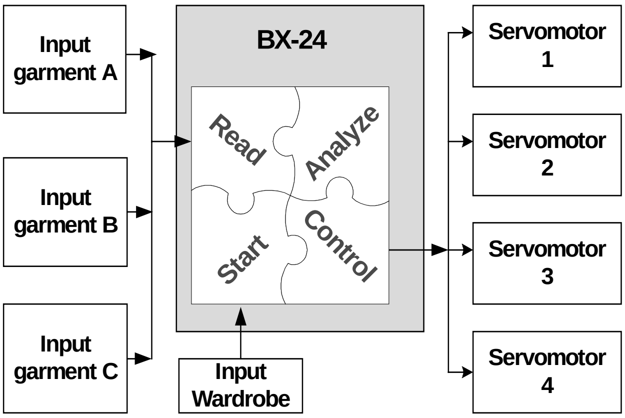

Figure 3. Program and system structure overview.

When using the artifact, you start with opening the wardrobe doors, in order to pick a set of clothes that appeals to your personality (Fig. 1, left). Next, you dress the paper doll on the mirror with the chosen clothes, thus implicitly expressing your personality type. After dressing the doll, you close the wardrobe doors and the analysis process starts. The result depends on the mix of clothes that was selected. After a short delay, one of the drawers is automatically opened, showing the analysis result on a sheet of paper inside the box. The personality is represented as an animal (cat, dog, horse, polar bear) accompanied by a short descriptive text.

DETAILED SYSTEM OVERVIEW

The Wardrobe

The wardrobe, 120 cm high, has two doors. Each door has an electronic magnet mounted at the bottom. It registers if the wardrobe doors are closed or open (Fig. 1, left).

The clothes

The wardrobe contains the clothes, held to the wall using Velcro. The mix of clothes is designed to fit different personalities. On the rear side of each clothes item an RFID-tag is placed (Fig. 1, left).

The dressing mirror

Next to the wardrobe is the oval dressing mirror, with the doll placed in the middle. Velcro pieces are mounted on the doll so the clothes can stick on it (Fig. 1, right). Behind the doll and the mirror, three antenna coils are placed, receiving the signal from the RFID-tags. The ID numbers, unique for each clothes piece, are used as input to the personality analysis algorithm.

Drawers

Below the dressing mirror there are four drawers with rails on the sides. Behind each drawer a small servomotor is placed. The motor has a drop shaped arm that can open the drawer. Each drawer contains a personality analysis result on paper. It describes a personality type by comparing you to an animal and its characteristics.

Computer system

A Basic-X 24 microcontroller, containing an Atmel AT90S8535 RISC core, controls the entire setup, including door detectors on the wardrobe, the multiplexed RFID reader circuitry and the four servo motors behind the drawers. The personality analysis subprogram is executed when the wardrobe is closed. The program reads the ID numbers from the tags and an algorithm translates this into one of four personality types.

THE UNIVERSEUM EXHIBITION

During an early stage a small, non-functional cardboard model was made and tested at the university. The final artifact was exhibited at Universeum between May 15-22, 2002. More than 800 people used the system and in general they understood the interaction with the system well and enjoyed using it. The procedure including opening the wardrobe, dressing the paper doll, closing the wardrobe and waiting for the result appeared to be a natural way to interact with the artifact. The way of passing input to the system by dressing the paper doll was highly appreciated by the visitors.

DISCUSSION AND FUTURE WORK

Further development of the technical part of the artifact is needed. The servomotors behind the drawers used at the moment are too weak, and one of them broke down during the exhibition. During the tests done at the exhibition we realized that many men found dressing a girl a bit awkward, and it would have been easier to have a male doll as well. The shoes were too small to be able to have RFID-tags on in the current setup. That was a pity because we think shoes are important in a persons personality. We did not design all clothes ourselves, nor did we do deep analyses of clothing styles. That was something we would like to put more time into. We also wanted to have a light effect on the drawers to show and emphasize that the analysis is done.

ACKNOWLEDGEMENTS

This project was sponsored by the PLAY Studio, Interactive Institute and the IT-University of Göteborg. Thanks to Peter Ljungstrand for all help.

REFERENCES

1. Hallnäs, L. and Redström, J. (2002). From Use to Presence; On the Expressions and Aesthetics of Everyday Computational Things. In *ACM Transactions on Computer-Human Interaction (ToCHI)*, Vol. 9, No. 2, June 2002, pp. 106-124.
2. Ishii, H and Ullmer, B. (1997). Tangible bits: Towards Seamless Interfaces between People, Bits and Atoms. In *Proceedings of CHI'97*. ACM Press.
3. Paper Doll. http://www.paperdoll.dk/MAYA/MAYAc.htm
4. Universeum. http://www.universeum.se/
5. Want, R., Fishkin, K., Gujar, A. and Harrison, B. (1999). Bridging Physical and Virtual Worlds with Electronic Tags. In *Proceedings of CHI 99*, pp. 370–377. ACM Press.

The Sentient Car: Context-Aware Automotive Telematics

Pablo Vidales and Frank Stajano

Laboratory for Communications Engineering
University of Cambridge
{pav25, fms27}@cam.ac.uk

ABSTRACT

We consider the implementation of a vehicle-based sentient space, which enables context-aware on-board systems with the ability to cooperate and adapt to changing conditions and events. This paper also describes the Sentient Car as a testbed, and the Adaptive Pollution Map as the first application that uses this platform.

Keywords: *Telematics, vehicle, context, mobility.*

SCENARIO

In-vehicle information systems of non-trivial complexity are already commonplace for electronic engine management, entertainment, voice communications and safety systems. GPS-based mapping and guidance systems have made their commercial debut and are already widely deployed in other markets such as Japan.

These subsystems will increase their usefulness once they cooperate with each other to form an integrated system, and once they adapt their behaviour to suit the situation. For example, the cellular telephone that is already present in many cars might be used by the mapping and safety subsystems as a long-range data communications channel for sending out traffic and pollution information about the area currently visited by the car, and for receiving such information for the area that the car plans to visit next. The route finder might then choose to avoid an otherwise convenient-looking route in favour of another that uses a less congested road. Similarly, monitoring driver behaviour (steering, braking, accelerating) and atmospheric conditions may allow the safety subsystem to give early warning of dangerous driving situations. Aircraft-style cockpit recorders ("black boxes"), the high-tech evolution of today's truck tachographs, are another plausible development, and their non-trivial implications on personal privacy are well worth investigating.

We are exploring the above scenarios in the field. We have equipped a vehicle with several sensing, computing and communicating subsystems, all interconnected to form an integrated sentient car system, and we are building a framework to allow each of these subsystems to act as a supplier or a consumer of contextual information about the car, the driver and the environment. In future experiments, a central server will also collate and redistribute the information from several cars. We are also interested in seamless context-driven handover between communication channels of different capacity, coverage and cost, such as GSM and 802.11b.

Figure 1 Sentient car dashboard

SYSTEM RESOURCES

We have equipped our experimental vehicle, a Ford minivan, with two dashboard LCD displays (figure 1), one for the driver and a larger one for the navigator, and with the following subsystems.

Sensing. We measure various air pollutants with sophisticated sensing equipment (figure 2). We also have a tap into the Electronic Control Unit of the vehicle to extract velocity, acceleration, temperature, steering wheel position etc. A GPS receiver gives the geographical position of the car to 10 m accuracy.

Figure 2 Pollution Sensors

On the software front we also sense the availability of the computing and communication resources: network connectivity, communication cost, communication bandwidth etc.

Computing. A full-size tower PC in the back of the van has interfaces to all the sensing and communicating devices, as well as ample storage and computing capacity.

Communication. A GSM cellphone and modem respectively provide voice and low-speed data connection with almost universal coverage. An 802.11b network card provides a much higher speed data link but with limited coverage.

The conceptual model of the sentient car system sees each of the sensors as an input to a "context server" that delivers context information to the on-board systems and to the communication resources manager. Each activity can then be optimally adapted to the current situation.

SAMPLE APPLICATION: ADAPTIVE POLLUTION MAP

One of the first applications we have developed on this platform senses the air quality around the car and overlays it on a map of the area that the car is visiting. In due course, multiple cars will upload such pollution data to a server that will then redistribute more complete and more up-to-date maps to all of them.

The client-server application is split into a client-side map viewer, running in the car, and a server-side map repository running at the fixed base station. We wish to explore the architectural issues that arise when maps have to be delivered over a channel of varying capacity and therefore we pretend that the whole map won't fit on the in-car computer's hard disk. This is not true for our current experimental setup, but it might be once we move to richer media including for example real time video footage of highway traffic (figure 3).

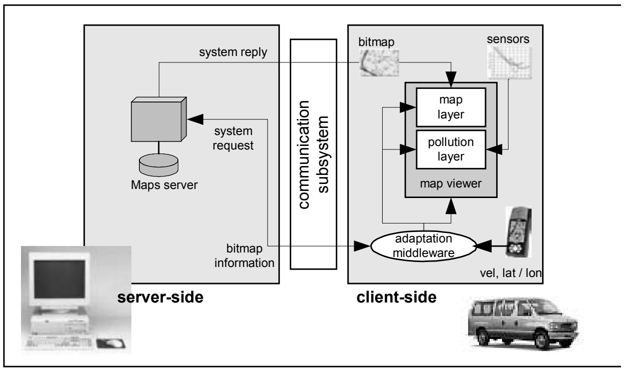

Figure 3 Application overview

The map viewer is arranged into layers. We currently have two: a bitmap layer showing the actual map and a pollution layer showing our air quality measurements. Many other layers could be added: for example a waypoints layer, a vector map layer, a traffic layer and so on (figure 4).

The portion of map that the viewer displays is affected by user input (scroll, zoom in, zoom out) and by the current context, which right now consists primarily of the car's location and velocity as reported by the on-board GPS. However the channel capacity (which goes up when the car is in an area of 802.11 coverage as opposed to just GSM) is another interesting input. As the car moves, we request a new bitmap to cover the relevant area, but we do so taking into account the time it will take to download it over the available channel.

For example, using 8 cm x 10 cm map tiles of 176 kb each, the time it takes to download one tile over the 9.6 kb/s GSM modem is 18.3 seconds. If we are travelling at 100 km/h heading NE, which tile should we request from the server, and at what scale, so that we are still in the middle of it when we have finished downloading and displaying it, 18.3 seconds later? This is a context-dependent calculation which is affected by position, velocity and channel capacity.

We are interested in the architectural issues of managing such updates in real time and with a pool of several cooperating cars all uploading and downloading information into a collective map database.

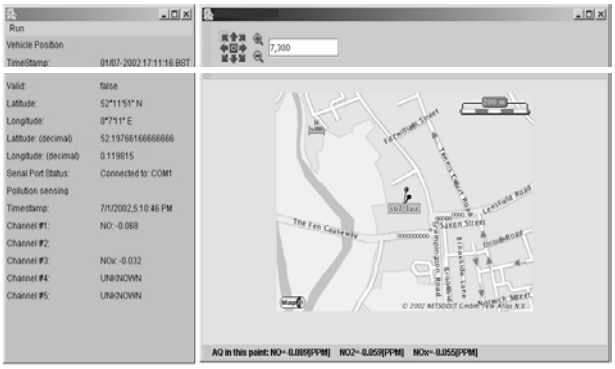

Figure 4 Adaptive Pollution Map

CONCLUSIONS

Sentient spaces are increasingly popular in building environments such as the office and home. The use of sentient spaces in cars will enhance the capacity of on-board systems dramatically, and will allow them to cooperate to provide and consume context information and adapt to frequently changing conditions. On-board systems must deal with more hostile heterogeneous conditions than apparent in sentient buildings, therefore knowing the context is even more important.

The Adaptive Pollution Map application shows the utility of sensing space around the vehicle in order to deliver service with different resources available.

ACKNOWLEDGEMENTS

Pablo Vidales has a scholarship from the Mexican government through the National Council of Science and Technology (CONACYT). The Sentient Car is funded by the Cambridge-MIT Institute (CMI).

Thin Silicon in a Novel 3-D Format for Implementation in Distributed Autonomous Micro Modules

Teresa Cussen, Bivragh Majeed, Kieran Delaney and Cian O'Mathuna
National Microelectronics Research Centre
Lee Maltings, Prospect Row, Cork, Ireland
353 21 4904088
{tcussen, bmajeed, kdelaney, cian.omathuna} @nmrc.ie

ABSTRACT
Silicon thinning, interconnection and packaging are key innovative hardware technologies that can be used to realise distributed autonomous micro-modules (DAMM) for future ad-hoc networks in ambient systems and intelligent environments. These would interact, respond and learn from their surroundings making integration of engineering, computer science and human intelligence a reality**.** This paper investigates thinning silicon sensors and packaging these sensors in a tetrahedral format. This form was chosen because it is vastly expandable and when miniaturised can be used as a building block for DAMM's, which can be designed to physically integrate into materials from which artefacts are fabricated.

Keywords
Thin silicon, 3-D packaging, flip chip, micro-modules.

INTRODUCTION

Ubiquitous computing represents the continuous trend in computing and communication technology towards smaller and more integrated information technology devices. The demand for smaller size, higher performance and increased functionality, means that interconnection and packaging technology has to improve significantly. The success of this integrated technology is dependent on the ability to embed DAMM's into our everyday surroundings. The optimal approach is to integrate DAMM's into an artifact in a manner coherent with its manufacture rather than adding it on after manufacture. These DAMM's will have computational power along with sensing and actuating abilities. The interface between the "real world" and the micro-nodes is established through sensor data. Interconnection and packaging technology has a key role to play in making these systems a reality. Advances in the packaging of sensors, in particular, strengthens this "real world" interface and facilitates the realisation of these DAMM's.

The most recent development in packaging to meet the growing demand for small highly integrated system has been chip scale packaging. These miniaturised packages combined the benefits of flip chip technology, i.e. shorter interconnection and area connection, with the considerable benefits of the package itself i.e. physical protection and easier assembly. But these technologies have lagged behind the growing demand for smaller and higher functionality devices. 3-D packaging has now emerged, as an innovative way to meet market requirements for the next generation of electronic products [5].

By using stacked thin chip packages it is possible to increase system integration while also reducing the length of the wire connects between the die thus reducing noise, increasing speed and reducing power. Many of the benefits of thin silicon packages contribute to the viability of distributed autonomous systems, such as those needed in ubiquitous computing, as these thin packages allow for easier integration of intelligent modules into artefacts.

Since standard silicon die thickness is in the order of 500 microns, reducing the height of a stacked package can be achieved by thinning the individual layers of the package [3]. Also, when silicon is thinned to below 50 microns it becomes flexible [1], [4], combined with a flexible substrate it can be a major advantage in the areas of wearable computing and embedded artefacts. One of the biggest advantages of 3-D packaging technology is that it increases silicon efficiency, which is the ratio of active silicon area to footprint area, an efficiency of greater than 100% is possible.

One of the earliest stacked chip scale packages consisted of bare die Flash memory and SRAM for use in mobile phones [3]. Today however industry has advanced to the point where it is considered feasible to stack 48 layers containing a total of 52 chips. The finished product would contain a processor, interface chips, DRAM, and 32 layers of flash memory [2].

There are three methods of 3-D stacking, which can be classified as follows: 1) wafer level stacking, which involves stacking thinned wafers on top of each other and interconnection is made through wafers via-holes, 2) chip level stacking, which involves stacking chips on top of each other and 3) package level stacking where packaged IC's are stacked to increase the silicon efficiency.

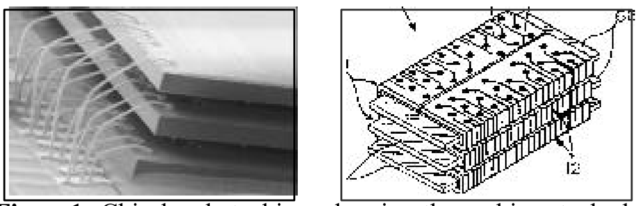

Figure1: Chip level stacking, showing three chips stacked and interconnected via wire bonding

This project deals with innovative chip level packaging. It employs thin silicon micro-sensors, which are packaged into a new shape using a flexible substrate. The aim is to produce the most efficient surface to volume ratio package, which can then be integrated with other similar packages to build DAMM's. The concept of packaging micro-sensors into tetrahedral shapes is very novel and has not been investigated before. This work has been done as a feasibility study for DAMM's at millimetre scale before moving to a highly miniaturised level.

SILICON THINNING

There are three main techniques used to thin silicon: wet etching, plasma etching and back grinding. For this experiment back grinding was chosen, as it is the most efficient method of removing the bulk of unwanted silicon. Back grinding refers to a mechanical process of removing silicon using an aluminium oxide powder (Al_2O_3). Al_2O_3 powder is mixed with water to make a slurry, which is released onto a grinding plate. Pressure is applied to the sample as it rotates on the surface of the grinding plate and the Al_2O_3 and H_20 slurry removes the excess silicon.

The test chip used consists of a heating element and diodes. The diodes act as temperature sensors where their forward voltage drop varies with changes in temperature. These changes in temperature are produced by applying a voltage to the on-chip heating element [6].

The electrical properties of a 225micron and an 80micron chip were compared to those of a regular 525micron chip to determine the effect of the thinning process. The result of this experiment is shown in figure2.

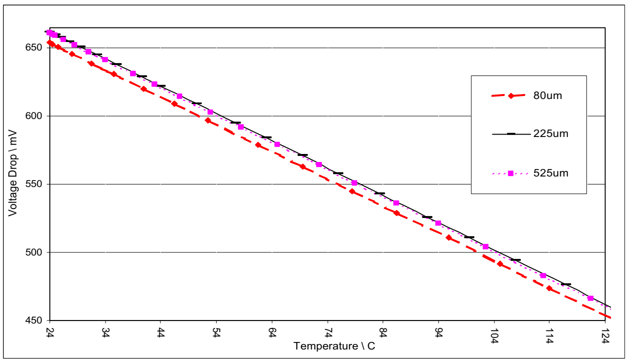

Figure 2: Forward voltage drop across diodes of varying substrate thickness as a function of temperature.

As the silicon substrate is reduced, figure3, so too is the chips ability to dissipate heat. This means that for the same heater voltage the thinner chips are much hotter than the regular chips. However, as shown in figure2, the voltage drop across each of the diodes is the same for the same temperature. This means that the thinning process does not adversely affect the operation of these temperature sensors.

Figure 3: Comparative profile of thin chip to regular chip.

PACKAGING

The packaging step starts with thinned silicon bonded onto a flexible substrate (flex). After etching of single sided flex the thin die are attached to the substrate using flip chip technology. Through holes, made by mechanical drilling, are necessary to access both sides of the flex. These holes are filled either with isotropic conductive paste or solder to obtain the connections. After testing, **t**he packaging step is completed by forming the flexible substrate into a shape of a tetrahedron.

The tetrahedron is the simplest of the polyhedrons with four equilateral triangles for faces. This is the least number of faces required to enclose a portion of three-dimensional space. The tetrahedron is the most stable structure when force is applied [7]. At the current size, it is possible to assemble tetrahedrons at a prototype level, and to validate their robustness and reliability. Repeatability at high volume is in question as nature of assembly technique will not scale easily, however, it is speculated that at the target level of miniaturisation, techniques such as self-assembly could solve the problem. Table 1 shows some of the important properties of a tetrahedron in comparison with other regular polyhedrons and some shapes into which tetrahedrons can be formed. Four tetrahedrons form another tetrahedron with a central space occupied by an octahedron, 5 tetrahedrons make a flying saucer structure and 20 tetrahedrons form a complex near spherical structure. These shapes are investigated for possible physical integration into an artefact. The most important property is surface to volume ratio, which this project plans to exploit in 3-dimensional packaging.

Properties	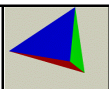 Tetra	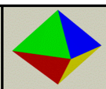 Octa	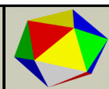 Icosa	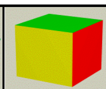 Cube	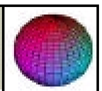 Sphere
Surface Area	1.73 a^2	3.46a^2	8.660 a^2	6 a^2	12.5 a^2
Volume (vol)	0.117a^3	0.471a^3	2.357 a^3	a^3	4.2 a^3
SA/Vol ratio	14.8/a	7.35/a	3.97/a	6/a	3/a

Shapes into which tetrahedron can be formed			
Properties	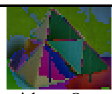 Tetra with an Octa centre	 Flying saucer	 Icosahedron
Surface Area	1.73 a^2	4.330 a^2	8.660 a^2
Volume (vol)	0.117a^3	0.5890 a^3	2.357 a^3
SA/Vol ratio	14.8/a	7.35/a	3.97/a

Table 1:Compartive properties of different polyhedrons

CONCLUSION AND FUTURE WORK

In order to develop and integrate DAMM's into every day surroundings a novel packaging technique is being explored in this current work. Future investigation will be directed at thinning and reducing the chip dimensions of more complex sensors including accelerometers. A comparative investigation will be carried out into heat transfer, connectivity reliability and stress analysis on different shapes and the effect on these properties, when thinned chips are packaged. The ultimate goal will be to miniaturize the dimensions of the DAMM to micron level so that they blend into typical materials used for manufacturing artifacts.

REFERENCES

[1] Chen, K. Y. Ultra-Thin Electronic Device Package *IEEE Trans. on Adv Pack*, Vol. 23, No. 1, Feb. 2000.

[2] Gann, K. D. Neo-Stacking Technology *HDI Magazine*, Dec 1999.

[3] Goldstein, H. Packages Go Vertical *IEEE Spectrum*, August 2000, pages 46-51.

[4] Landesberger, C. New Dicing and Thinning Concept Improves Mechanical Reliability of Ultra Thin Silicon *Adv. Pack. Mat. Proc. Prop. & Interf., Proc. IEEE pp92-97.*

[5] Morihiro, K. The Dawn of 3D Packaging as a system-in-package *Sharp Publication.*

[6] O'Mathuna, S.C. Test chips, test systems and thermal test data for multi-chip modules in the ESPRIT-APACHIP project *IEEE Trans. Comp. Pack. Manu. Tech*., A. 1994; Vol. 17: Sept. 1994.

[7] Wenninger, J. M. Polyhedron Models *Cambridge University Press*, 1996.

Toward A Better User Experience in Tele-education – Recent Advance in Smart Classroom Project

Weikai Xie, Yuanchun Shi and Guanyou Xu
Department of Computer Science, Tsinghua University, Beijing, 100084
{xwk@media.cs, shiyc@, xgy-dcs@}.tsinghua.edu.cn

ABSTRACT
Smart Classroom is an augmented classroom in which Tele-education functions can be seamlessly integrated with the face-to-face education environment. Some recent achievements of this project is described here. The first is the development of the Virtual Assistant module, which can enhance the speech interaction between the teacher and the system. Following is the integration of a SmartBoard screen and the development of a laser pointer tracking module, which replaces the former Virtual Mouse. Finally, a module called Smart Cameraman is developed. This module can automatically select the video to send to the remote students, from the cameras in the classroom that is most appropriate to the current context of the class

Keywords
Smart Space, Tele-education, Intelligent Environment

INTRODUCTION

Among the promising application domains of UbiComp is the education. The Classroom 2000 project [1] at GIT has been devoted to the automated capturing of the classroom experience. Likewise, the Smart Classroom project at our institute is focused on Tele-education. Most currently deployed real-time Tele-education systems are desktop-based, in which the teacher's experience is totally different from teaching in a real classroom. For example, he/she should remain stationary in front of a desktop computer instead of moving freely in the classroom. The goal of Smart Classroom project is to narrow the gap between the teacher's experience in Tele-education and that in the traditional classroom education, by means of integrating these two currently separated education environments together. Our approach was to move the user interface of a real-time Tele-education system from the desktop into the 3D space of an augmented classroom (called Smart Classroom) so that in this classroom the teacher could interact with the remote students with multiple natural modalities just like interacting with the local students.

An overview of the project and the description of the first phase work can be found in [2]. To summarize, the following features have been implemented in the first phase of the Smart Classroom: 1) A Biometric characteristic based authentication mechanism. The teacher is authorized to use the classroom if and only if he/she passes the face and speaker identification process. 2) The Mediaboard facility, whose interface is a wall-size project screen. The teacher can display slides on it to the students. Remote students can view the content of this board with the corresponding client software. 3) The Studentboard facility, which is another wall-size project screen functioning as the window to the remote students. The teacher and local students become aware of the presence of remote students through this facility. 4) The Virtual Mouse facility. Through a hand-tracking module, the system allows the teacher using his finger to drive the cursor on the Mediaboard. 5) The speech interface. The teacher can use voice commands to control some functions of the room, and the room will inform the teacher the occurrence of some events with synthesized voice. Fig. 1 is a snapshot of a prototype system deployed in our lab. Further information and a demo video can be found on our website [3]. In this paper, we present some recent achievements of this project.

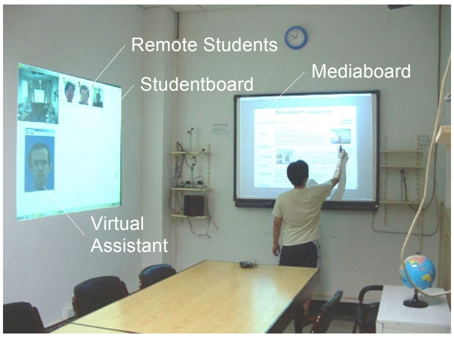

Fig. 1 A snapshot of the Smart Classroom

VIRTUAL ASSISTANT

As mentioned above, the teacher in the Smart Classroom could use speech to interact with the system. However, after some evaluation, we found it seems somewhat strange for teachers to speak and listen to a lifeless room, so we introduced a Virtual Assistant role into the room to embody the speech interaction with the teacher on behalf of the room. The Virtual Assistant encompasses a face-animation engine as well as the former stand-alone Speech Recognition module and Text-to-Speech module. It is represented as an animated human head displayed on the Studentboard (as showed in Fig. 1) with its face and lip animation synchronized with the synthesized voice. Through this facility, the metaphor of speech interaction for teachers was turned from speaking with a room to speaking with an assistant.

FROM VIRTUAL MOUSE TO SMARTBOARD AND MAGIC LASERPOINTER

At the beginning of the project, we intended to develop a mechanism that functions as a mouse but without the need of approaching to a desktop computer. Therefore we developed the Virtual Mouse module, which is based on a hand-tracking algorithm. The teacher's hand movement along the vertical plane was tracked when he standing steps away from the Mediaboard, and could be turned into the movement of the cursor on the board or strokes on the slides. However, the user experience turned out to be frustrating for two reasons: 1) The skin-color modal and appearance modal of hand used to locate the teacher's hand are intrinsically un-robust in a setting like Smart Classroom where illumination is variable due to the use of front-projection projector, The case is even worse when the light of the projector irradiates on the teacher's hand directly. 2) While the metaphor here, using teacher's finger as an extended pointer device, is natural for driving the cursor, but not for draw scribbles on the slides, for one is difficult to keep his/her hand movement smooth without a firm support point of his/her arm.

According to this lesson, we gave up the idea of Virtual Mouse and developed other two mechanisms instead, which together can achieve the same goal but more reliable. First, to allow the teacher to stably draw scribbles on the slides, we adopted a commercial touch-sensitive projector screen called SmartBoard as the screen for Mediaboard, which can digitalize the strokes drawn on that. This way the teacher can directly write or scribbles on the board with provided stylus pens and the strokes will be overlapped on the slides. The Second mechanism is as follows

Magic LaserPointer

It is a laser pointer tracking module through which the teacher can interact with the Mediaboard or Studentboard at a distance. It is reasonable to develop an interactive tool based on a laser pointer because it has been used widely in presentations and the spot of the laser pointer is relatively easy to track. We use a combination of following clues to locate the spot: 1) its intensity in red channel is high 2) it has a small appearance 3)it has a continuous movement trajectory. Fortunately, The developed algorithm has shown its reliability in practice. The located spot of laser pointer is interpreted differently, depending on whether it is on the Mediaboard or Studentboard. If on the Mediaboard, it is interpreted as the mouse movement events on the local Mediaboard, and at the same time, a red spot will be displayed on the remote students' screen, to indicate the position the teacher is currently pointing to. While if on the Studentboard, it is used to select a remote student when the spot moves onto the icon representing this remote student. If successfully selected, the icon will be highlighted as an indication (see Fig 2). Together with the speech interface of the system, the teacher can carry out a rich set of tasks with only a laser pointer. For example, the teacher can point at a hyperlink on the slides on the Mediaboard and says "Jump to this page". Then the Mediaboard will load this page. Or the teacher can point at a remote student's icon on the Studentboard and says, "Please give your answer". Then the system will switch the floor to this remote student.

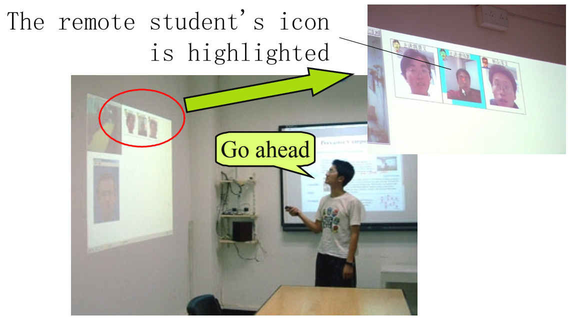

Fig. 2 Select a remote student with a laser pointer

SMART CAMERAMAN

In our real world, when attending in a classroom, a student can adjust the focus of his/her sight from time to time as the context of the class changes. For example, when the teacher is writing a formula on the blackboard, the student will focus on the formula. Or while the teacher is showing a model of DNA, the student will focus on the model on the teacher's hand. To provide the remote students with the same experience, we developed a facility called Smart Cameraman, which could infer the current focus in a classroom based on perceived information, then from the video output of the cameras in the classroom, automatically selects the most appropriate one in the current scenario, and transmits it to the remote students. The context in the classroom is modeled as an enhanced Finite Sate Machine, each state representing a possible context of the classroom, including TeacherWritingOnBoard, TeacherShowingAModel, RemoteStudentSpeaking and Others (Their names reflect their intended meanings). The state transition is trigged by events perceived in the system, typical events include : 1) the SmartBoard detects a contact on it. 2) the teacher picks up a model from a table or put it back, which is recognized by a gesture-recognition module. 3) the floor is given to a remote student. In order to prevent the state from changing o frequently (it may results from either the temporary disturb in the system or the indeterminate nature of some events), we associate each state with an inertia value that decreased as time passing and associate each event a probability value. Whether the state transition will take place depends on a weighted consideration of the remnant inertia of the current state and the probability of the occurred event.

REFERENCES

1. Gregory D. Abowd. Classroom 2000: An Experiment with the Instrumentation of a Living Educational Environment. *IBM Systems Journal, Special issue on Pervasive Computing*, Volume 38, Number 4, 508-530,

2. Weikai Xie, Yuanchun Shi and Guanyou Xu. Smart Classroom - an Intelligent Environment for Tele-education. *In Proceedings of The Second Pacific-Rim Conference on Multimedia (PCM 2001)*, Beijing, China. Springer LNCS2195, 662-668.

3. http://media.cs.tsinghua.edu.cn/~pervasive

Transmission Frequency Optimization for Ultra-low Power Short Range Wireless Communications

D C Yates, A S Holmes and A J Burdett
Department of Electrical and Electronic Engineering
Imperial College of Science, Technology and Medicine
Exhibition Road, London SW7 2BT
david.yates@ic.ac.uk, a.holmes@ic.ac.uk, a.burdett@ic.ac.uk

ABSTRACT

Analysis is introduced which determines the optimal transmission frequency for maximum power transfer across a short-range wireless link. This essentially consists of a comparison between two transmission methods known as near-field and far-field transmission. Constraints on antenna dimensions and the required transmission distance strongly influence the choice of frequency. Preliminary results for coil antennas of varying dimensions have been presented in the form of a surface plot. This illustrates the regions of superior power transfer for the two transmission methods depending on application parameters, thus enabling an optimal frequency to be chosen.

Keywords

Wireless, ultra-low power, near-field, far-field, antennas

INTRODUCTION

The ubiquitous computing paradigm envisages dense wireless networks connecting nodes of varying computational power. Central to this vision is the apparent invisibility of many of these devices to the user, who is aware of, but is not inconvenienced by their presence. This requirement places stringent limits on the acceptable size and weight of devices and encourages the embedding of network nodes into everyday objects. Even more challenging is that apparent invisibility demands minimal power consumption. A short battery lifetime for the nodes of a dense wireless network is completely impractical for the user, yet the battery capacity is limited by the size restrictions. Ultimately these devices should be self-powered, perhaps using either solar cells or vibration-to-electrical energy converters. This becomes conceivable when nodes begin to consume less than 100 microwatts [4].

For those nodes that act as sensors, have limited computational power, or whose main function is simply to relay the incoming data to the following node it is the RF communications that will likely dominate the power consumption. Equally the main limitation to size reduction will be the required antenna dimensions. Electrically small antennas (that is antennas, whose dimensions are much smaller than the wavelength) have exceptionally poor efficiency, inferring that small antenna dimensions require a high transmission frequency. It must be realized however that the power dissipated by the electronics increases with frequency [6]. The optimization of the above factors to minimize power consumption is thus necessary in the pursuit of an ultra-low power radio link for ubiquitous computing networks.

ANTENNA FIELD REGIONS

Three field regions surround a transmitting antenna [3], two of which need to be considered here. The reactive near-field contains stored energy in either an electric or magnetic field depending on antenna type. Conversely the propagation of energy as electromagnetic waves takes place in the far-field.

The field region in which the receiver lies is essentially determined by the transmission frequency for a particular transmission distance. The relationships governing the power transfer from transmitter to receiver differ substantially depending upon the field region. These must be fully analyzed to determine the application parameters for which low frequency near-field transmission outperforms high frequency far-field transmission in terms of transmitter to receiver power transfer.

MODELLING POWER TRANSFER

Far-field transmission can be modelled using the well-known Friis transmission formula [2] [3]:

$$\frac{P_{RX}}{P_{TX}} = p\eta_{TX}\eta_{RX}D_{TX}D_{RX}\frac{\lambda^2}{(4\pi x)^2} \quad (1)$$

x is the transmission distance, λ the wavelength and p is the relative antenna polarization factor. η and D represent the standard antenna parameters of radiation efficiency and directivity respectively. Subscripts TX and RX distinguish between transmitter and receiver parameters.

A similar power transfer relationship (shown below) has been derived for the near-field case. The validity of this equation requires poor coupling between transmitter and receiver, which is generally the case since the transmission distance is significantly larger than the antenna dimensions for most applications.

$$\frac{P_{RX}}{P_{TX}} = p\frac{N_{TX}^2 N_{RX}^2 \pi^2 r_{TX}^4 r_{RX}^4 \omega^2 \mu_0^2}{16 R_{TX} R_{RX} x^6} \quad (2)$$

for two antenna coils of N_{TX} and N_{RX} turns with radius r_{TX} and r_{RX} respectively. ω is the angular frequency and μ_0 is the permeability of free space.

Antenna Modelling

To evaluate these two expressions the antenna parameters for the two transmission methods have to be modelled. Preliminary analysis has concentrated mainly on loop (coil) antennas due to their benefit in the design of ultra-low power transmitter architectures [7]. Of great importance are the two dissipative elements in the coil - the loss resistance, R_{LOSS}, and the radiation resistance, R_{RAD}, which combine to from R_{TX} and R_{RX} in (2). Antenna efficiency, η, and directivity, D, depend on these parameters as follows:

$$\eta = \frac{R_{RAD}}{R_{RAD} + R_{LOSS}} \quad (3) \qquad D = \frac{E^2 L_e^2}{4R_{RAD}} \quad (4)$$

where E is the incident Electric field strength and L_e is the effective antenna length.

The radiation resistance determines the amount of power transferred to the far field for a particular antenna input current. This has been modelled in MATLAB using an analytic equation derived in [1]. The loss resistance must be modelled by taking into account the proximity and skin effects, which explain how increasing frequency causes the cross section of the conductor, in which the current flows, to reduce. These effects have been modelled by utilizing equations and tables of parameters, derived and computed by S Butterworth and summarized in [5].

COMPARISON

Determining the regions in which a particular transmission method outperforms the other is a complicated process and depends upon the interaction of many parameters. Thus, only a brief overview can be given here.

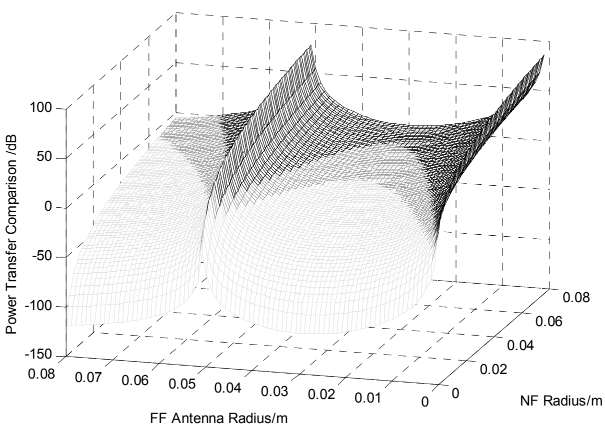

Fig 1: Power transfer comparison between near and far field transmission for varying antenna radius

Figure 1 is a comparison of near-field transmission (at 50MHz) with far- field transmission (470MHz) for a distance of 50cm, varying the near and far-field antenna dimensions (shortened to NF and FF respectively on the axis labels). The black region represents the area for which near-field transmission is superior. The darker gray area denotes the region where far-field transmission operates more efficiently than near-field transmission by a margin of 10dB or less. The use of near-field transmission should be strongly considered in this region since the lower operating frequency results in reduced power dissipation in the electronics [6]. To a first order approximation the power consumed by an analogue circuit can be considered to be directly proportional to frequency [6]. Depending upon the dominance of the electronics in the equations governing transmitter power consumption, the far field case could perform worse than illustrated in figure 1 by up to about 10dB for the frequencies used in this comparison. The lighter gray represents the area where far-field transmission is superior by a margin of 10dB or more.

CONCLUSION

The above surface plot (figure 1) suggests that near-field transmission should be employed once the allowable antenna radius exceeds 0.05m. Such graphs must be treated with caution, because the conclusions can alter sharply depending on other parameters. Distance is of particular importance, since near-field transmission decreases with $1/x^6$ compared to $1/x^2$ in the far-field case. It is also imperative that the optimal frequency within the two transmission methods for particular antenna dimensions is first evaluated in order that the correct comparison is made. This can also be achieved using the modelling introduced here.

ACKNOWLEDGMENTS

This project is funded by the DC-Initiative as part of the ORESTEIA project.

REFERENCES

1. Adachi, S., Kasahara, T., Mushiake, Y. A Loop Antenna in a Compressible Plasma. IEEE Trans. Antennas and Propagation, May 1969, AP-17, 396-398
2. Friis, H. T. A Note on a Simple Transmission Formula. Proc. IRE Waves Electrons, May 1946, vol 34, 254-256
3. Johnson, R. C. Antenna Engineering Handbook, Third Edition, McGraw-Hill 1993
4. Rabaey, J. M. et al. Design Methodology for PicoRadio. Design. Automation and Test in Europe, 2001. Conference and Exhibition, Proc. 314 -323
5. Terman, F. E. Radio Engineers Handbook, McGraw-Hill 1943
6. Vittoz, E. A. Low Power Design: Ways to Approach the Limits. International Solid State Circuits Conference, 1994. Digest of Technical Papers, 14-18
7. Ziaie, B., Najafi, K., Anderson, D. J. A Low-power Miniature Transmitter Using a Low-loss Silicon Platform for Biotelemetry. Proc., 19th Int. Conf., IEEE/EMBS, 1997, vol. 5, 2221 - 2224

Websign II: A Mobile Client for Location based Services

Geoff Lyon, Mehrban Jam, Cyril Brignone, Salil Pradhan
Hewlett Packard Laboratories
1501 Page Mill Road, Palo Alto, CA 94304, USA
{gml, jam, cyrilb, salil} @hpl.hp.com

ABSTRACT
The WebsignII mobile client device has been designed to allow nomadic users to interact with localized services. Implemented as an expansion pack for a personal digital assistant, the client provides a looking glass linking users to web based services at their current location. The client is also able to identify and resolve physical objects within close proximity of the user.

Keywords
Mobile Client, Location Awareness, Service Discovery

INTRODUCTION
The Websign concept contains both systems infrastructure and client access devices, and is designed to unite physical and virtual spaces by binding physical locations and objects to virtual information services [1]. This augments the nomadic user's reality with pointers to websites and services that are specific to the users immediate location and surroundings.

Figure 1: Jornada 568 based Websign Client

Users interact with the Websign client, a personal digital assistant (PDA) equipped with a custom expansion module as shown in figure 1. This expansion module contains global positioning system (GPS) and compass circuitry that determine the location and orientation of the PDA. The client retrieves hyperlinks to services in proximity to the user using a wireless data link; these hyperlinks are typically filtered against the user's profile or immediate interests to present the most appropriate data set for their present context. As the user rotates, hyperlinks to the various services in the directional vicinity are highlighted on the device. By clicking on one of these, a link to the appropriate website or local service is established. This type of interaction is suitable for outdoor use and where the user is still some distance from the points of interest (with respect to the accuracy available with GPS). Use of the system indoors and where objects are very close together is still challenging. To assist in these situations, WebsignII adds barcode and RF-ID reader technology to the expansion module. This allows the user to scan physical objects or reference labels and then resolve these to the appropriate online information or service [2,3].

CLIENT HARDWARE
The client side hardware consists of a PDA and a prototype expansion pack; solutions have been developed for the both the HP Jornada and the HP iPAQ. Figure 2 shows a block diagram of the Jornada based design.

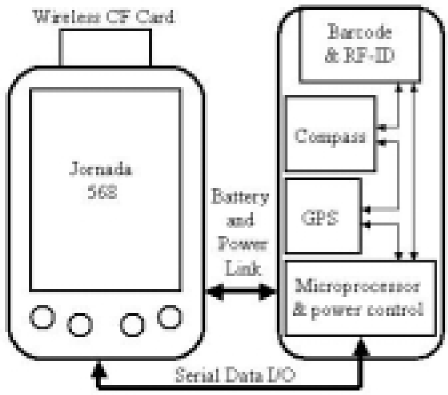

Figure 2: Main components of the client device

Control of the various sensor modules (GPS, Compass, Barcode and RFID) is performed by an embedded microprocessor; an ARM7 based device with integrated memory and peripheral circuitry, resulting in a single chip solution. Communications with the PDA uses its serial port interface. Embedded software in the expansion pack parses command requests from the PDA, performs the required device operations then returns the resulting data. Power management is a key feature of the expansion pack design and is crucial to user acceptance of the client. Independent power control of each sub-system is provided to minimize operational power consumption together with power sensing circuitry to match the expansion packs state to that of the PDA. When the PDA is in power suspend mode or is powered down by the user, the expansion pack goes into a similar state with a current consumption of less than 1uA. Table 1 shows the current consumption under various operational conditions, where Base represents the microprocessor circuitry.

Operating condition	Current consumption
Base	20mA
Base + Compass	32mA
Base + GPS	105mA
Base + Barcode-RFID	200mA
Standby / Switch Off	<1uA

Table 1: expansion pack operational current consumption

CLIENT SOFTWARE

The client's software has been written to simplify application development under the WinCE operating system. All of the functionality has been encapsulated into an ActiveX Control. User applications (developed using Visual Basic, Visual C++ or as a sequence of web pages) all invoke the client functions through a generic ActiveX API. Once invoked, the client software manages local device interactions, local data storage and wireless communications with the Websign infrastructure.

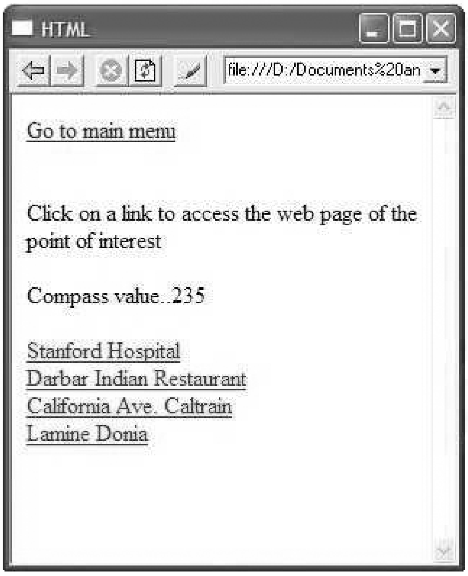

Figure 3: Client user interface example

Figure 3 shows an example user interface, implemented as a web page using HTML. As the device is rotated the compass value alters and the list of associated hyperlinks to places of interest changes based on the orientation of the device. Once the Websign device has obtained a GPS fix, it downloads the appropriate Websigns for that area. To avoid overloading the user with information, two forms of filtering are applied to the pointers. A profile or interest filter reduces the data set to those matching the user's needs. A vector filtering operation reduces the displayed pointers to the device's immediate location and orientation, producing a set of pointers that are close to the user and in the direction that the Websign device is being held.

DISSCUSSION AND FUTURE WORK

The positioning and directivity attributes of the WebsignII client allow it to augment the users present environment and establish associations between physical locations and online information or services. Barcode and RFID reader technologies allow additional associations with physical objects. This combination of features enables the creation of Websign based applications, targeting both ubiquitous computing scenarios and enterprise system solutions.

The client device provides dependable outdoor use and a limited degree of indoor operation. A number of indoor positioning systems based on RF and optionally ultrasound have been suggested [4] and the addition a more ubiquitous solution would enhance Websign's utilization for indoor service discovery. We are also working to develop adaptive horizon filtering algorithms to automate the acceptance range of service links based on the client's context.

REFERENCES

1. Pradhan S. (etal): Websigns - Hyperlinking physical locations to the Web, IEEE Computer vol.34 no.8 (2001) 42-48.
2. Kindberg T.: Implementing physical hyperlinks using ubiquitous identifier resolution. HP Labs Technical Report, 95R1 (2001).
3. Want R. (et al): Bridging Physical and Virtual Worlds with Electronic Tags, Proc. ACM Conf. Human Factors in Computing Systems, ACM Press (1999) 370-377.
4. Hightower J. & Borriello G.: Location systems for ubiquitous computing, IEEE Computer vol.34 no.8 (2001) 57-66.

A Framework for Intelligent Instrumented Environments

Christoph Stahl
Dept. of Computer Science
Saarland University
Im Stadtwald, 66123 Saarbrücken, Germany
stahl@cs.uni-sb.de

ABSTRACT

In this paper, I will present an abstract of my doctoral research activity towards a framework for intelligent instrumented environments that provides abstract layers for sensor input, knowledge representation and human-computer interaction. I will motivate my research by an illustrated example, identify core features of the framework and describe the components of an intelligent environment. Besides the framework itself, my doctoral work is focused on the sensor and mobile assistance components.

Keywords

Pedestrian Navigation, Intelligent Environment

INTRODUCTION

The research activity in the AI-Lab of Prof. Wahlster is focused on the design of intelligent human-computer interfaces and user modeling. Intelligence stems from the dynamic adaptation of graphical presentations according to a user model and the display capabilities of the outputdevice. The graphical presentations are combined with speech output and allow the user to interact with the appli-cation using pointing gestures. Recently our group has developed the hybrid ressource adaptive navigation assis-tance system REAL [1]. The prototype was designed to pro-vide travelers in the rather complex Frankfurt airport with navigational instructions, either by a 3D route visualization on large public displays or by guiding them through incre-mental descriptions using a mobile PDA.

Future research aims to develop a multimodal dialogue manager, allowing the user to ask the system for assistance to solve higher-level tasks, like shopping or arranging appointments. Besides providing avigational assistance, the components have to offer access to location based services knowledge and keep track of the user's goals and schedule. Mobile and stationary components are to be integrated to form an intelligent environment, which is instrumented by various sensors to locate people and objects. In the background, a plan recognition process will constantly monitor and reason about all user activities to proactively remind the user of their goals and to offer assistance.

EXAMPLE SCENARIO: EXPLORING CEBIT FAIR

This section will give an example on how a mobile collaborative pedestrian assistance system, embedded in an intelligent environment, could help you in the near future to successfully explore a complex environment like the CeBit fair in Hannover, Germany.

Suppose that you and some colleagues have been assigned to present your research. This requires one of your group to attend the booth all day. Your secondary, contradictory goals are to gather information, meet business partners and attend to events at certain times and locations. These goals can only be achieved by efficient collaboration.

On arrival, your mobile PDA guides you to your booth using the IR beacons deployed all around the hall. You agree with your colleagues to attend to the booth in the afternoon and for now to enjoy your surroundings. You use your PDA to explore the nearby halls for their exhibitions, and visit some well known companies to see their new products. Instead of collecting paper material to memorize the exhibits, you simply bookmark the location and enter some comments. Meanwhile your colleague occasionally checks your items and copies points of interest to her agenda.

At 10:33 your PDA alarms you of an incoming phonecall from a dotcom CEO who wants to discuss a common project. Since he has no organizer, you ask him to use the public information displays instead. He identifies himself with his ticket, which contains an radio tag, and you see his location. Now you ask the intelligent environment to suggest a restaurant close to both of you, and your personal electronic assistant does the reservation for three persons, since you will need some technical advice from your colleague at the meeting. The assistant puts the appointment goal, including the occasion, time and destination, to all agendas.

Around 11 AM your PDA alerts you that its power is low. You have 15 minutes left, enough to query the environment where to buy AA batteries. The system suggests a nearby shop and puts the shopping goal in your agenda.

On the run you stop to take part in a sweepstake. Unfortunately the prizes are given out at 5 PM. Since you have an agreement with your colleagues to mind the booth in the afternoon, one of your friends will have to act in your place. You enter a localized event with an alarm set to 5 PM into your agenda and direct it to you and your friends.

Now you hurry to buy the batteries, but your PDA alarms you again. It has recognized from the goals outlined in your agenda that you're heading in the wrong direction, and so directs you to the shop. Meanwhile the environment found a shopping goal in your colleagues agenda. "Would you mind buying some chewing-gum for Tom ?" Since you will meet him soon at lunch, the environment found a reasonable plan to ask you to buy it for him. You agree on this and continue to explore the hall.

In the background, your assistant is aware of your appointment. At 11:40 it's convinced that you aren't and reminds you. The visualization of your estimated walking range in the next 10 minutes makes you think about catching a bus, but you can see from the PDA that it's behind schedule. You check the driver's online camera and see a crowd blocking the bus, so you request a shuttle service. At 12 PM you hand over your colleague's chewing-gum, enjoy your meal and agree on the dotcom's project.

REPRESENTATION FRAMEWORK

The previous example introduced manifold services which one might expect from an intelligent environment. In order to realize such a large distributed system, we need to decompose it into components and we have to design and develop a representation framework which allows for the integration of many different components and sensors. The environment's intelligence is based on knowledge about the spatial structure in the surroundings , the location of people and objects within, the user's goals, their possible actions and their plans. Therefore the framework has to provide an ontology and *knowledge representation layer* as a common 'language' among the environment's entities. It also has to provide a *sensor input layer* to feed the knowledge base with observations. In turn, the gained knowledge has to be distributed to various dialog-, planning- and reasoning pro-cesses. To avoid a data avalanche, it seems reasonable to use temporal,spatial and semantic scopes to constrain the information, which has to be communicated.

In addition, the framework has to represent the *human-computer interaction layer*. All available hardware devices in the environment are listed and classified into
- input (speech, touch, vision) and output (audio, video)
- public (shared by many users) and personal

channels. This approach allows dialogue sessions across multiple devices. For example, a PDA may be used as an input channel and a large public display as an output channel. We will adopt the Context Toolkit [2] as the sensor input layer. Since it is typical for a complex and changing environment to have multiple dialogue and reasoning sessions and PDAs entering and leaving, we need further knowledge representation and distribution concepts.

COMPONENTS OF AN INTELLIGENT ENVIRONMENT

Now, that we have explained the framework, we will identi-fy the environment's components and their features.

Dialogue Manager

The dialogue manager is the core component for human-computer interaction. It will handle requests to:

- Provide navigational assistance (How do I get there ?)
- Solve different tasks (I need..!, where can I get. ?)

It is designed to parse multimodal user requests and allows for the combined use of natural language and pointing gestures on a touchscreen. In response, the system will dynami-cally create multimedia output including text, speech and graphics and adapt it to fit the interface's capabilities. On large displays 3D presentations will be shown, whereas on PDA's automatically generated sketches are used [3]. The content is also adapted to the user's language and know-ledge, which is represented in a usermodel.

Stationary Public Displays

Our current information kiosk RANA[1] provides navigational assistance using 3D animations, which adapt to the users time pressure by applying different presentation styles to provide navigational guidance. In order to gain access the knowledge stored in the user's profile, the display should be able to identify the user. To solve this problem various techniques are proposed in the sensors section.

Location-Aware Mobile Assistants

The environment may also be accessed by PDAs. The combination of indoor and outdoor localization capabilities using infrared-transmitters and GPS for incremental navigational guidance have already been demonstrated in [1] and an adaptive route planning mechanism is in development. Additionally, the mobile assistant should allow the user to represent their goals through the use of templates. This has the advantage of a common semantic framework for agenda representation and plan recognition. It is also desired to support location-stamped memos. Sharing this information within the environment allows for many different collaboration features.

Sensors

To act and respond 'intelligently', the environment has to know about the location of persons and objects, the user might be interested in, so we will utilize various technologies. Audio volume gives a cue of a persons presence, and dialogue speaker-clustering allows user classification into stereotypes. Vision systems may allow face recognition or user tracking. Persons will be equiped with radio tags for reliable identification. All these sensors produce a vast amount of information, therefore the analysis and evaluation of the data should be done at sensor locations, using embedded hardware, so that only abstract facts will be sent to the environment's knowledge base.

Plan Recognition

User actions will be monitored by a plan recognition process, which also has access to the environments knowledge base. The user's behaviour is matched against the goals known from their profile, thus the system will be able to proactively help to complete plans by providing selected information.

REFERENCES

1. Baus, Krüger, Wahlster: A ressource-adaptive mobile navigation system. Proccedings of IUI2002: International Conference on Intelligent User Interfaces 2002, ACM Press, 2002.
2. Salber, Dey and Abowd: The Context Toolkit: Aiding the development of context-enabled applications. In Proceedings of CHI'99, pp. 434-441
3. Stahl, Krüger, Baus: Location Dependent Generation of Sketches for Mobile Indoor Route Descriptions. To appear in ECAI 2002 workshop notes on Artificial Intelligence in Mobile Systems (AIMS).

Amplification of Reality: Computation as Material for Design

Jennica Falk
Story Networks group, Media Lab Europe
Sugar House Lane, Bellevue, Dublin 8, Ireland
jennica@mle.ie

ABSTRACT
As our physical environment becomes saturated with computational technology, our perception of and interaction within the tangible domain will inevitably alter. Computation is shifting from being a tool to a material for designing not only digital worlds, but also the physical world, a transformation I refer to the *amplification of reality.* An applied and speculative approach reveals the interaction design implications, and I am currently using an interactive narratives and games domain, live action role-playing, as area for exploring and informing my ideas and prototypes. One of the challenges I am currently facing relates to evaluations of my ideas and implementations.

MOTIVATION
In our interaction with, and within, the world, the properties of our physical environment and objects around us play important roles to enhance and aid communication. Objects and locations are also integral to our perception and understanding of the physical space, and we have developed very sophisticated skills for manipulating and organizing the physical world. These premises form the foundation for much of the research in areas such as tangible user interfaces, context-aware applications, and ubiquitous computing, where the primary objective is to move computation beyond the confining desktop computer, and provide a seamless interfaces between users, everyday objects, and the surrounding environment. Although the mappings between physical and digital media are somewhat trivial to accomplish technologically, at least on a coarse scale, a more difficult task is to form an understanding of the social and cultural implications of such a design practice. How will an increasingly computational physical world, affect our social and cultural environment? What kinds of new interaction paradigms can be designed? How is the physical world designed as an interface that effectively communicates functionality within these paradigms?

Approach
I have taken a speculative and applied approach with the purpose of negotiating these questions. Through the design of functional prototypes, which are informed by existing theories and principles of interaction design, and through contextual studies of human-world interactions, the philosophical and theoretical lens of my research is reduced to practice. I propose three areas of inquiry that will provide a foundation for my doctoral thesis, where the main area focuses on interaction design theory, particularly emphasizing the design of interfaces that span both the digital and the physical worlds. The supporting technical area provides a practical focus, specifically on the design of prototypes. The supporting contextual area provides a parallel focal point on

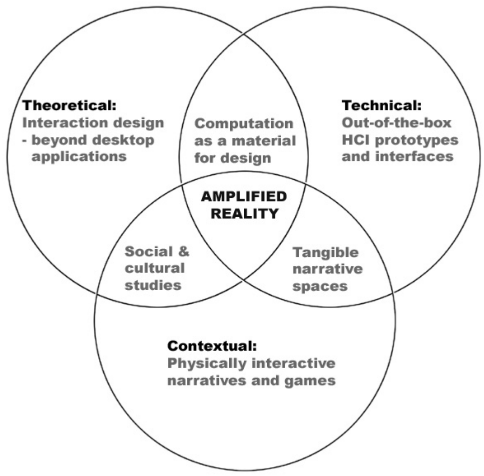

Fig 1. The intersecting research areas

a specific human engagement with the physical world, physically interactive narratives and game play. Figure 1 shows a diagram outlining the intersecting areas. The following sections will examine them in detail.

THE AMPLIFICATION OF REALITY
As a student of *New Informatics* [2], which roughly translates into an applied computer science field, I am in the area of the design and use of information technology. Of particular interest is how computers are changing the way we interact with each other as well as how they change our perception of and relationship with the world. I aim to study and contribute to the transformation of computing from being a tool with which we accomplish tasks, to a material for designing the physical world. I refer to this transformation as the *amplification of reality*, which emphasizes *expressive* and *embedded* computational properties [3]. To tackle this research problem, I have divided the task into three areas of investigation and an additional three areas of application, each described below (this model is an adaptation of the general examinations proposals that are practiced at many universities, for instance in North America).

Theoretical area: Interaction design theories
The theoretical area is primarily a literature review of existing theories of *interaction design.* Of particular relevance is ubiquitous computing with its attempt to move computation into the comfortable attention space of our everyday environment, and tangible user interfaces, which explores how computation can be rendered graspable and embodied in physical forms. A variety of related areas, such as context-aware computing, location-based story-telling, ad-hoc net-

working, wireless computing, etc., are additional topics of interest in this literature review.

Technical area: Prototyping and implementing

Prototypes facilitate communicating and verifying ideas. The supporting technical area has an implementational focus, specifically on the design of interfaces that take computing beyond the desktop application. The technical area offers me opportunity for hands-on experience with enabling technologies such as sensors, actuators, and microcontrollers, as well as opportunity to experiment with the design of form and function.

Contextual area: Interactive narratives and games

I have chosen the specific area of *physically interactive narratives and games* to inspire and inform the prototyping process. What I mean by physically interactive narratives and games are activities in which intimate relationships emerge in the interaction between participants, artefacts, and the variety of contexts that the story or game provides. Of particular interest to me is live action role-play (LARP), which in its basic definition and purpose, is a dramatic narrative form in which participants portray fictional characters to create a web of stories that emerge in their situated and improvised interaction. LARP games suggest a number of interesting relationships from which to extrapolate, between players, plot, locations, and objects in the game. Parallel observations may be made from for instance improvisational theatre or children's play.

APPLICATIONS OF RESEARCH AREAS

The intersections of the three research areas outlined, present opportunities for applied and speculative investigations. In the relationships between theory, practice, and application, a clearer understanding of amplified reality takes form. The following sections suggests the direction of such a discussion.

Computation as a material for design

In the shift from desktop computing to ubiquitous and tangible computing, the computer changes from being a metaphor of interaction, to a medium of interaction. In other words, computation is transformed from being a material for designing not only the digital domain, but also the physical world. Much like wood or steel or paint are used as materials in the design of certain artifacts or effects, computation can be viewed as a material for design. With computation, new expressions and functionality can be superimposed on, and embedded in, the physical world.This is a philosophical as a practical subject, where This notion has been previously explored by Redström in his doctoral thesis [4].

Social and cultural studies

A closer look at objects and the fictitious functionality assigned to them in role-playing games, confirms their importance to the progress of the game as well as to the interaction between players. By studying and participating in LARP games, I am hoping to make a number of observations, e.g.:

- How do physical objects serve as mediators and facilitators of game plot and progress?
- How are different physical locations serving as context for interaction and how do they affect interaction?
- How are abstract and imaginative concepts, such as magic, immortal interventions, death, ghostly presences, etc., communicated within a LARP setting?

Of specific interest to me is to look at how players are interacting *through* rather than *with* this game world and the objects within it, and to reveal design implications for when we in a similar way assign computational functionality to (or amplify) our physical world.

Tangible narrative spaces

Where the contextual and the technical areas meet, specific applications and prototypes emerge. *Tangible narrative spaces* are story or game spaces in which users interact with the physical environment and tangible interfaces in order to interact with narrative content. There environments will employ novel input and output channels, new ways of interacting with and manipulating digital content, and interfaces that are distributed over a large physical space.

I have explored the notion of tangible narrative spaces in some previous projects. *Pirates!*, a collaboration project between the PLAY research group and Nokia Research Center, is a multi-player game for handheld computers, which requires players to explore the physical space in order to explore the virtual game space [1]. The *Tangible MUD* project illustrates how tangible user interfaces can extend a virtual game world, allowing everyday physical objects to function as both input mechanisms as well as displays of game output. I am currently working on a prototype that combines these two projects, creating a physically distributed interactive and tangible story space.

CHALLENGES

Of specific concern to me at this point, is how to evaluate and verify my ideas and implementations. I am not convinced that metrics and quantitative methods are appropriate for evaluating a mainly speculative approach to designing computational technology. To me, and indeed many others who address their work in a similar way, such an approach is very subjective. How do we know when we have made a real contribution? My feeling is that while it is important to frame our work in the context of other's work, what is more valuable is to challenge it by articulating the questions that emerge, rather than solely focusing on the replicable answers, and that we therein may find the contributing factors of this research.

ACKNOWLEDGEMENTS

To my advisors, Professor Bo Dahlbom, Glorianna Davenport, and Hiroshi Ishii, thanks for challenging me and inspiring me. To Johan Redström, and colleagues at PLAY and Media Lab Europe, thanks for all encouragement.

REFERENCES

1. Björk, S., Falk, J., Hansson, R., and Ljungstrand, P. (2001). Pirates!: Using the Physical World as a Game Board. In *Proceedings of Interact'01*, Tokyo, Japan.
2. Dahlbom, B. (1997) The New Informatics, In *Scandinavian Journal of Information Systems*, vol. 8, nr 2.
3. Falk, J., Redström, J., and Björk, S. (1999). Amplified Reality. In *Proceedings of HUC'99*, Karlsruhe, Germany.
4. Redström, J. (2001) *Designing Everyday Computational Things*, Gothenburg studies in Informatics, (Report 20).

Designing Pervasive Computing Technology – In a Nomadic Work Perspective

Jannie Friis Kristensen

Department of Computer Science, University of Aarhus, Aabogade 34, 8200 Aarhus N, Denmark

jannie@daimi.au.dk

ABSTRACT

In my thesis work I am investigating how the design of pervasive/ubiquitous computing technology, relate to the flexible and individual work practice of nomadic workers. Through empirical studies and with an experimental systems development approach, the work is focused on: a) Supporting interpretation and inclusion of implicit and invisible as well as explicit and visible characteristics of artifacts, users and use practices. b) Identifying breakdowns in human-computer interaction situations, with particular emphasis on the computation that happens "behind the scenes" in the pervasive computing environment, and how that computational process at a sufficient level is made intelligible, visible, accountable and negotiable to the human participant.

Keywords

Nomadic work, HCI, pervasive computing.

INTRODUCTION

Studying modern lifestyle and work organization in the light of the vision of creating ubiquitous/pervasive computing environments, has inspired me to focus at specific design challenges posed by the combination of disappearing/invisible computation processes, human computer interaction and specific context issues posed by nomadic and mobile work. The objectives of the visions serve as the general perspective through which I approach technological development and qualitative assessment of supporting nomadic and mobile work.

I begin with an introduction of the research areas of nomadic work and ubiquitous/pervasive computing, second I introduce the empirical studies of my research, and conclude with a presentation of work and results so far.

NOMADIC AND MOBILE WORK

For much of the 20th century, the available communication technologies, construction practices and business processes dictated a world of industrial work mainly confined to scientifically managed office and production buildings. As a result of the general socio-economic and technological development in recent years, we have witnessed radical changes in work-tasks and work organization. These changes are partly characterized by an increasing uncoupling between work and a specific workplace, creating a new generation of so-called 'nomadic workers', characterized by much more flexible, transforming and individual patterns of work, than seen in the early days of the industrial society. Another result of this development is a redefinition of the relationships between work, home and public spaces.

The nomadic and mobile work context has been a central concern for parts of the research communities in HCI and CSCW for the last ten years. Many projects and research groups are focusing on conceptualizing and describing nomadic work through empirical studies and theoretical distinctions, others have focused on concrete design work on services and devices supporting mobile and nomadic work. Examples are Luff and Heath [5], Belotti and Bly [1], Kristoffersen and Ljungberg [4] and Fagrell et al [3]. In parallel we have seen a range of workshops and conferences relating to technical, methodological and social implications of relating to nomadic and mobile work.

UBIQUITOUS AND PERVASIVE COMPUTING

Mark Weiser [7] introduced the notion of ”ubiquitous computing”, describing the vision of the disappearing computer. He was advocating for the computer technology to become part of our existing environments and tools, away from the stationary and personal desktop computer. Also Alan Kay and Don Normann [6] have in different forms followed this line of thought. Alan Kay introduced the notion of transparency, and was an advocate for introducing more natural interaction modes in relation to computers. Also Don Normann has participated in this discourse for decades, and has recently in his book “The Invisible Computer, contributed by introducing the concept of ”Information appliances”. He argues to concentrate on the tools and the tasks at hand, and ”hide the motor” and the computational processes inside.

Nomadic Work and Pervasive Computing

A central concern in supporting mobile and nomadic work is creating an information infrastructure that can support the nomadic/mobile work practice. A reliable communication network through which information is transferable and accessible, is only one component of such system, further components are a network of reliable access points, tools and services. Such a comprehensive network should constitute the link between the physical and social sphere of work activities, and the information sphere in which data is stored, processed, uploaded and downloaded. The visions of ubiquitous and pervasive computing, presents alternative ideas and discussions in this area. In these visions, computing is no longer a discrete activity bound to a desktop, it is an environment where we expect all the “known” computing devices to access information and work together in one seamless, integrated system. The pervasive computing environment is envisioned to help us manage information quickly, efficiently and effortlessly.

EMPIRICAL STUDIES

Empirical studies form the imperative background of my thesis work, based on work in two different large-scale projects: The EU project WorkSPACE [2] and a project for the county of Aarhus, in relation to Center for New Ways of Working. These studies form the imperative background for describing, discussing and understanding challenges in relation to designing pervasive technology environments for nomadic work contexts. I have studied the work of landscape architects, architects and engineers, focusing on the use of objects and tools, special aspects of nomadic and mobile work, environments and rooms, as well as collaboration and distributed work. Material has been collected over a period of one year so far, through interviews, workshops and "fly on the wall" studies, documented by still-pictures, video, transcripts of selected material, workshop material and notes.

THREE FOCUS AREAS

The pervasive computing visions are perfectly focused on supporting mobile and nomadic work practices, but this combination also introduces a number of new challenges and problems. The challenges emphasized in the empirical studies and analysis of the thesis work, has led to a focus in three areas:

- Empirical Translation & Design Context
- Interface and Interaction Design
- Experimental Systems Development

Future work of the project will be focused on 2 and 3.

Empirical Translation & Design Context

The empirical studies highlight two important shifts to address in the design process: 1) From some-place to some-where; Computers have moved from the workplace to home use, onto the streets and into our everyday lives. This also means that the computer – or rather the computing power - has moved from the desktop PC and a fixed context, to a range of old and new devices and environments, working in constantly transforming and largely unpredictable contexts. This shift invokes a second design shift from focus on computers or single devices to a broader focus on computing. 2) From visible to embedded; Hiding more of the computation certainly gives apparent advantages, but it also introduces new complexities in relation to human-computer interaction. This shift also invokes a second shift in relation to analysis and gathering of empirical material, in placing attention on the challenges posed in describing and translating implicit and invisible characteristics of the use context.

Interface and Interaction Design

The central focus of my investigations in this area, are interaction and control issues introduced in relation to the implementation of embedded, invisible and intelligent computation. The central questions addressed are; how do the actors, (human/non-human) access knowledge about a situational context, transform that into an interpretation of the situation, convey this interpretation into actions, and finally how do they make negotiation of that understanding and the corresponding action possible?

I have developed the concept of "situated interpretation" to describe the complex and multifaceted task of acting a situational context, an understanding that defines the initial basis for acting, secondary for negotiating and thus forms the basis for meaningful interaction. The concept is currently being developed and introduced in the design process of pervasive computing environments. It is the goal that analysis and evaluation of the situated interpretation processes underlying the concrete human-computer interactions in a given situation, can help form the basis for creating easy-to-use interfaces that are intelligible, accountable and sufficiently negotiable for the human user.

Experimental Systems Development

In the projects we have performed numerous experimental design sessions together with end-users, including evaluations and modifications of prototypes. The methods chosen and developed for these sessions have been focused at developing a useful translation from highly flexible, creative and individual work praxis' to the design of general technological environments and artifacts. Collecting empirical material through the use of pictures diaries, and analyzing empirical material in terms of "abstractions of actions" are examples of methods developed in the project, based on concerns about preserving creativity and individuality, and the balance between understanding and taking into account the present work situation on one side, and creating something new on the other. Due to space limitations I will not refer these methods in detail in the present paper.

REFERENCES

1. Bellotti, V. and S. Bly (1996) Walking away from the desktop computer: Distributed collaboration and mobility in a product design team. I Proceedings of CSCW'96, K. Ehrlich & C. Schmandt (Eds.), ACM Press, pp. 209-218.
2. Büscher, M., Krogh, P., Mogensen, P., & Shapiro, D. (2001) Vision on the move - technologies for the footloose. Appliance Design. Vol 1 (1), pp 11-14.
3. Fagrell, H., Kristoffersen, S. and Ljungberg, F. (1999) Exploring Support for Knowledge Management in Mobile Work. I Proceedings of ECSCW'99, S. Bødker, M. Kyng & K. Schmidt (Eds.), Klüwer, pp. 277 - 290
4. Kristoffersen, S., Ljungberg, F. (2000) Mobile Use of IT. Proceedings of IRIS 22, T.K. Kakola (ed.), Jyvaskyla University Printing House.
5. Luff, P. and C. Heath (1998) Mobility in Collaboration. In Proceedings of CSCW'98, S. Poltrock and J. Grudin (Eds.), ACM Press, pp. 305-314.
6. Norman, D. (1998) The Invisible Computer. MIT Press, Massachusetts Institute of Technology, Cambrigde.
7. Weiser, M. (1991) The Computer for the 21'st Century. In Scientific American, September 1991, pp. 933-940.

Digital Communication through Touch

Kurt Partridge
University of Washington
Box 352350
Seattle, WA 98195 USA
+1 206 543 5143
kepart@cs.washington.edu

ABSTRACT

Touch-mediation communication is a local communication technology in which a user directly controls the transmission of information between a wearable device and a ubiquitous device by his or her touch. Based upon intrabody communication, touch-mediated communication naturally limits the zone of communication to devices of interest and does not require conscious attention from the user. Implementations are power-efficient compared to other local communication mechanisms. I plan to construct two touch-mediated communication devices: a shoe and a bracelet. Each device will be evaluated on how well it correlates touch and communication.

Keywords

Communication, Bluetooth, IrDA, RFID, Intrabody Communication, Personal Area Networks

INTRODUCTION

Many applications require local exchange of digital information between a wearable device and a device in the environment. Everyday examples include electronic access cards and remote control devices. Ubiquitous computing applications will need local communication to identify users, transmit commands, and collect data.

Many communication media and technologies have been used for local communication, such as visible light for barcodes, infrared signals for IrDA, radio frequencies for Bluetooth, magnetic fields for RFID, and physical contact for smartcards and iButtons.

This work concerns a different approach: the transmission of imperceptible low-power electrical signals through the human body. The feasibility of "intrabody communication" (IBC) has already been demonstrated [5, 2]. I hope to show how to make IBC *touch-mediated*. That is, communication should be possible if and only if the same person simultaneously touches both IBC transceivers.

WHY USE TOUCH FOR LOCAL COMMUNICATION?

Touch-mediated communication offers benefits that other local communication media lack:

- The zone of communication is limited to the devices of interest. Other devices that happen to be nearby but are not touched cannot, in theory, interfere with or participate in the communication. If Alice operates device D, and D wants to communicate with Alice's PDA, there is no chance of accidentally mistaking Bob's PDA for Alice's. Media that enable communication between all nearby parties, such as radio waves, cannot distinguish Alice and Bob. Such media are also more easily snooped.
- Touch-mediated communication is not physically imposing. Some mechanisms, like barcode scanners, fingerprint readers, IrDA, and RFID require the user to hold a device or click an activation button. Touch-mediated communication leaves users' hands free to move easily between devices.
- Touch-mediated communication is also not mentally imposing. Touch is something that users would "have to do anyway" [1]. To borrow Weiser's terminology, touch easily "disappears." Touch's invisibility enables applications that would otherwise take too much user effort. For example, a wearable could keep a database of all touched objects. The user could query the database to determine how his or her usage of time, or an application could use the database to predict future context. If collecting the data required a conscious action on the part of the user, the database would be less complete and less accurate because of the overhead to add an entry.

There are also disadvantages to touch-mediated communication. First, data rates for IBC are low—our current prototype operates at 44kbps. The theoretical maximum speed is higher, but still under a megabit per second. Data rate issues are minor, however, because vertical handoffs can easily increase transmission capacity. Second, touch-mediated communication is inappropriate for devices beyond arm's reach. While some applications do need across-the-room communication, many others do not. And third, unintended actions may occur if a user unintentionally touches an object. If the object is another person, an unexpectedly long communication channel may be possible since the signal can flow through multiple people that are all touching each other.

PREVIOUS WORK ON INTRABODY COMMUNICATION

An IBC system is best understood as a single circuit that includes the person, transmitter, and receiver. The person behaves like a moderately high-valued resistor in series with small capacitors. The transmitter generates a signal with a frequency between 10kHz and 10MHz and a peak-to-peak amplitude from 5V to 50V. The receiver sees a greatly attenuated version of the transmitted signal, which it amplifies and otherwise processes before passing it up to the application.

The signal attenuation arises mostly not from dissipation of power inside the person's body, but rather from the lack of a common ground for the transmitter and receiver. The two circuits' grounds are coupled capacitively through the environment, using either nearby conductors or the miniscule capacitance of the surrounding air. As the user moves, the impedance of the ground connection varies. Consequently, the strength of the received signal varies.

Variations in received signal strength cause problems for touch-mediated communication. Ideally, a touch-mediated communication receiver would pick up a strong signal if the user touched another transceiver, and no signal otherwise. But if the grounding connection is too weak, the receiver may not detect the signal. Conversely, if the grounding connection is too strong the receiver may detect the signal when it should not. We have observed these effects in both informal experiences and formal experiments [3]. Poor connections are common with portable devices, which must couple circuit grounds through the air around the person. Undesired connections are common when two transceivers come within a few tens of centimeters of each other.

PLAN OF ACTION

I plan to investigate two techniques to address these shortcomings of IBC as a mechanism to support touch-mediated communication. Both techniques amplify the signal so it can be received even with a poor grounding connection. The signal is then ignored if touch is deemed not to have occurred.

The first approach uses more sophisticated signal analysis. The rate of change of signal strength, rather than a fixed threshold, may indicate when a user touches and stops touching a device. A low-frequency component may be produced from user "fidgeting," which changes the system geometry and therefore the received signal strength.

The second approach uses additional sensors to provide sufficient information to detect touch. We have constructed a vibration sensor [4], although other sensors would also work, such as electrical power-line noise, or the controls on a device's physical user interface.

Although these techniques should improve the correlation between touch and communication, problems may still arise, particularly if multiple people are present. For example, if a user without an IBC transceiver touches an IBC-enabled device, and another user with an IBC transceiver stands nearby, the device may believe that the nearby user is touching it. Even more pathologically, the user with the transceiver may touch the user without one.

Technical solutions to these problems are challenging. However, if signal analysis cannot distinguish these situations from genuine touch, feedback to the user may provide enough warning to avoid problems in a practical system.

I plan to construct a shoe and a bracelet to explore these ideas. Prior work [5, 3] has shown that transceivers located in shoes usually produce the greatest received signal strength, because the signal can couple through conductive surfaces in the floor. But couplers in shoes have particularly low signal strengths when communicating with portable devices. Bracelets have fewer problems with portables, but the overall weaker signal strength and potential for transceiver proximity makes them more susceptible to communication without touch.

I also plan to write software to simplify the construction of touch-mediated applications. An initial code base has already been developed.

EVALUATION

I will evaluate the proposed techniques for both devices in a setting where conditions can be reproduced, and in a practical setting where users are likely to perform unpredictable actions.

ACKNOWLEDGMENTS

Bret Sherman and Boris Startsev assisted with the application code base.

REFERENCES

1. Hinckley, K., and Sinclair, M. Touch-Sensing Input Devices. CHI 1999.
2. Matsushita, N., Tajima, S., Ayatsuka, Y., Rekimoto, J. Wearable Key: Device for Personalizing nearby Environment. ISWC 2000.
3. Partridge, K., Dahlquist, D., Veiseh, A., Cain, A., Goldberg, J., Foreman, A., and Borriello, G.. Empirical Measurements of Intrabody Communication Performance under Varied Physical Configurations. UIST 2001.
4. Partridge, K., and Borriello, G. A Hybrid Sensor Approach to Touch-Aware Communication. UW CSE Technical Report 02-07-02, 2002.
5. Zimmerman, T. Personal area networks: Near-field intrabody communication. IBM Systems Journal, 35(3):609—617, 1996.

Facilitating the Capture & Access of Everyday Activities

Khai N. Truong
College of Computing & GVU Center
Georgia Institute of Technology
801 Atlantic Drive
Atlanta, GA 30332-0280 USA
+1 404 545 1036
khai@cc.gatech.edu

ABSTRACT
There are situations in life when people need to recall specific details from previous experiences. To complement what can naturally be remembered, people will use tools to help create artifacts that preserve information pertaining to an experience. Unfortunately, the act of generating these records can sometimes conflict with people's abilities to fully appreciate the experiences. To assist in this process, many automated capture and access applications have been developed to relieve us from the burden of manually documenting these experiences. In this research, we develop an infrastructure to facilitate the construction of this class of applications for a variety of situations and domains.

Keywords
Ubiquitous computing, capture and access application, infrastructure, application development.

INTRODUCTION

Ubiquitous computing is a vision of technology seamlessly integrated into the environment to assist humans in their everyday lives [7]. One of the services envisioned is the automated capture of everyday experiences made available for future access. Automated capture and access applications leverage what computers do best – record information. In return, humans are free to fully engage in the activity and to synthesize the experience, without having to worry about tediously exerting effort to preserve specific details for later perusal.

As ubiquitous computing becomes more and more pervasive, it is important to understand instances where capture and access services are desirable and when it is inappropriate. One way of learning how people can benefit and use automated capture and access services is by actually evaluating its authentic use in these environments. Over the years, researchers have constructed capture and access applications for classrooms [1], [4] and meeting rooms [3], [5]. However, there still exist many more domains that could potentially benefit from the same kinds of services, such as the home and the car.

To facilitate research in ubiquitous computing, advances are necessary to improve the tools available for those with creative applications to realize their visions. Toolkits and infrastructures can provide levels of abstractions as well as reusable building block that can aide in the development process of a system. In this research, we have developed INCA (an infrastructure for capture and access applications) aimed:

1. to lower the barrier for building capture and access applications,
2. to make complex applications more evolvable, or capable of being fine tuned to meet the users' true needs, and
3. to support the construction of a wider variety of applications previously unexplored.

CAPTURE AND ACCESS REQUIREMENTS AND ISSUES

In his 1945 Atlantic Monthly article, Vannevar Bush described his vision of the *memex*, a generalized capture and access application [2]. He noted that a "record … must be continuously extended, it must be stored, and above all it must be consulted." The point of every capture and access application is to preserve some portion of an experience so that it can be ***accessed*** in the future. While the specifics of why users want capture and access services are domain dependent, from the users' perspective, all capture and access applications need to support:

- the preservation of details from an experience,
- the marking of associations between these captured artifacts, and
- the availability of all information on a topic from previous experiences when needed.

To meet these broad goals, application developers need to provide support:

- for ensuring that information is captured,
- for associating and aggregating related information,
- for keeping information available for later use, and
- for providing information when users want it.

AN INFRASTRUCTURE FOR CAPTURE & ACCESS

Capture and access applications are typically comprised of a confederation of heterogeneous components that must work seamlessly together. As a result, application developers are often forced to exert large amounts of effort and time at tediously creating the "glue" that allows these independent, distributed, and heterogeneous systems to work together. However, these are *accidental* tasks – tasks not directly related to the developer's primary goals in the development of the application. This points out the potential for lower level support to aid in the construction of capture and access applications.

Many capture and access applications have been built as one-off prototype demonstration vehicles, not meant to be extensible over time. However, the architectural similarities that exist across these applications point to the potential creation of reusable components to reduce future development time and effort. An examination of this class of applications reveals similarities in functionality, where some part of the system is responsible for:

- the ***capture*** of information,
- the ***storage*** of the information,
- the ***transduction*** of information into different formats,
- the ***access*** of the information, and
- the ***integration*** and ***synchronization*** of related information when multiple streams of information are captured.

We have developed an infrastructure to leverage on these natural separation of concerns common across capture and access applications, while abstracting away some of the complexities that arise in the development of this class of applications, such as:

- ***Networking***: providing developers with a single, repeatable design process for building both networked or stand-alone applications in the same fashion.
- ***Information management***: generically capturing, storing, and delivering all information streams as raw bytes tagged with contextual metadata attributes.

Additional features include support for observing the capture and access system's run-time state as well as the ability to control capture and access allowing end-users and application developers to address privacy concerns. Through these separations of concerns, key architectural abstractions, a toolkit of reusable services, and other additional features, this infrastructure addresses the *accidental* tasks in the development process of the applications themselves. As a result, developers can then focus on solely building applications to meet the users' needs.

RESEARCH HYPOTHESES

As part of this research, we will use the infrastructure to evaluate the following hypotheses:

- The infrastructure will allow application developers to more easily build and evolve capture and access applications.
- Designers can build solutions that are reusable by other application developers.
- INCA can support the development of a wider variety of applications previously unexplored.

SUMMARY AND FUTURE RESEARCH

Our research is aimed at providing application developers with an infrastructure that makes it easier to build capture and access systems. We have completed the first version of the capture and access infrastructure and have made this distribution publicly available. We have provided the infrastructure to members of our research group as well as several other research institutes. We will investigate the infrastructure's ease-of-adoption by programmers. Feedback from the developers will be elicited to understand the kinds of applications that they are actually able to build and the difficulties they encounter.

Additionally, we have done an extensive study of related work, and have described a set of dimensions for capture and access applications [6]. Using the results from this study we are able to identify the design space for capture and access applications and the holes in this design space. We will re-implement several applications from the literature to demonstrate INCA's support for aiding in the construction of existing applications. We will demonstrate INCA's support for building a wide range of applications by varying the extremes for all dimensions of the design space; thereby proving its potential support for all possible applications.

ACKNOWLEDGMENTS

I would like to thank my advisor, Gregory D. Abowd, for all his support, encouragement and guidance.

REFERENCES

1. Abowd, G.D., Classroom 2000: An experiment with the instrumentation of a living educational environment. IBM Systems Journal. **38**(4) (1999), 508-530.
2. Bush, V. As We May Think, in *Atlantic Monthly*. (1945)
3. Chiu, P., et al. NoteLook: Taking Notes in Meetings with Digital Video and Ink. In *Proceedings of ACM Multimedia'99* (Orlando FL, 1999), 149-158.
4. Davis, R.C., et al. NotePals: Lightweight Note Sharing by the Group, for the Group. In *Proceedings of CHI'99* (Pittsburgh PA, 1999), 338-345.
5. Minneman, S., et al. A confederation of tools for capturing and accessing collaborative activity. In *Proceedings of ACM Multimedia'95.* (San Francisco CA, 1995), 523-534.
6. Truong, K.N., Abowd, G.D., and Brotherton, J.A. Who, What, When, Where, How: Design Issues of Capture & Access Applications. In *Proceedings of UBICOMP 2001* (Atlanta GA, 2001).
7. Weiser, M., The computing for the 21st century. Scientific American. **265**(3) (1991), 94-104.

Ubiquitous Computing: Transparency in Context-Aware Mobile Computing

Louise Barkhuus, Ph.d. student
Department of Design and Use of IT
The IT University of Copenhagen, Glentevej 67, DK-2400 Copenhagen, Denmark
barkhuus@it.edu

June 30, 2002

Ubiquitous Computing

With the increasingly distributed use of computers and the wide range of available computing devices, our society is experiencing a previously unknown level of mobile computing. The adoption of mobile communication technologies such as mobile phones and PDAs is an example of the mobility that constitute our society today and the pervasiveness that mobility contribute with, is evidence of a changing computing paradigm.

Ubiquitous computing describes a widely networked infrastructure of a multitude of computing devices. It moves the interaction beyond the desktop and into the real world with a special attention to activities of everyday life. According to Mark Weiser the vision is to get the computer "out of the way, allowing people to just go about their lives" [5]. The criteria of transparency is then fundamental to the paradigm of ubiquitous computing. The transparency imply more than just a user-friendly interface; the technology should facilitate the task in a non-intrusive way and in this way "hide" the underlying technology for the user. The questions to pose however, is how this transparency is acquired and if mobile devices at their present state are to any degree transparent?

In order to study the transparency within mobile devices I introduce the concept of context-awareness, a central aspect of ubiquitous computing. The context-aware computing describes a scenario where the computing device knows its own present context and acts accordingly. This scenario, however, is highly complex since it requires a closed community where the computing takes place, contrasting mobile computing that requires a high level of mobility. The solution at the present state of consumer information appliances and mobile devices is that of *context-dependent* devices, devices that let the user define the context and then act accordingly. One example is the mobile phone that offers profiles for high, low and soundless settings. The two concepts of context-awareness and context-dependency are close and often overlapping however, in my research they will be defined separately in order to be comparable.

Problem Statement

The main question within my research is how the user experience of transparency within a ubiquitous computing environment can be affected by context-aware mobility compared to context-dependent mobility. Context-aware mobility is defined as mobile applications that change according to context where context-dependent mobility just requires some alteration by the user. The area of ubiquitous computing in question will be limited to mobile devices and their use in a consumer context.

What I hope to accomplish with my research is to contribute to a theoretical framework of how mobile technology can acquire the status of "transparent" as people use the technology as tools for specific tasks. I will attempt to provide further knowledge into the area of user centered ubiquitous computing by demonstrating how transparency is acquired or perhaps not acquired in relation to use of mobile technology. The theoretical framework will consider the aspect of context as one attempt to acquire transparency.

Theoretical Background

Drawing on Weiser's definition of ubiquitous computing [4], [5], I intend to use a conceptual framework that includes his definition of transparency in combination with other research conducted in the area. As for context-aware applications, there are still few, if any, outside research laboratories at the present state and moreover, most research falls into the category of location-based services [1]. The theoretical framework of context-awareness is therefore limited but researchers such as Dey, Abowd, Salber [1] as well as Schilit and Theimer [2] have all made extensive work within the area. Finally, the situated actions as defined by Lucy Suchman [3] is an area that I base my research on. The concept is relevant to the use of mobile devices in that the situations of use are difficult to specify in opposition to plans and fluctuate according to context. Thereby they create a flow of different situations where the necessary actions are not interrelated.

Planned Research

To study the transparency of context-aware and context-dependent technology, the level of transparency within specific applications should be measured. Both context-aware and context-dependent applications should be considered in order to compare these. Here, the method of a case study is proposed as one way of acquiring empirical data. The case study will provide data of the actual use of mobile devices and applications in real settings.

I am currently doing a preliminary, literary study of the two concepts transparency and context-awareness/dependency. The study will conclude on the past and current use of the terms in relation to mobile computing and an operational definition is to be developed as a foundation for my further research. I am defining the concepts in regard to other researchers' definitions as well as my own considerations, in order to define usable concepts that can be measured and thereby studied as part of my research.

There are two case studies that I intend to carry out. First, I will attempt to find if the present mobile devices are to any degree transparent in their use. The study will rely on user logs, interviews and observation in order to determine the users perception of their devices. The user group will include mobile phone user, PDA users and combined users. The goal of the study is to compare the perceived transparency (defined according to trends in ubiquitous computing) in mobile phones with that of PDAs and by this attempting to define their present level of transparency according to my definition. Second, I plan to do a large-scale case study within "Crossroads Copenhagen", a newly formed research collaboration, which includes a project of location-based mobile services and technologies. Here, the aspect of context-awareness is essential and actual large-scale implemented applications are candidates for further study. The collaboration is presently in its initial phase and I am currently attempting to structure a case study that takes place in this context.

References

[1] A.K. Dey, G.D. Abowd and D. Salber: "A Conceptual Framework and a Toolkit for Supporting the Rapid Prototyping of Context-Aware Applications", *Human-Computer Interaction* vol.16, 2001, pp. 97–166.

[2] B. Schilit and M. Theimer: "Disseminating active map information to mobile hosts". *IEEE Network* 8(5), 22–32.

[3] L. Suchman: *Plans and Situated Actions.* Cambridge: Cambridge University Press, 1987.

[4] M. Weiser: "The Computer for the 21st century", *Scientific American,* 265(3), 1991, pp. 66–75.

[5] M. Weiser: "Some Computer Science Issues in Ubiquitous Computing, *Communications of the ACM,* 36(7), 1993, pp. 74–84.

Using Autonomous Agents to Maintain a User Presence in Ubiquitous Collaborative Environments

Marcela Rodríguez
Computer Science Department, CICESE, Ensenada, México
marcerod@cicese.mx

ABSTRACT
Handheld computers are becoming a natural medium to tap into an ubiquitous computing infrastructure. Handhelds, however, most often operate disconnected from the network thus reducing the opportunities for computer-mediated collaboration with colleagues or computational resources. We propose the use of autonomous agents that represent the user and are able to make decisions on his behalf while he is disconnected or inactive. We present an extension to the COMAL handheld collaborative development framework to support autonomous agents that can act on behalf of the user in pervasive computing environments.

Keywords
Ubiquitous computing, autonomous agents, CSCW.

INTRODUCTION

Due to the personal nature of handheld computers, most of its applications today are oriented to single users and require limited or no connectivity at all. However, the users of these devices work in collaborative environments in which the need to exchange information and share ideas with others is very clear. Thus, as these devices become more pervasive and support network connectivity, we expect them to become major players in future ubiquitous collaborative environments [5]. Furthermore, the interconnected infrastructure as a whole should be able to sense the context in which a specific situation is taking place and adapt to it according to its location of use, the people and objects that are around, and changes of those entities over time [4].

Increasingly, handhelds provide alternatives for network connectivity, yet they are most of the time inactive or disconnected from the network, which severely limits their use for ubiquitous collaboration. In these circumstances, autonomous agents that act on behalf of the user and reside on a desktop computer or trusted server, might be able to maintain a limited user presence and execute actions on his behalf. To explore the potential of such an approach, we have extended COMAL (Collaborative Mobile Application Library) [1], a handheld collaborative development framework, to support the development of autonomous agents in ubiquitous collaborative environments.

UBIQUITOUS COLLABORATIVE APPLICATION DEVELOPMENT FRAMEWORK

We analyzed use scenarios of autonomous agents in ubiquitous collaborative environments, which helped us determine key issues to be addressed in the development of the framework. Some of these use scenarios are presented next:

A user is co-authoring a research paper with a colleague and he needs to incorporate his final contributions. However, the latest version of the paper is currently locked by his co-author who is attending a conference. The user sends a message to his co-author's agent who will decide, based on the context, whether or not to liberate the resource.

When the user arrives at a conference with multiple simultaneous tracks, he fills a form on his PDA in which he specifies his main interests. He connects his handheld to a point of presence to send his profile, which will launch an agent in the conference server to build for him a recommended personalized schedule. The schedule is stored in a server to be downloaded to the handheld the next time he is connected to the network.

Finally, at the attendee's handheld may appear agents which represent devices available in the private network of the conference site, such as, a public printer. The attendee can become aware of the printer status and send a paper that he wants to share with a colleague he met at the conference.

Based on the scenarios just presented, we identified the following requirements for the application development framework:

1. Support for disconnected mode of operation
2. A consistent development API
3. Support for different communication modes
4. Agents are created automatically or explicitly by the user
5. Agents should be able to communicate with its user and with other agents in other desktops.
6. Additionally, agents may represent users as well as devices or services available on ad-hoc networks.

The first three requirements are satisfied by COMAL [1], an architecture and a set of application libraries designed to build Shared Objects for Portable Environments, or SOPE applications. COMAL libraries exist only in the desktop and handhelds computers. Developers using this framework can then concentrate more on building the ubiquitous collaborative solution itself, rather than doing low level programming. To satisfy the last three requirements, we have extended COMAL with SALSA (Simple Agent Library for Sope Applications). This extension allows users to implement simple agents on top of SOPE's to act on behalf of the user while he is disconnected, or for representing devices that are within reach of the user.

COMAL currently supports the Palm Computing Platform with PalmOS.

SALSA: EXTENDING COMAL WITH AUTONOMOUS AGENTS

SALSA is located on top of the Desktop COMAL library. An agent programmed using SALSA includes components for perception, action and reasoning as illustrated in Figure 1. The perception module accesses the Application Database, and feeds the reasoning subsystem, which governs the actions, including deciding what to perceive next. The application database is a component of the COMAL framework and includes information synchronized with the handheld. The actions that can be triggered by an agent include sending a message to the user or creating a new agent with a predefined subgoal that can be automatically launched or sent to the user for him to decide whether and when to execute it. The reasoning component can be implemented as a simple action/perception rule or with a more sophisticated algorithm. This is left to the user based on his application's logic.

SALSA also provides the set of infrastructure services to support agent interaction, as well as conventions that allow them to understand and communicate with each other and with the user. To implement an agent, SALSA defines an abstract *Agent* class, which controls the life cycle of an agent. The library provides a set of primitives through which the desktop application can *create* and *kill* an agent, *communicate* with other agents residing on another desktop through XML messages, evaluate user's interests and execute tasks. Once the agent is created, it is registered in an Agent Directory which provides a yellow-page service that allows the application to identify which agents provide what services; a white-page services that allow agents to locate each other.

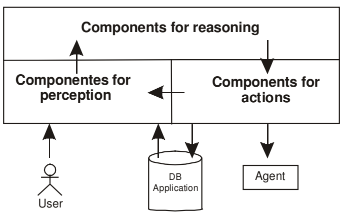

Figure 1. SALSA's Agent Architecture.

A SAMPLE APPLICATION

Figure 2 illustrates the use of the SALSA primitives and the COMAL library to implement a Personal Conference Scheduling Application based on the use scenario described previously. We can appreciate how the agent interacts with the system components. The messages shown in bold correspond to the agent library while the others are part of COMAL. When the user has connected his handheld to an access point to *send* his profile (see Figure 3), an agent is *create*d and *activate*d in the conference server. The agent *evaluate*s this information to generate a schedule, which is store in the application database and will be integrated in the calendar application of the handheld the next time it is synchronized using a point of access.

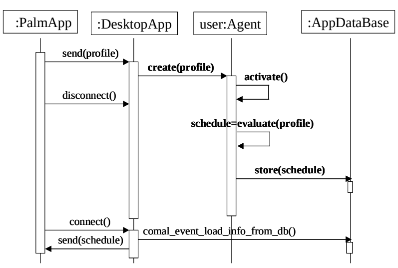

Figure 2. Behaviour of the conference scheduling agent.

Figure 3. PDA's Form to specify user interests.

RELATED WORK AND DISCUSSION

The proliferation of different computing devices such as handhelds is stimulating the development of frameworks that allow developing mobile collaborative systems for wireless or ad-hoc networks. Moreover, there is a migration of technologies originally developed for PC's to the realm of handhelds and wireless networks. Collaborative Systems Development APIs and Agent technology is following this downsizing trend, and many projects are under way to enable these systems on handheld devices. One of these projects is the LEAP [2] project, which implements an agent platform that runs seamlessly on both mobile and fixed devices. Nevertheless, this platform does not provide any means for implementing service-level interoperability or agent autonomy.

The DACIA [5] and Ycab [3] projects proposed frameworks that allow developers to create collaborative applications that can run on both mobile and fixed devices. YCab's architecture avoids a single point of failure by distributing the control over multiple nodes of the ad-hoc network. In DACIA, while the user is disconnected, a mobile application can be parked, which can continue to interact with other parties on behalf of the user. In contrast, our project allows the creation of autonomous agents that represent the user while he is disconnected.

The projects just described, do not contemplate the need to write data on the featured applications of the mobile device, such as the Calendar or Date Book applications. Also, they do not support communication via the PDA's infrared port as allowed under our approach.

REFERENCES

1. Alba, M., Favela, J. (2000). *Supporting Handheld Collaboration through COMAL* Proc. of CRIWG'2000, 52-59
2. Bergenti, F, and Poggi, A. (2001). *LEAP:A FIPA Platform for Handheld and Mobile Devices.* Proc. of ATAL'2001,436-446
3. Buzko, D. et al. (2001). *Decentralized Ad-Hoc Groupware API and Framework for Mobile Collaboration.* Proc. of Group'2001, 5-14.
4. Dey, A. (2001). *Understanding and Using Context.* Personal and Ubiquitous Computing, Vol. 5, No. 1,.4-7.
5. Litiu, R., Prakash, A. (2000). *Developing Adaptive Groupware Applications Using a Mobile Computing Framework*. Proc. of CSCW'2000, 107-116.

Using Internet Services to Manage Massive Evolving Information for Ubiquitous Computing Systems

Emre Kıcıman
Stanford University
emrek@cs.stanford.edu

ABSTRACT

A major challenge in building ubiquitous computing systems is the large variety of heterogeneous devices. Building applications that cope with this heterogeneity requires managing massive amounts of quickly evolving information, mapping among the various semantically-equivalent functionalities of devices. We advocate using Internet services to store and collect these mappings. We are implementing this hypothesis in the IMHome project, focusing on the spontaneous interaction of media creation, storage, and playback devices. To evaluate IMHome, we are developing metrics to measure the dependability of IMHome-based applications.

Keywords

Internet services, Ubiquitous computing, Massive evolving information, Heterogeneity, Dependability

INTRODUCTION AND MOTIVATION

One of the major challenges in building a ubiquitous computing (ubicomp) system is ensuring that it can scale to the large degree of heterogeneity present in ubicomp environments [3, 1]. The difficulty is that an application must explicitly know how to interoperate with the massive variety of devices in real-world ubicomp environments. Even more challenging, new devices are being developed every day.

For example, consider a ubiquitous computing application that plays the first thirty seconds of every song in a user's music collection—perhaps as part of a "name that tune" game. This is not complicated functionality and seems quite easy to build—except that this application must be able to control everything from CD players to MP3 players to tomorrow's next-generation audio player, all with slightly different capabilities and potentially significantly different control interfaces. And while it is easy to build an *ad hoc* application for any specific stereo system, it is practically impossible to extend the same *ad hoc* application to support an open-ended and evolving variety of audio players. The application simply has to know too much.

MASSIVE EVOLVING INFORMATION

This heterogeneity problem is a member of a small but important class of ubicomp challenges we call *massive evolving information* (MEI) problems. We characterize this class of problems with two properties:

1. For any specific, static scenario, a simple *ad hoc* solution exists, based on *a priori* knowledge about the scenario.
2. A solution in a general, dynamic environment is theoretically achievable. Practically however, the information required is too massive and/or changes too rapidly to be feasibly stored and managed in a local environment.

Other MEI problems include associating annotations (such as product reviews) with physical artifacts, and mapping machine-level identifiers to human-understandable names.

To create a general solution to an MEI problem, we must either reduce the amount of information that must be known (such as through standardization) or we must find a way to store and manage large, dynamic data sets in ubicomp environments. Standardization efforts have been made for years, but have so far failed to stabilize due to changing requirements and technologies. And though we can store increasingly large data sets in ubicomp environments, we still have issues distributing updates at the scale required by many MEI problems.

USING INTERNET SERVICES TO MANAGE MEI

In contrast to our ability to store and update massive evolving information in ubicomp environments, there are well-known techniques for scaling Internet services to manage large amounts of data, large numbers of queries, and rapid updates. Today's Internet services routinely manage massive amounts of data.[1]

By building an Internet service to manage our massive and evolving information and accessing it remotely from our ubicomp environments, we reduce the particular MEI problem to a domain-specific issue of how to generate the large data set and keep it up-to-date. Though this is potentially still a difficult issue, it is generally more tractable.

One example that illustrates how Internet services can be used to solve MEI problems is found in today's Internet CD databases.[2] The problem of linking a music CD to its album information and related music is an MEI problem, since 1) it is trivial to build a system which stores related information for a static set of albums; and 2) the rapid rate at which new CDs are published makes it difficult to build a local system with knowledge of all CDs.

[1] At the extreme, the Internet Archive (archive.org) manages over 100TB of data. Major search engines routinely manage multiple terabytes of data.

[2] Popular versions of this service include http://www.CDDB.com and http://www.FreeDB.org/

A partial solution is to embed this information in the CD itself, but that precludes linking, say, an artist's first releases to their later albums. Instead, the solution used today is to use an always-accessible Internet service to store and manage information related to CDs, and distribute the work of entering new CD information across the users of the service.

RESEARCH HYPOTHESIS

We hypothesize that by using Internet services to manage massive, evolving data sets, we can build a useful solution for solving the semantic heterogeneity problem in ubiquitous computing. Specifically, we propose to use an always-accessible Internet service to manage approximate-equivalence mappings among the functionalities of and control commands for various heterogeneous devices.

Using Internet services to logically centralize the storage of this mapping information leaves us with two subproblems. First, to use approximate-equivalence mappings among devices' functionalities, there must be agreement on the mechanics of accessing those functions. Current trends in self-describing data formats and communication protocols such as XML and tuplespaces are providing these standard mechanisms. Additionally, their clean separation of syntax from semantics makes it easier to apply mediation techniques to transform among formats [2].

Secondly, as new devices are built and new classes of functionalities are developed, we must update the equivalence mappings in our Internet service. The simplest method for doing this is to provide the end-user of the system with the option of manually providing a mapping when one is missing. To help willing users determine the correct mappings, we can provide them with human-language descriptions of devices and functionalities as extra context. Once a new mapping is created, it can be pushed back to our Internet service to be used by others.

IMHOME

We are testing our hypothesis in the context of IMHome, a project focusing spontaneous interoperation of devices for media creation, storage and playback in the home. Relevant details of the IMHome architecture include its use of a *task abstraction* to organize cooperating devices and services; a simple, self-describing *property-based control interfaces* for devices and tasks; and *property mapping services*, deployed as Internet services, to programmatically transform the control interfaces of devices and tasks.

EVALUATION

The problem of evaluating systems software in ubiquitous computing environments extends well beyond the scope of this thesis summary. However, since our community has not yet settled on specific evaluation methods, we believe it is worth discussing.

Though researchers in the systems software community have traditionally concentrated on the performance aspects of systems, their focus is now moving towards dependability [4]. This new focus has special significance for ubicomp, where a dynamic and fragile environment poses unique challenges to building robust systems [5].

We believe that an important axis for measuring systems software contributions to ubicomp is whether they improve the dependability of ubicomp applications. Contributions which enable new functionality should ideally improve dependability as well; any contribution that compromises dependability should be met with skepticism.

By enabling applications to more robustly adapt to ubicomp's heterogeneous environment, we expect our framework to improve the overall dependability of ubicomp applications. To verify this hypothesis, we are developing metrics for measuring the dependability of IMHome applications based on its task and control abstractions.

SUMMARY AND FUTURE RESEARCH

The goal of our research is to provide a useful solution for spontaneously using heterogeneous devices together. We are focusing our efforts on media devices in the home environment and are completing a first prototype of the IMHome system. We are currently integrating multimedia devices into this system and developing sample applications. We have recognized that the management of massive evolving information is key to practically solving issues of semantic heterogeneity, and are beginning to build a prototype Internet service for managing the equivalence-mappings of IMHome device functionalities. We plan to evaluate our system based on how much it improves the dependability of ubicomp applications.

ACKNOWLEDGEMENTS

I would like to thank Atsushi Shionozaki for his feedback on the IMHome project. This work has been supported by a Stanford Graduate Fellowship, an NSF Graduate Fellowship, and a gift from the Sony Corporation.

REFERENCES

1. W. Edwards and R. Grinter, At Home with Ubiquitous Computing: Seven Challenges. In *Third Intl. Conference on Ubiquitous Computing (Ubicomp2001)*, 2001.

2. E. Kıcıman and A. Fox, Using Dynamic Mediation to Integrate COTS Entities in a Ubiquitous Computing Environment. Proceedings of *The Second International Symposium on Handheld and Ubiquitous Computing 2000.*

3. T. Kindberg and A. Fox, System Software for Ubiquitous Computing. *IEEE Pervasive Computing Magazine* 1(1):7081, January 2002.

4. D. Patterson, et al. Recovery Oriented Computing: Motivation, Definition, Principles, and Examples. *UC Berkeley Computer Science Technical Report* UCB//CSD-02-1175, March 15, 2002.

5. S. Ponnekanti, B. Johanson, E. Kıcıman and A. Fox. Designing for Maintainability, Failure Resilience, and Evolvability in Ubiquitous Computing Software. In Submission to *Operating Systems Design and Implementation 2002.*

iClub, An Interactive Dance Club

Joshua Samberg and Armando Fox
Computer Science Dept., Stanford University
Stanford, CA 94305-9040 USA
+1 650 723 9558
gfunkj@stanford.edu, fox@cs.stanford.edu

Maureen Stone
StoneSoup Consulting
stone@stonesc.com

ABSTRACT

We describe iClub, an interactive dance club application developed to run in Stanford University's Interactive Workspace ("iRoom") environment. Because of the properties of iROS, the Interactive Workspace Operating System on which iClub is built, it is easy to extend iClub with new A/V modules, integrate new physical input devices such as those being developed by iStuff, and tolerate transient failures of hardware or software during an iClub session.

Keywords

Interactive workspace, ubiquitous computing, entertainment, music

A DANCE CLUB USING INTERACTIVE WORKSPACES TECHNOLOGY

iClub is an attempt to experiment with some of the ways Interactive Workspace technology might be used in an entertainment setting in the near future. The iClub is an application developed to run in Stanford University's Interactive Workspace ("iRoom") environment [5] and is built using iROS, the open-source Interactive Room Operating System [4] developed as part of the Interactive Workspaces research agenda. As shown in the video, iClub combines a computer-controlled DJ with visuals synchronized to the music, the ability for guests to influence music playback (speed, low pass audio filter, interactively-triggered sound effects, and voting on which song to play next), and the ability to integrate physical UI devices such as sliders and wands (no wands in the video) into the environment.

Because of the properties of iROS, it is easy to extend iClub with new A/V modules, integrate new physical input devices such as those being developed by iStuff [1,7], and tolerate transient failures of hardware or software during an iClub session. iROS was specifically designed to accommodate legacy and COTS (commercial off-the-shelf) devices, so everything in iClub runs at user-level on unmodified Windows 2000 or Linux PC's.

COORDINATION-BASED PROGRAMMING IN iROS

The programming model for iROS is one of ensembles of independent software entities that communicate by passing messages called *events.* Events are passed using a logically-centralized broadcast-like communication substrate called the *Event Heap* [3], which combines the properties of a tuple space [2] with the semantics of publish/subscribe. Entities retrieve or subscribe to events using templates that contain precise values for fields to be matched and wildcards elsewhere. The Event Heap thus provides *referential decoupling*: the intended recipients of an event are determined by the contents of the event itself rather than being directly named. This single property enables both application-level multicast ("this event is for all display controllers managing Bubbles") and location-independent naming ("this event is for the DJ, whichever machine it may be running on").

iClub ARCHITECTURE AND IMPLEMENTATION

The DJ

The iClub is implemented using OpenGL and the iROS client libraries, in C++ with a few portions in Java. The iClub architecture is centered around the DJ, a process that selects MP3 files to play based on recent input from the Voting modules. Voting can be done from handheld or pocket PC's with Web browsers or at fixed kiosks in an iClub. The DJ extracts and broadcasts a *beat clock* event synchronized to the song currently playing; visuals such as the Disco Ball and Dancing Dolls use this to synchronize their displays to the music. The DJ also broadcasts periodic *frequency histogram* events, used by other visuals such as the Ring of Fire and Graphic Equalizer Display. This architecture makes it easy to extend the iClub with new visuals that respond to the same events.

Guest Interaction

Each touch-sensitive display screen is managed by a Display Controller process that subscribes to any event having that display as a target. The Controllers manage the Dancing Dolls and Bubbles associated with each display. When a guest touches a Bubble, a BubblePop event is published. The DJ picks up this event and uses it to play a sound effect, while the Mermaid picks it up to help her keep track of the total number of Bubbles in the iClub. A guest can also touch the Turntable to change the speed of the currently playing song. If a Bubble is about to move off the edge of a display, that display's Controller publishes a BubbleExit event, which is picked up by the Controller of the display the bubble should reappear on. (The Controllers all read the display interconnection geometry from the same file at startup.) New guest-controllable behaviors can be added to iClub just by making each behavior emit the appropriate event for consumption by other modules in the room, such as the DJ or visuals. .

Integration of physical UI using event interposition

The original iClub design allowed guests to control the Low Pass Filter using a software slider widget. We wanted to integrate the physical iSlider, a prototype from the iStuff project [1] that generates its own event type. The referential decoupling supported by the Event Heap allows us to write a simple Listener that listens to iSlider events and emits the corresponding Low Pass Filter events for the DJ.

Partial failure semantics

A main goal of iROS was to provide robustness against transient failures [6]. Since each iClub module is an autonomous process, the failure of any module has no effect on others, and when the failed module is restarted it will resynchronize to the beat. (Automatic restarting will be a feature of iROS by the end of 2002.) If the DJ or the Event Heap crashes, the other modules "freeze" until the DJ or Event Heap is restarted, then continue where they left off.

SUMMARY

The iClub demonstrates that Interactive Workspace technology and iROS can serve as the basis of entertainment-based applications; although iROS was designed to be quite general, it is safe to say its designers did not foresee its use in this application domain. iClub demonstrates the effectiveness of several key iROS design goals: the ability to integrate heterogeneous hardware and software modules to create new behaviors; the ability to rapidly integrate new UI's (including physical UI's) to control existing behaviors; robustness to transient failures; and extensibility via incremental integration of new components.

iClub was implemented in about ten weeks by undergraduates Joshua Samberg, Xiaowei Li, Yeon Jin Lee, and Kabir Vadera as part of a senior programming project course. This suggests that iROS's other major goal, ease of use and expressiveness for application developers, has been realized.

REFERENCES

1. Borchers, J., Ringel, M., Tyler, J., and Fox, A. Stanford Interactive Workspaces: A Framework for Physical and Graphical User Interface Prototyping. *IEEE Wireless Communications*, *Special Issue on Smart Homes*, 2002 (in press).

2. Gelernter, D., and Carriero, N., Coordination Languages and their Significance, *Commun. of the ACM* 32(2), February, 1992.

3. Johanson, B. and Fox, A. The Event Heap: A Coordination Infrastructure For Interactive Workspaces. In *Proceedings of Fourth IEEE Workshop on Mobile Computing Systems and Applications (WMCSA 02),* Callicoon, NY, June 2002.

4. Johanson, B., Ponnekanti, S., Kiciman, E., Sengupta, C., and Fox, A. System Support for Interactive Workspaces. Stanford Computer Science Technical Report CS-2001-01.

5. Johanson, B., Fox, A., and Winograd, T. The Interactive Workspaces Project: Experiences with Ubiquitous Computing Rooms. *IEEE Pervasive Computing Magazine* 1(2), April-June 2002.

6. Ponnekanti, S., Johanson, B., Kiciman, E., and Fox, A. Designing for Maintainability and Failure Resilience in Ubiquitous Computing System Software. Submitted for publication.

7. Ringel, M., Tyler, J., Stone, M., Ballagas, R., and Borchers, J. iStuff: A Scalable Architecture for Lightweight, Wireless Devices for Ubicomp User Interfaces. Poster in adjunct proceedings of *Ubicomp 2002*.

Private and Public Spaces –
Video Mediated Communication in a Home Environment

Yngve Sundblad, Staffan Junestrand, Sören Lenman, Björn Thuresson, Konrad Tollmar
CID – Centre for User Oriented IT Design
KTH
S-100 44 Stockholm, Sweden
{yngve, stefanju, lenman, thure, konrad}@nada.kth.se

ABSTRACT

This video demonstration is based on scenarios of a family's everyday activities supported by video mediated communication (VMC). It was recorded in 1999 at comHOME, a concept dwelling of the future, built at and with Telia Network in Stockholm. The principal issue explored in the comHOME project, and in the video, concerns various aspects of private and public spaces using VMC. The design concept is based on the integration of different comZONES (communication zones), where the resident can be seen and/or heard.

The architectural space, then, in combination with information and communication technology (ICT) solutions forms an interface to the digital world. A main observation from the making of the video is that it is a very good complementary method in a complex design-process because of the focus on the user perspective.

Keywords

Video mediated communication, home architecture, privacy, boundaries, design methods

INTRODUCTION

Several trends indicate that VMC (video mediated communication) will become an important part of communication in our homes, [1]. VMC can support and complement a wide range of activities in that context, e.g., studies, care of the elderly, professional work and leisure activities. However, dwellings are ill suited for VMC due to, e.g., unsatisfactory lighting conditions, floorplan layout and spatial design. Also, current VMC solutions for collaborative work are not well adapted for the home.

Other concerns are social and emotional aspects of communication, which is a requirement for a domestic environment.

comZONES

The comHOME dwelling at Telia Networks in Stockholm, which was used for recording the video, is a laboratory and a showroom for a dwelling of the future, [2]. The comHOME project covers several aspects of future dwellings, such as making the apartment smart, but the primary goal has been to develop and integrate VMC solutions into a home. The design of the dwelling is based on the idea of creating different comZONES to support the demands for both private and public digital spaces within the home environment [3]. The comZONEs have the following characteristics:

- In an inner, public zone, one can be both seen and heard.
- In a middle, semi-public zone, one can be seen but not heard.
- In an outer, private zone, one can neither be seen nor heard.

These spatial characteristics may also vary over time, depending on the scenario of use.

The principal architectural issue was the establishment of the mental and physical boundaries between the public and the private in the comZONES, i.e., to uphold the demand of neither being seen nor heard - when so desired. The comZONES are expressed by technical solutions such as screens and cameras, but also through the use of architecture - spatial forms, colours, light and materials. Thus, the architectural space in combination with ICT solutions forms an interface to the digital world.

The creation of the different comZONES in the rooms of comHOME is a major technical undertaking, needing control of focus depth and field of view for video space and the control of the audio space, which is more complex, however. The fairly precise video space cannot be matched with equally well-defined boundaries in audio space.

New technology that might solve part of this problem is array- microphones, spatially directed loudspeakers, and real-time image and audio manipulation that can filter background actions and sounds.

An additional means of protecting privacy while maintaining continual contact is to, in some situations, replace VMC with a shared 3D digital environment. Here, rendered user representations provide an abstraction of information that can act as a filter for what is kept private or made public, [4].

THE SCENARIOS IN THE VIDEO

The video shows four scenarios, each one demonstrating the idea and some functions of a comZONE.

workPLACE

The workPLACE is a place for professional communication, located in the combined home office and bedroom. A table with two sideboards and a lowered ceiling with integrated lighting spatially defines the inner public zone, where the resident can be seen and heard. In the video Christine moves in and out of this public zone illustrating the function of not being heard and seen when in the private zone. The use of a shared 3D digital environment for communication (DE) is also illustrated at the workPLACE. Problems of privacy intrusion while a participant is in the public zone are dealt with by abstraction of information, while still providing pertinent information. The DE also serves as a vehicle for initiating richer forms of communication, such as VMC.

videoTORSO

The videoTORSO, a flat screen that can be twisted between vertical and horizontal position by voice command, is a setup for informal everyday communication in the kitchen. The public zone is normally located in a defined area around the videoTORSO and the user must step up to it in order to be heard. But the public space could also be tracking a user who moves around in the room. The scenario in the movie shows Tony giving a voice command to the videoTORSO to establish a call and then moving up to the public zone. The presence of Christine and the daughter illustrates the semi-public zone along the kitchen fittings.

Upper left: workPLACE, upper right: videoTORSO, lower left: comTABLE, lower right: mediaSPACE.

comTABLE

The comTABLE in the kitchen contains a computer as well as a touch screen, a camera, a microphone and loudspeakers in a mobile frame at the rear end of the table. The use of this table is two-fold. In an upright position, it enables a virtual guest to participate in a dinner through VMC. Secondly it could be used for, e.g., reading a digital morning paper, or doing online ordering of groceries. By placing the camera in the frame the syntax for adjusting the comZONE becomes clear – fold up the display for a camera view around the table – fold down the display, and the camera will be turned off, although the image appears, as shown in the scenario.

mediaSPACE

The mediaSPACE in the living room extends the physical room by connecting to a distant space, presented on two 80–inch screens, mounted side by side and seamlessly integrated into one wall. Thus, this comZONE creates a larger social space. The mediaSPACE is primarily a public zone and is limited by a curtain on its back wall. The fact that this room is a public space when VMC is active creates a different set of problems. It becomes a challenge to both keep an overview and to provide close-ups within the scene. The video illustrates how Christine and Tony participate in a public event and then simply shut the system off for returning to their private sphere.

CONCLUSIONS

The most interesting experience from making the video is that it has been a very good complementary method in the design of the comZONES because of the focus on the user perspective. Also, writing the script helped us realize alternative, often better, ways to use the systems than the ones we had imagined during design. E.g., that the arrangement of the multiple screens in the workPLACE is a very complex issue that has to be studied further, and that more flexibility is needed, e.g., to choose portrait or landscape formats when using the videoTORSO, the comTABLE or the workPLACE.

ACKNOWLEDGEMENTS

This work has been performed in collaboration between three partners:

• CID, supported by NUTEK/Vinnova, KTH, Telia and 15 other industrial and user organisation partners, [5]

• S-lab at Telia Networks in Stockholm, whose director, Lars Lindblad, and personel have contributed with building the comHOME dwelling and paving the way for this study

• The Smart Things and Environments for Art and Daily Life Group at the Interactive Institute, whose director, Ingvar Sjöberg, has given guidance and financial support to making the video.

Emanuel Hägglund at Gluggen Production has given advice and produced the video.

REFERENCES

1. Kraut and Fish, Prospects for Videotelephony. In Finn, E.K., Sellen, A.J. and Wilbur, B.W. (Eds.) Video-Mediated Communications, New Jersey: LEA, 1997.

2. Junestrand, S. and Tollmar, K. Video Mediated Communication for Domestic Environments – Architectural and Technological Design. In Streitz, N. et al. (Eds.) CoBuild’99. Proceedings LNCS 1670, Springer, Heidelberg, Germany, 1999.

3. Hall, E. T., The Hidden Dimension, Man’s use of Space in Public and Private, The Bodley Head Ltd, London, 1966.

4. Lenman, S. 3D Digital Environments for Social Contact in Distance Work. In Electronic Proceedings of Webnet'99 World Conference, Honolulu, Hawaii, Oct. 24-30, 1999.

5. Sundblad, Y. and Lenman, S. Five Years' Experience from CID - an Interdisciplinary Competence Centre for Design of Usable IT Applications. BIT, Behaviour & Information Technology, vol. 20, no. 2, pp. 81-89, March-April 2001.

Roomware® – The Second Generation

Norbert Streitz, Thorsten Prante, Christian Müller-Tomfelde, Peter Tandler, Carsten Magerkurth

Fraunhofer Gesellschaft für angewandte Forschung e.V.
IPSI – Integrated Publication and Information Systems Institute
AMBIENTE – Workspaces of the Future
Dolivostr. 15, D- 64293 Darmstadt, Germany
+49-6151- 869 {919, 924, 863, 997}
{streitz, prante, tandler, magerkur}@ipsi.fraunhofer.de
http://ipsi.fhg.de/ambiente

ABSTRACT

In this paper, we provide background information on the video shown at UbiComp'02. This video is a revised version of the one shown at CHI'02 [3] and provides an extensive presentation of the second generation of roomware® components. Compared to the first generation, the redesign, resp. new design and implementation resulted in a comprehensive environment consisting of several different roomware components and software facilitating new forms of human computer interaction and cooperation. The second generation consisting of: DynaWall, InteracTable, CommChairs, ConnecTable, and the Passage mechanism, together with the corresponding software BEACH, PalmBeach, and MagNets.

Keywords

Roomware, human-computer interaction, gesture-based interaction, pen-based systems, CSCW, shared workspaces, ubiquitous computing, wireless networks, interactive walls, tables, and chairs, mobile devices.

INTRODUCTION

At CHI'99, we presented i-LAND - an interactive landscape for creativity and innovation [1]. It is an environment that provides a set of artifacts (roomware components) in combination with software for supporting individual as well as group work (asynchronous and synchronous cooperation) in meeting room scenarios. i-LAND was the first generation of roomware. Now, three years later, we present the second generation of roomware and software. The second generation is the result of a redesign of our previous roomware (DynaWall, InteracTable, CommChairs) and software (new BEACH user interface), and of the design of new roomware (ConnecTable) and new software (PalmBeach, MagNets). The new set of roomware components was built in cooperation with partners in the R&D consortium "Future Office Dynamics" (FOD).

FOD was founded by us in order to meet the high interest resulting from the exhibition of and reports about the first generation roomware and software. Fig. 1 shows an overview of the second generation of roomware.

Figure 1: The Second Generation of Roomware®

ROOMWARE®

We define roomware as the result of integrating information and communication technology in room elements such as doors, walls, and furniture [1,2]. It is part of our general approach that the "world around us" is the interface to information and for the cooperation of people. It requires an integrated design of real and virtual worlds. In this approach, the computer as a device disappears and is almost "invisible" but its functionality is ubiquitously available via new forms of interacting with information. Thus, the roomware approach moves beyond the limits of standard desktop environments on several dimensions.

This approach has a wide range of implications for designing the workspaces of the future and for the future role of office buildings in general as we see it changing to what we call "cooperative buildings" [1,2]. There is no space to describe the implications here in detail but we are addressing them in the presentation.

DynaWall®

The DynaWall is an interactive electronic wall, representing a touch-sensitive vertical information display and interaction device that is 4.50 m wide and 1.10 m high. The availability of sufficient display space enables teams to display and to interact with large information structures collaboratively in new ways. Two or more persons can either work individually in parallel or they share the entire display space. The size of the DynaWall opens a new dimension in human-computer interaction: information objects can be taken from one position and put elsewhere on the display („take and put") or thrown from one side to the opposite side ("shuffle") - in analogy to using physical objects - even with different accelerations.

InteracTable®

The InteracTable is an interactive table for informal group discussion and planned cooperation. It is 90 cm high with a display size of 63 cm x 110 cm. The horizontal workspace is realized with a touch-sensitive plasma-display (PDP) that is integrated into the tabletop. People can use pens and fingers for gesture-based interaction with information objects. They can create and annotate information objects that can also be shuffled and rotated in order to accommodate different view orientations that cause problems when working on horizontal interaction areas.

CommChair®

The CommChair combines the mobility and comfort of armchairs with the functionality of a pen-based computer. It has an independent power supply and is connected to all other roomware components via a wireless network. The BEACH software provides a private workspace for personal notes and a public workspace that allows moving them to other roomware components, for example to the DynaWall. Using the CommChair, one can interact remotely with all objects displayed on the DynaWall.

ConnecTable®

The ConnecTable is a modular version of the CommChair but with new functionality. It is a mobile, networked and context-aware information appliance that provides affordances for pen-based individual and cooperative work. In order to dynamically enlarge an interaction area for the purpose of shared use and vice versa, a flexible coupling of displays has been realized that overcomes the restrictions of display sizes and borders. Two ConnecTable displays dynamically form a homogenous display area when moved close to each other. The appropriate triggering signal comes from built-in sensors allowing users to temporally combine their individual workspaces to a larger shared one by a simple physical movement in space. No additional login is needed. Connected ConnecTables allow their users to work in parallel on an ad-hoc created shared workspace as well as exchanging information by simply shuffling it from one display to the other. Correct view perspectives are facilitated in the same way as at the InteracTable.

Passage

The "Passage" mechanism provides an intuitive way for the physical transportation of virtual information structures using arbitrary physical objects, so called "Passengers". The assignment is done via a simple gesture moving the information object to (and for retrieval from) the "virtual" part of the so called "Bridge" that is activated by placing the Passenger object on the physical part of the Bridge. No electronic tagging is needed. Passengers can be viewed as "physical bookmarks" connected to the virtual world.

SOFTWARE

The innovative functionality of the roomware components results from the combination of hardware features and software features enabling new types of interactions and cooperation.

The cooperative hypermedia environment **BEACH** provides new forms of very intuitive human-computer interaction based on using only fingers and pens and new ways of cooperative sharing for multiple device interaction. It provides a modeless user-interface allowing to scribble and to gesture (for commands) without having to switch modes. The incremental gesture recognition detects the type of input and provides feedback via different colors or sounds. Although the DynaWall is composed by three separate displays, people can work cooperatively on it because BEACH makes the physical borders disappear by providing one seamless shared workspace.

MagNets is a special application of BEACH providing a creativity tool where information objects are represented by cards that are "magnetic" in the sense that they are repelling or attracting in order to structure networks of ideas.

PalmBeach supports mobile work on the road by providing a card-based creativity tool on personal digital assistants (PDAs). Information created on the PDA can be intuitively "beamed" to the DynaWall thus integrating small handheld devices into our suite of roomware components.

ACKNOWLEDGMENTS

We thank our partners, in the R&D consortium "Future Office Dynamics" (FOD) (www.future-office.de), especially Wilkhahn and Wiege, for their substantial contributions to the development of the second generation of roomware components. We also thank all students of the AMBIENTE division (www.darmstadt.gmd.de/ambiente) for their contributions to the implementation of the software.

REFERENCES

1. Streitz, N. et al. (1999). i-LAND: an interactive landscape for creativity and innovation. In *Proceedings of CHI'99* (Pittsburgh, May 15-20, 1999). ACM Press, 120-127.
2. Streitz, N., Tandler, P., Müller-Tomfelde, C., Konomi, S. (2001). Roomware: Towards the Next Generation of Human-Computer Interaction based on an Integrated Design of Real and Virtual Worlds. In: J. Carroll (Ed.), *Human-Computer Interaction in the New Millennium.* Addison-Wesley, 553-578.
3. Streitz, N., Prante, T., Müller-Tomfelde, C., Tandler, P., Magerkurth, C. Roomware: The Second Generation. In: *Video Proceedings and Extended Abstracts of ACM CHI'02.* ACM Press, New York, NY, 506-507, 2002.

The Inspiration Watch: An Augmented Shopping Device

Rachel Murphy
Hewlett Packard Laboratories Bristol
Filton Road, Stoke Gifford,
Bristol BS34 8QZ, UK
racmur@hplb.hpl.hp.com

Abigail Sellen
Hewlett Packard Laboratories Bristol
Filton Road, Stoke Gifford,
Bristol BS34 8QZ, UK
abisel@hplb.hpl.hp.com

ABSTRACT

The "Inspiration Watch" is a conceptual design for a wirelessly connected watch that displays images on a small, colour LCD screen. The watch is intended to aid shoppers in deciding on gifts to buy for friends and family. The concept was inspired by observations made during a study of consumer's shopping behaviour. After generating an explanatory animation of the concept, a short study was conducted to gain user feedback.

KEYWORDS

Wish-lists, shopping, images, display, consumers.

INTRODUCTION

Conventional approaches to gift buying often require consumers either to ask for ideas from the person that they are buying for, or to ask a friend or family member for suggestions. These approaches can have shortcomings in terms of getting it wrong or "spoiling the surprise". Alternatively, gift buying can be aided by accessing on-line "wish-lists" or using wedding list type services to find out exactly what people want. However, these methods are usually only linked to one shop or support one type of gift buying. They also leave little to the shopper's own imagination in thinking through gift ideas.

In this paper, we develop an alternative approach to gift buying which attempts to address these issues. In a recent investigation [1], we explored the value of shopping in physical shops compared with online shopping. For example, we found that people quite naturally browsed shops to get ideas of what to buy a person as a gift:

"… I might ask people what they want for Christmas. Well.. I wouldn't ask them directly. I might ask another member of my family what they want. Other than that I just spend a day or something having a look around…"

A quotation from another interview in the same study helped describe these thought processes in a bit more detail:

"Well at Christmas in my head I would have a list of people and I would be thinking about possible things for each person and then I might just see something and I might think oh I'll buy that for that person if it was something I'd think they'd like".

Taken together, these insights helped highlight an area that we felt would benefit from a technological solution

CONCEPT

The aim of this new design idea was to exploit the accessibility, adaptability and storage capabilities of the Internet, whilst recognizing the good reasons why people still enjoy shopping within a physical environment. The "Inspiration Watch" is a wirelessly connected wearable device that allows the wearer to view visual images representing the interests of friends and family. This is intended to inspire the wearer to think of appropriate gifts to buy whilst shopping. Images could also be accompanied by appropriate text, sounds, or meta-data, to provide richer descriptions.

The idea behind this concept is that shoppers access friend's or family's information from an online image database similar to HP's "cartogra.com" photo archive website [2]. An image library for the Inspiration Watch allows people to set up a personal database by either letting them place their own pictures in the database or by allowing them to select images they like from an existing database. Adding to one's own image database could be done through the Inspiration Watch itself. A wearer would be presented with a series of random images. The wearer clicks on the images he likes, adding them to his own database. This might be easily done whilst on the move, taking advantage of "dead-time" (for example time wasted waiting for a bus or standing in a queue).

Gift buyers are then able to access each individual's image database through their own Inspiration Watches by clicking through photos of their friends and family, selecting the image of the chosen person and scrolling through their set of images.

EVALUATION OF CONCEPT

In order to carry out some initial evaluation of this concept, a short photo animation was generated using Macromedia Director to help communicate the idea of the Inspiration Watch (Figure 1). This movie was shown to 12 participants: 4 women and 8 men ranging in age from 20 to 50 years old. After showing the movie, participants were asked a series of questions about what aspects they liked or didn't like about the concept, what they would change and how useful they might find such a device. Responses were recorded and main findings were summarised.

Figure 1. Still image from the Inspiration Watch animation

Results

The concept appealed in general terms both to men and women. Men in particular were enthusiastic due to the difficulties they said they had with gift buying. Many different suggestions were made for improving the concept, however. The main comments were as follows:

Adding a Reminding Feature

The most frequent comment was that people said they would like to be reminded of forthcoming special events for which gift buying would be necessary. For example, they wanted to be reminded when a birthday was coming up by an alert on their watch:

"You could customise it for a birthday kind of thing or have some sort of alarm system [telling you] you need to buy a birthday present for this person or a card for that person."

One suggestion was that a person's picture would appear on the watch a few days before her birthday (or other event). This would allow the wearer to go directly to that person's image database and plan a shopping trip by selecting that picture.

Enhanced Browsing

In terms of browsing through people's image databases, some participants said they wanted images not at random, but rather categorised in a simple way: i.e. jewellery, animals, hobbies or holidays. They felt this would be valuable especially for navigating through large numbers of images:

"If you happen to be in a department store or happen to be in a record shop, what type of music… a drill down menu I suppose, would be useful"

However, other participants were averse to any changes that would complicate the interface for the most simple form of the watch idea.

Form Factor

The watch form factor received a mixed reaction. For example, some people did not routinely wear a watch and wanted access through a mobile phone instead. In fact, almost all participants felt that a mobile phone would be a more convenient, easy, and realistic way to access an image database whilst shopping.

Additional Information

When asked what additional information they would like to have access to on the watch, almost all participants felt that text would give deeper meaning to the images. Also, in support of the mobile phone form factor, adding text to a picture via a mobile phone would be easier than adding text via a watch. Suggestions were also made about having access to information such as what a person might *not* want, what a person already has, what the gift buyer might have bought this person in the past, and other relevant details such as the size of clothing the person wears. A suggestion was also made about the value of the gift buyer being able to input images into other people's databases as ideas occurred to them.

Casual Capture

A final issue that was raised was in relation to the creation of image databases and whether or not people would be happy to spend the time setting one up. One solution to this could be by having an embedded camera in the viewing device that would allow users to capture images of things that they like. These images could then be sent directly to their image database.

CONCLUSION

This paper has shown how animation of new design concepts can be used as a quick way of gaining user feedback about ideas still in their initial stages of conceptualization. In this case, this idea points to ways in which new and existing mobile devices and services (such as the Nokia 9210 colour screen GPRS mobile phone and third-generation cell phone networks) might be used for a new class of application to augment "bricks and mortar" shopping.

Next steps for this project are to create a number of inspiration image databases using real image databases created by users. These can then be tested in a real shopping environment to find out exactly how useful they actually could be.

ACKNOWLEDGMENTS

Many thanks to Kate Shaw who conducted the initial shopping study, and to Sarah Beech who starred in the animation.

REFERENCES

1. Shaw, K. & Murphy, R. (2001). *Dot.Com shopping in the real world: Opportunities for integrating digital and physical marketplaces*. HP Tech Report.
2. *http://www.cartogra.com*

UbiControl: Providing New and Easy Ways to Interact with Various Consumer Devices*

Matthias Ringwald
Distributed Systems Group,
Department of Computer Science,
Swiss Federal Institute of Technology (ETH) Zurich,
8092 Zurich, Switzerland
mringwal@inf.ethz.ch

ABSTRACT
This paper describes a system for using a PDA to easily control consumer devices found in our environment. By attaching a laser-pointer to the PDA, the user can simply point to a device in sight and request a user interface description. The user can then control the selected device in a web browser like fashion that facilitates spontaneous interaction with consumer devices.

Keywords
Interaction, consumer devices, user interface description, virtual counterparts, home-automation, X10, Slink-e, selection by pointing, laser-pointer, solar panel

INTRODUCTION
We assume that most consumer devices like TV sets, VCRs and stereos will still exist as separate devices in the near future. Each of those devices comes with its very own remote control that in principle requires reading the manual or spending some time to figure out how to operate the device. A standard PDA could often provide a better and user-friendly interface to those devices [3].

In this paper we want to show how a PDA can be used to interact with consumer devices in our environments as depicted in Figure 1. The user points with an attached laser-pointer to a device she wants to control. By doing this, the device is selected for further interaction.

The concept of device selection by pointing was proposed before in [2]. We used that concept to build a home-automation system that allows the integration of various different consumer devices that are controlled over a central server through a PDA. The real world device the user points to with the laser-pointer is used as a reference for the server to know which device to control.

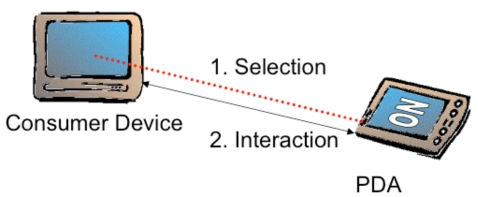

Figure 1: Selection and interaction with a consumer device by use of a PDA

* The work presented in this paper was supported (in part) by the National Competence Center in Research on Mobile Information and Communication Systems (NCCR-MICS), a center supported by the Swiss National Science Foundation under grant number 5005-67322.

SELECTION
Interaction with a previously unknown device requires finding this device first. This might seem simple at first, but much has been done in recent years to provide a middleware that allows device discovery like Jini [8] or Salutation [7]. In those systems, devices have to register with a central lookup service to allow other devices to be found.

When the user enters a new room, her PDA could register with the discovery system and get a list of all available devices that provide some kind of service. But for the user to control any of them, she mentally needs to connect the physical device to its virtual representation. This is not an easy task as the list doesn't contain any location information, so even finding out what devices are in the same room as the user is usually not possible. Even if the physical location is known to the devices, and the list of available devices could then be reduced to all devices in the same room, this approach still leads to problems if there are multiple instances of the same type (e.g. two TV sets) in one room, requiring the user to figure out what symbolic name belongs to which device to access it through the PDA.

Using a laser-pointer to point to the surface of a device is a straightforward and out-of-band solution to the device discovery problem, allowing us to do without a discovery service.

USER INTERFACE
To allow interaction with new consumer devices without any previous setup of the PDA, a user interface description for every device is stored on the server. After device selection through the laser-pointer, the interface description is downloaded to the PDA.

Much work has been done on how to specify a user interface that abstracts from the presentation device [1]. Common to those is the separation of functionality and presentation.

For the user interface description, we have used a structured text file to describe the attributes of the actual device and its control widgets to get started. The attributes for a TV set could be power mode, volume, current channel number and current channel name. The control widget should then be a power button, channel number label, channel name label, volume plus/minus buttons and channel plus/minus buttons. In those files, the layout is fixed for a PDA form-factor. The graphical user interface is rendered using the current state of the device on the PDA.

Interaction with the device happens in a web-application-like fashion. If the user presses a button on the PDA, a request with the encoded command is sent to the server. This command is processed and a result is sent back to the PDA that triggers an update of the displayed interface. The current state of the device is always displayed on the PDA as an indirect feedback to the user action.

Figure 2: The Visor prototype

INTEGRATION OF EXISTENT CONSUMER DEVICES

As most consumer devices do not have a network connection, we chose a server based approach to control them using existent home-automation systems. This allows the PDA to communicate with the server instead of the real devices.

To abstract from the used home automations system, we used virtual counterparts for all consumer devices. Those virtual counterparts live on the server. As the real devices do not provide any status information by themselves, we used the virtual counterparts also to keep track of the real device's state and to integrate some control logic. For example, a TV set counterpart could support symbolic station names in addition to the actual channel number.

THE IMPLEMENTATION

For the control PDA we used a Handspring Visor and added a Xircom 802.11 card for wireless connectivity. A standard 802.11 base station was used in infrastructure mode to provide TCP/IP access to the Visor. We used HTTP for the communication between the PDA and the server. The laser-pointer is connected over a one transistor driver to the serial transmit pin on the HotSync connector. The power supply of the 802.11 card was used to power the laser-pointer. Figure 2 shows our extended PDA.

We used a Dell PC with a 800 MHz Pentium 3 processor and 256 MB memory running Windows 2000 for the server. To control the consumer devices, we connected the PC to a X-10 system [9] and a Slink-E device [4]. The X-10 system is a home-automation system that is able to send on/off commands to a device over the power-line. We used it to control the lighting. The Slink-e device can record and play back infrared signals like those used for consumer devices like TV sets, VCRs, DVD players and stereos, so we installed infrared transmitters in all rooms to control those devices.

To allow the selection of the consumer devices, we attached a solar panel as a receiver for the laser beam to every device and connected them to the server.

For the virtual counterparts, we used the Context Toolkit [6] and implemented them as Context Widgets.

For a more detailed version of the UbiControl system see [5].

DISCUSSION AND FUTURE WORK

We have shown that a PDA with an attached laser-pointer is well suited to select and control existent consumer devices. To achieve this functionality we used a server and provided virtual counterparts for the consumer devices. Those counterparts run on the server and can control the real devices over several proprietary home-automation systems. We further demonstrated a way to utilize a user interface markup language to provide a user interface on the mobile device.

As this approach requires a server and the wiring of those devices, we will investigate ways to build a distributed system wherein the PDA can communicate directly with the devices and no server installation is necessary.

ACKNOWLEDGMENTS

The video was recorded with the help and the narration of Peter Jensen from Georgia Institute of Technology, Atlanta, USA and was edited by Firas Awais of 4as Grafix, Magstadt, Germany.

REFERENCES

1. Marc Abrams, Constantinos Phanouriou, Alan L. Batongbacal, Stephen M. Williams und Jonathan E. Shuster. UIML: An Appliance-Independent XML User Interface Language. *WWW8/Computer Networks 31(11-16)*, 1999, p. 1695–1708.
2. Michael Beigl. Point & Click - Interaction in Smart Environments. *In Lecture Notes in Computer Science*, Volume 1707, Springer-Verlag, 1999, p. 311–314.
3. Jeffrey W. Nichols. Using Handhelds as Controls for Everyday Appliances: A Paper Prototype Study. *In ACM CHI'2001 Student Posters,* 2001, p. 443–444.
4. Nirvis Systems Inc. Slink-e home automation device. http://www.nirvis.com/, 1999.
5. Matthias Ringwald. Kontrolle von Alltags-gegenständen mit einem PDA. Diploma thesis, University of Karlsruhe, November 2001. http://www.inf.ethz.ch/ mringwal/publ/diplom.pdf.
6. Daniel Salber, Anind K. Dey und Gregory D. Abowd. The Context Toolkit: Aiding the Development of Context-Enabled Applications. *In ACM CHI 1999*, 1999, p. 434–441.
7. Salutation Consortium. Salutation. http://www.salutation.org/, 2002.
8. Sun Microsystems, Inc. Jini Connection Technology. http://www.sun.com/jini/, 2000.
9. X-10.ORG. X-10 Technology and Resource Forum. http://www.x10.org/, 1997.

POSTER REVIEWERS

Stavros Antifakos, ETH Zürich
Michael Beigl, TecO, University of Karlsruhe
Staffan Björk, PLAY, Interactive Institute & Chalmers University of Technology
Christian Decker, TecO, University of Karlsruhe
Jennica Falk, Medialab Europe
Lalya Gaye, Future Applications Lab, Viktoria Institute
Jussi Holopainen, Nokia Research Center
Jan Humble, Swedish Institute of Computer Science
Henrik Jernström, PLAY, Interactive Institute
Magnus Johansson, IT University of Göteborg
Rikard Lindgren, Telematics Group, Viktoria Institute
Sara Ljungblad, Future Applications Lab, Viktoria Institute
Peter Ljungstrand, PLAY, Interactive Institute & Chalmers University of Technology
Sus Lundgren, PLAY, Interactive Institute
Peter Lönnqvist, FUSE Group, KTH
Florian Michahelles, ETH Zürich
Magnus Nilsson, IT University of Göteborg
Claudio Pinhanez, Pervasing Computing Group, IBM T.J. Watson Research Center
Johan Redström, PLAY, Interactive Institute & Chalmers University of Technology
Dimitris Riggas, Computer Technology Institute
Matthias Ringwald, ETH Zürich
Tobias Rydenhag, IT University of Göteborg
Johan Sanneblad, Future Applications Lab, Viktoria Institute
Martin Strohbach, University of Lancaster
Tobias Zimmer, TecO, University of Karlsruhe
Karl-Petter Åkesson, Swedish Institute of Computer Science

VIDEO JURY

Marcus Bernander, Allgon AB
Staffan Björk, PLAY, Interactive Institute & Chalmers University of Technology
Lalya Gaye, Future Applications Lab, Viktoria Institute
Lars Erik Holmquist, Future Applications Lab, Viktoria Institute
Maria Håkansson, Future Applications Lab, Viktoria Institute
Peter Ljungstrand, PLAY, Interactive Institute & Chalmers University of Technology
Magnus Nilsson, IT University of Göteborg
Tobias Rydenhag, IT University of Göteborg

AUTHOR INDEX

UBICOMP2002