Privacy, Security and Trust within the Context of Pervasive Computing

edited by

Philip Robinson
Harald Vogt
Waleed Wagealla

Privacy, Security and Trust within the Context of Pervasive Computing

THE KLUWER INTERNATIONAL SERIES IN ENGINEERING AND COMPUTER SCIENCE

Privacy, Security and Trust within the Context of Pervasive Computing

edited by

Philip Robinson
University of Karlsruhe, Germany

Harald Vogt
ETH Zürich, Switzerland

Waleed Wagealla
University of Strathclyde in Glasgow, UK

Springer

eBook ISBN: 0-387-23462-4
Print ISBN: 0-387-23461-6

Created in the United States of America

Visit Springer's eBookstore at: http://ebooks.springerlink.com
and the Springer Global Website Online at: http://www.springeronline.com

Contents

Preface

Pervasive Computing is sometimes labeled as another passing "technology hype", while some people in society admit fear of the possibilities when computers are integrated into our everyday lives. Researchers are busily investigating solutions to the security requirements identified by businesses and consumers, with respect to confidentiality, privacy, digital rights maintenance and reliability of information systems.

The question of trustworthiness of spontaneously invoked interactions between devices as well as of exchanges with previously unknown human principals and with entities from unknown organizations or domains has also been raised. Furthermore, sensor networks and powerful embedded computers facilitate the computation of people's location, activities, conditions and other properties that would not have been immediately available to information systems in the past. While these seem like relatively disparate problems, in reality we form notional mappings between these problems and hence solutions. For example, some authors refer to trusting the context as opposed to trusting a person or thing. The assurance of security within a context has then been identified as a property in the function of trusting the context. Furthermore, people tend to exchange private information with those they trust, and within an environment where trust is somehow provable. What we believe is that an investigation of the interfaces between the notions of context, privacy, security and trust may result in deeper understanding of the "atomic" problems, but also lead to more complete understanding of the social and technical issues in pervasive computing.

The goal of the workshop was not to focus on specific, even novel mechanisms, rather on the interfaces between mechanisms in different technical and social problem spaces. 21 people from different parts of the world took part in the one-day discussion, including PhD students, seasoned and junior researchers.

This workshop promises to be a lasting experience and we encourage researchers to participate in future events. We hope that you will find its proceedings useful and valuable to read.

August 2004

Philip Robinson, Harald Vogt, Waleed Wagealla
Workshop Co-chairs SPPC 2004

Acknowledgments

We would like to thank all authors for submitting their work, and all members of the Program Committee, listed below, for their cooperation and time spent reviewing submissions. Finally, we thank Kelvin Institute, Microsoft Research, and SAP for financially supporting the publication of proceedings for the SPPC workshop 2004.

Program Committee

I

THE INFLUENCE OF CONTEXT ON PRIVACY, TRUST AND SECURITY

This page intentionally left blank

SOME RESEARCH CHALLENGES IN PERVASIVE COMPUTING

Philip Robinson[1], Harald Vogt[2], Waleed Wagealla[3]

[1] *Telecooperation Office, University of Karlsruhe, Germany*
philip@teco.edu

[2] *Department of Computer Science, ETH Zürich, Switzerland*
vogt@inf.ethz.ch

[3] *Department of Computer and Information Sciences, University of Strathclyde in Glasgow, UK*
waleed.wagealla@cis.strath.ac.uk

Abstract The topics of privacy, security and trust have become high priority topics in the research agenda of pervasive computing. Recent publications have suggested that there is or at least needs to be a relationship of research in these areas with activities in context awareness. The approach of the workshop, on which this proceedings reports, was to investigate the possible interfaces between these different research strands in pervasive computing and to define how their concepts may interoperate. This first article is therefore the introduction and overview of the workshop, providing some background on pervasive computing and its challenges.

1. Introduction

We are currently experiencing a bridging of human-centered, socially oriented security concerns with the technical protection mechanisms developed for computer devices, data and networks. The foundations of this bridge started with the Internet as people, both purposely and accidentally, provided gateways to their personal computers and hence information. With enterprise-scale and even personal firewalls, providing a rule-controlled entry point into network domains, as well as cryptographic means of ensuring secrecy, many attacks on computer applications and data were circumvented, given that people behind the virtual

walls adhered to policy. Pervasive computing however moves these resources from behind these centrally configured virtual walls, allowing mobility, distribution and dynamic interconnection, in order to support more advanced services and modes of usage. Living in a world where the walls, cars, stores, clothing and cafés are automatically aware of the context and hence needs of owners, users and (potential) patrons, due to embedded computers, sensors and advanced networking, can be sometimes intriguing; yet on other occasions society questions the state of their privacy, becoming insecure and untrusting with respect to technology.

On April 20th 2004, as part of the Pervasive Conference in Vienna Austria, about 21 international researchers and technologists came together to discuss this matter. Rather than looking at specific pervasive computing technology or security mechanisms, the goal was to gain an understanding of the relationships between context-awareness, privacy, security and trust, as these are the nuts and bolts that hold the society-technology bridge in place. By way of introduction, the publication begins with a brief overview of the State of the Art in Pervasive computing, in order that the motivations of the workshop are better understood. The workshop's themes and motivations are discussed in section 3, while section 4 provides an outline of the results of this workshop.

2. The State of the Art in Pervasive Computing

The term "Pervasive Computing" emerged from research at IBM during 1996 - 97, embracing the vision of computing services available anytime, anywhere and on demand [10]. Advances in global and mobile wireless technologies, giving new meaning to electronic business, remote workers and collaborative enterprises, motivated this. This is reflected in the current wave of standardization activities surrounding Web Services, where enterprises open-up their computing infrastructure at the service level and provide remote interfaces. Five years earlier, Mark Weiser at Xerox PARC was leading research labeled as "Ubiquitous Computing (UbiComp)", and expressed its concepts in his 1991 paper: "The Computer for the 21st Century" [13]. UbiComp's initial focus was not on making infrastructure available everywhere but preached ubiquity as a notion similar to the availability of natural resources and utilities such as electricity, water and air. Today we are noticing a convergence of themes such that the technical infrastructure advancement principles of Pervasive Computing complement the user centric opinions of the UbiComp community. The major difference in philosophies has been that Pervasive Computing was started with the initiative to exploit the existing wide-

scale deployment of computing technology, while UbiComp's initiatives were to effectively make this complex mass of technology transparent to the human user's, especially those with limited technical "know-how".

For the purposes of the workshop themes and this publication, we consider Pervasive Computing to be comprised of five research areas - mobile computing, wireless networking, embedded computing, context awareness with sensor technology, and human computer interaction (HCI). An overview of these is given below, including the context within which they were discussed during the workshop. There are additional terms that may contribute to the vision of Pervasive Computing, but we have selected the ones with which we have most often encountered during workshops, conferences, seminars or discussions with other researchers in the field. In addition, other terms tend to be an overlap of these five themes e.g. "Wearable Computing is an overlap of Mobile Computing, Embedded Computing and HCI". "Nomadic Computing" is an overlap of Mobile Computing and Wireless Networking.

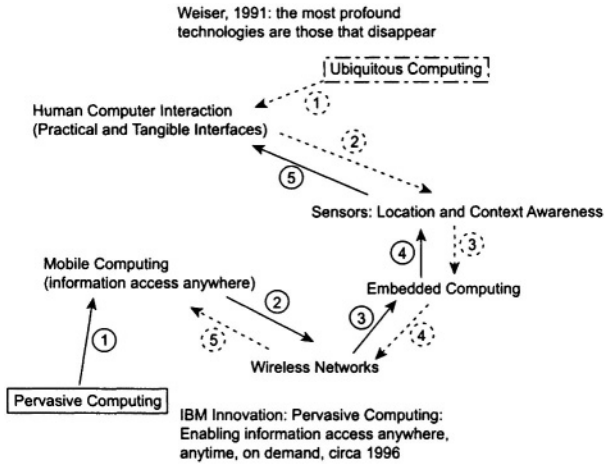

Figure 1. Advances in both Pervasive and Ubiquitous Computing (UbiComp) show a convergence of the communities. UbiComp was initiated with a user-centric methodology, while Pervasive was based on a bottom-up strategy for exploiting technology

We therefore consider Pervasive Computing to embrace the five areas of research stated in Figure 1 above. There are additional terms that may contribute to this vision, but we have selected the ones with which we have most often encountered during workshops, conferences, seminars or discussions with other researchers in the field.

2.1 Mobile Computing

Mobile Computing allows people to be on the move and still continue working with their familiar user interface and applications. Initially this meant carrying a large case, heavy yet lower-quality monitor and a large battery source. However, today's PDAs (Personal Digital Assistants), Laptop Computers and even some Mobile Phones are capable of supporting the basic applications that users need - word processing, communications, timetable, calculator, address book and so on. Display, Microprocessor, Ergonomic, Energy and Material research have all contributed to what we refer to as a mobile computer today. Other phrases that refer to mobile computing are Nomadic Computing, where the term "Nomad" implies no real fixed place of abode, and Wearable Computing, where the feedback and control interfaces of the computing devices are built-in to the garments of the user. For example, spectacles become displays, the CPU (Central Processing Unit) is the size and form factor of a Walkman, and a T-shirt becomes a router in a personal network [8]. These small, luggable, concealable and wearable computers have however been the targets of theft, such that individuals and companies have suffered loss of expensive equipment and, moreover important, sometimes sensitive information.

2.2 Wireless Networks

Wires tend to be intrusive, as they require planning and coordination during installation, alterations in the aesthetics of the environment and hinder versatile movement. For these reasons, wireless protocols have been developed to support long-range (e.g. GSM, GPRS), local-area (e.g. IEEE 802.11), and short-range (e.g. IrDA, Bluetooth) communications. Along with the nature of the data this imposes differences in the security requirements for applications that employ these protocols. The issues with security in wireless environments are well known, as the medium is generally more widespread, shared and it offers many more points of contact. Wireless networks are therefore more prone to eavesdropping and other malicious attacks because of these characteristics.

2.3 Embedded Computing

Embedded computers are small, typically single-purposed (as opposed to general purpose) machines that are built-in to larger systems, devices or objects. The particular function that they perform must be done without having the concerns of scheduling and preemption that would be the case in multitasking operating systems. Embedded computers

may have their own power supply, memory, custom OS, and network interfaces. Embedded computing has been considered as contributory to Pervasive Computing, while many Pervasive systems are built by creating a distributed network of micro nodes each with a special purpose. There is still a need however to coordinate and make sense of the interaction between these small computers by a more powerful system. However, as these embedded systems are so small and resource-limited, they do not support large-scale crypto protocols. Nevertheless, they may store data fragments that may be reconstructed by any system capable of coordinating their interaction. There is therefore some concern that Pervasive Computing systems may ignore privacy, security and trust requirements at the very low level, either because it is too complex or technically infeasible.

2.4 Context Awareness with Sensor Technology

One of the more significant contributions of Pervasive and Ubiquitous Computing has been the work in the area of location and context awareness. Research in this area suggests that computer systems need to be more informed about their environment and that of their users, in order to enhance their performance and manner in which they provide computational services. The way this is done is by having various sensors distributed in the environment, including temperature, light intensity, movement and location, and then aggregating the information from these sensors to produce some representative value of the situation. The computer systems that receive this situation data can then adapt in order to better serve the circumstance. For example, if there are many people congregating outside of an empty meeting room, the computer system that automatically administers this meeting rooms may be enabled to sense the situation and try to appropriately prepare the environment for such a meeting. The major issue with these smart, sensing and adaptive environments is the degree of personal information to which they require access. This may be obtained from the RFID (Radio Frequency Identification) tags the people are wearing or some form of tracking system. While the users enjoy the benefits, they may remain incognizant of ensuing threats to their privacy by other parties also tapping into their situation traffic.

2.5 Human Computer Interaction (HCI)

HCI research has been recognized for more than a decade now, however, it was initially focused on the selected placement and font of text, as well as the rendering of graphics and widgets on a graphical user in-

terface in a manner that matched the human user's perception of what these objects should represent. Today HCI has moved beyond the computer screen and back into the real world, where computer interfaces are being realized by manipulation of directly physically graspable objects [5]. Moreover, it can be understood that the digital media is being captured in the form of physical objects. This therefore suggests that the availability and controllability of digital information must be reflective of how the associated physical objects are handled and managed.

2.6 A Pervasive Computing Environment

Having defined the major contributing themes to Pervasive Computing, in this section we propose a model that moulds these themes together and provides a single architecture for a "Pervasive Computing Environment". It is a five-layered model representing different levels of computational abstraction from the perspective of the human. The top layer is referred to as the "physical layer", as this comprises the physical artifacts, affordances and norms with which a human user is inherently familiar. With HCI in mind, the goal of is that the human need only be concerned with the handling and resultant feedback of the physical layer. That is, the human may or may not be aware of the reception of a computational service, but is aware of changes in state of physical objects with which he or she interacts. The second and third layers are for translation between the physical and computational layers of the model. The second layer is called the "Perceptive layer", while the third is called the "Analog/Digital conversion layer". The Perceptive layer is composed of sensors (for taking input from the physical layer) and actuators (for providing output to the physical layer, prompting it to actualize its state). The analog/digital layer then does the concentrated task of converting between analog and digital signals, such that there is comprehension between the real world and the so-called "virtual world". We have also decomposed the computation and communications layers into primary and secondary functions. The primary functions of computation and communication are those concerned with the coordination functionality of the environment - such as communication protocols and operating systems. The secondary functionality is the actual applications that are implemented within the environment - these would include Office-ware, Meeting Rooms, Smart Homes and others that already exist on the market or are still in development. Orthogonal to each layer is a "Utilities" component. This represents the power and administration required to drive and manage the operations of constituents of each layer. The utilities component is therefore particularly sensitive when considering that

attacks that compromise the utilities of an environment typically make the system unavailable, unless the appropriate back-up mechanism is implemented.

We suggest that this model can be used as a generic reference when discussing any form of pervasive computing environment. Examples include Smart Spaces [11], Adaptive Environments [7], Augmented Worlds [9] and Ambient Spaces [1]. These are all specializations of the model, depicted in figure 2, where the constituents of each layer may be configured to meet the particular requirements of the system environment.

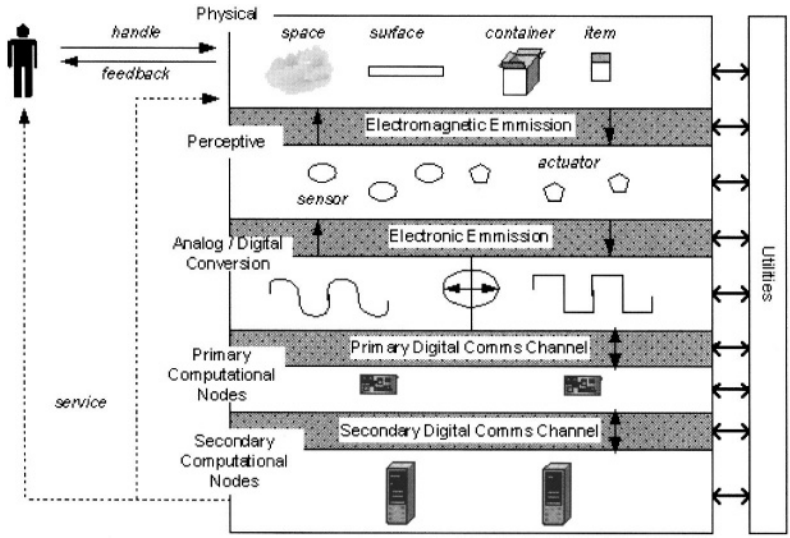

Figure 2. Depiction of a Pervasive Computing Environment

When considering context-awareness, privacy, security and trust, it is recognized that these have implications for and dependencies on each of these layers. Context awareness cannot function if the infrastructure for perception, conversion and computation does not dependably function. Dependability is a property of trust, and it is an assumption upon which many security and privacy systems are based. There are of course systems of adaptation that propose compensation measures for loss or lack of utility or computational power, which may count towards a higher assurance of dependability. Although dependability is and was not a central focus of the workshop, it has aided in motivating the themes addressed.

3. Workshop Themes and Motivation

The motivation for this workshop was derived from consideration of everyday situations. For example, when someone asks to momentarily use an office space, what goes through the mind of the owner? The owner may be concerned that this arbitrary person may make an overseas call and therefore leave an unwanted expense behind. Additionally, this person may browse high profile or confidential documents lying on the table or even look at the numbers stored in the phone. From an even more retrospective standpoint, the owner may have concerns about why their office was selected and how the inquirer gained the knowledge to support this decision. As the reality of pervasive computing becomes more and more apparent, these requests become more subtle, frequent and potentially impacting. Even if one concurs that this is a case of extreme paranoia, it is not easy to comprehensively reason about these concerns.

Consider the future. Devices embedded in the smart environments and worn on our bodies will communicate seamlessly about any number of different things. In such kind of interactions, huge amounts of information will be shared and exchanged. Even though they may be the means of enjoying context-based and other advanced services, there is an increased risk involved in some of these interactions and collaborations, if collaborators are about to use our private possessions. Questions naturally arise: do you want this information shared? How can you trust the technology - yours and the environment's? What does the environment itself do, and how can you secure the access to private information, even though you may want to share it in certain contexts? This further illustrates how combined assessment of the interrelationships between trust, security, privacy and context aid in confident decision making. In every-day life we do not treat these concerns in isolation; we actually make spontaneous decisions that are based on maintaining a "comfortable" balance. Even though we do not completely understand these basic building blocks, the potential trade-offs are intuitively understood.

3.1 Context Awareness

Dey defines "context" as: any information that can be used to characterize the situation of an entity. An entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and applications themselves [4]. We have adopted this definition but first some of the terms need to be clarified for appropriate use in a context where security, trust and privacy are important. For example, the term "information" is very broad, but we wish

to refer to this as "evidence", which has a stronger semantic affiliation i.e. supports an argument. This also stresses the urgency that context is something that may have to be proved in some situations. Therefore, to "characterize a situation" implies that we are supporting arguments for the conditions of the situation. This implies that there must be some premises or rules used to come to these conclusions. Additionally, the terms "entity" and "user" always require further clarification. We therefore want to stick to the terms "subject" and "target", without making any assumptions about their nature i.e. physical or electronic. Therefore the sentence "... between a user and an application" would be simply replaced with "... between a subject and a target". The term "relevant" is also ambiguous, based on assumptions and is subjective. We therefore strike it from our definition as we deem that context-awareness should be a pursuit of facts.

Our definition would therefore read as: *Context is any evidence that can be used to support arguments for the conditions of the situation of any subject or target, which influences their interactive behaviour.*

Privacy, security and trust may hence be representative of the rules that influence the interactive behaviour between a subject and a target, or the post-assertion of the validity of the interaction and resultant context. Context is therefore the knowledgebase that supports the reliable derivation of meaning in an environment, while context-awareness is the ability of an entity to adapt to changing "meanings" of information.

3.2 Privacy

Technical solutions to the privacy problems in ubiquitous computing cannot stand on their own to protect our privacy rights. Privacy protection has always been the subject of legislation, since there is an inherent conflict in service provisioning: personal data must be collected in order to adapt services to the users' needs and preferences, but once given away, there is no technical procedure to revoke it or detain somebody from passing it on. Technology makes collecting data easy but cannot help protecting it against abuse. Thus traditionally, solutions rely on binding the collector of personal data by law to certain procedures, for example obfuscation (by anonymizing the collected data) or deletion after a certain time period.

However, data collectors must be enabled to meet the standards set by jurisdiction and market forces, and technology can help in this regard. This potentially leads to systems that are both easy to implement, and therefore cost efficient and widely usable, and compliant to privacy

standards. This is where a great part of privacy research in pervasive computing is aimed at.

Pervasive computing technology is often described as the ultimate tool for constant surveillance of large parts of the population, since ultimately all actions are reflected in some networked computing device, allowing putting together personal profiles in unprecedented detail and accuracy. Users might become unaware of this fact as computers become "invisible" and careless as they become unavoidable anyway. Ronald L. Rivest put it this way: "What was once forgotten is now stored forever. What was once private is now public."

Public concerns about the privacy problems of pervasive computing are nowadays preceded by the potential dangers of RFID technology, which is seen by many industries as a potential means for improving the efficiency of doing business. Object identification on an object level may be abused for creating profiles and exploiting user behaviour. While these concerns might sometimes be exaggerated, they are fundamentally valid. It seems however that the combination and ubiquity of small computing devices, wireless communication and sensors holds potential for far greater dangers to privacy to come.

3.3 Security

A system is generally called secure if there are measures taken to avoid a "bad" outcome, where the definition of bad greatly depends on the application scenario. The accepted concepts of security include availability, authenticity, authority, integrity, confidentiality and reliability, with their proportionate significance depending on the task at hand. A great deal of security mechanisms supporting these concepts have been developed, especially since the growth of the Internet, and have gained wide acceptance in military, business and consumer applications. Examples range from tamper resistant devices, cryptography and security protocols to intrusion detection systems. All these techniques will be crucial for securing pervasive computing systems, but existing incarnations are not all equally applicable. Security mechanisms for pervasive environments must be

- scalable to the small resource provisions of "invisible" computing devices,
- able to deal with devices and environments of unknown origin,
- and adaptive to the dynamics of mobile and socially motivated computing.

Developing such techniques is the challenge of research in this area. This does not dismiss the large resource of past work in cryptography, security policies and physical security. It really calls for additional methodologies for comprehending, implementing and integrating security at and between the different layers of pervasive environments.

3.4 Trust

Trust is multidisciplinary concept, which has been used in the fields of sociology, psychology, philosophy, and most recently in computing. Within these disciplines, trust is defined and treated from different angles that show its utilizations and applications. Although, there is no consensus about a definition of trust, there is a general agreement on its properties as a subjective and elusive notion. In these proceedings, contributions are concerned about the utilizations of trust in pervasive computing. The application of trust in computing is widely acknowledged by the term trust management [2]. This term has emerged as a new concept in computing, where it supports descriptions on how to facilitate trust relationships between entities. The establishment of trust enables systems to exchange information even without the intervention of administrators to authorize these interactions.

The application of trust management systems and models in pervasive computing is about how to grant users access to resources and information based on their trustworthiness rather than the application of conventional techniques that map authorizations to access rights. The view of trust management systems is that trust would be used as a measure for how much resources or what types of information are permitted or would be disclosed to others. This seems to fit the domain of pervasive computing quite well, since there is no fixed infrastructure and entities are not attached to specific domains, from which information about identities could be obtained. There are also potential interactions with huge numbers of autonomous entities, and these interactions are triggered and established in an ad-hoc manner. Therefore, to facilitate interactions in pervasive computing, trust management is considered to be the most appealing approach to reasoning about potential users' trustworthiness for granting them access to the required resources. Trust management aids in taking autonomous decisions on whom to trust and to what degree. These decisions embody reasoning about the trustworthiness and the risk involved in these interactions between entities.

To illustrate the exploitation of trust, let's consider the example of an interaction between the agents of two users (systems working on the users' behalf) that will be carried out by using their PDAs. Assume

that agent A wants to share or to get access to B's resources or stored data. The first task for B is to reason about the trustworthiness of A. This reasoning is mainly based on the accumulated trust information either from previous interactions (if there are any) or from trusted third parties (aka recommendations). There are situations, in which there is inadequate information for reasoning about trust. In this case, B would either run a very restricted risk analysis, or accept the interaction on the basis of trusting dispositional factors. However, reasoning about trust when adequate information is available is much easier in comparison to the situations of no prior information. This is why some of the proposed trust management systems incorporate solutions for uncertainty. There are some other factors that influence greatly the establishment of trust, namely contextual information about the interaction, and privacy concerns. The combination of the trust reasoning and other factors (context and privacy) will help immensely in taking decisions regarding interaction requests. This shows how trust would facilitate establishing interactions especially under the described possible complex circumstances. Therefore, trust must be balanced against other factors: users desire to participate in interactions and to share information; and users' concerns about security and privacy that would deter them from participation in interactions.

It is very clear from the above discussion that interactions are established on the basis of the individual's trustworthiness rather than a fixed security policy of access right roles. The collected evidence or information, that will be made available after the interaction is finished, would serve as solid ground for possible future reasoning and decisions. This is why trust is considered as a dynamic parameter that evolves over time.

The proposed trust management systems for pervasive computing are promising and encouraging [3, 6], but little is mentioned about implementation of these models and their validation, which would be necessary for their adoption. Moreover, the mechanisms for trust management introduced some questions about their computational cost and complexity, for which studies on techniques that help keep the overhead and complexity low, are still welcomed.

4. Outline of Proceedings

In the workshop's call for papers we posed many questions about the possible interfaces between context, security, privacy, and trust. We, as organizers and program committee members, felt that addressing the concerns of security and privacy in pervasive computing would come out clearly if interfaces were defined and considered within the proposed pro-

tocols, models, and architectures. The interfaces and their dependencies serve as a good research issues to tackle and to propose models that identify coherent solutions.

The contribution we received, in terms of submitted papers, from the workshop's participants helped in addressing and proposing solutions that would advance the developments in pervasive computing. Accordingly, the organizational of the workshop day and these proceedings are divided into four main sessions. Each one of them is devoted to the discussion about interfaces and relationships, as it has been illustrated in Figure 3. The discussion is not merely on the internal properties of individual themes, but on the properties of the interfaces from abstract view. The sessions during the workshop day were:

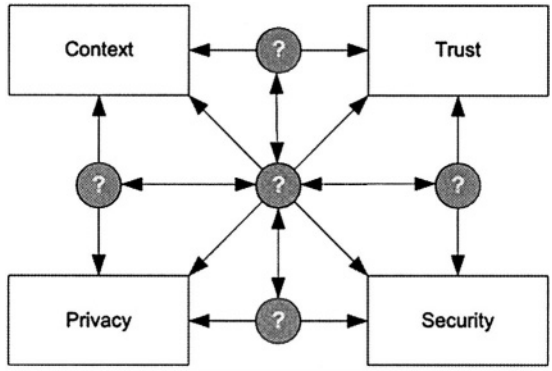

Figure 3. The view on possible interfaces between context, trust, privacy and security

1 *The Influence of Context on Privacy, Trust and Security.* The effect of context is foreseeable when discussing the concerns of security and privacy. The importance of context stem from the fact that all its information are necessary to reach a useful decision in the face of the complexity environments of pervasive computing. These decisions are essentially for granting access to resources or information and they vary according to the relevance context. Context, as parameters or information, will guide and ease the view about security, since security polices and conditions can be adjustable and contextualized. The combination of context information with systems/applications data, trust information, recognition and identity, and security policy gives a clear view of the environment. The influence of context can also be seen as adjustment/self-tuning for privacy, trust, and security, in the sense that context information determines how much information could be revealed

and to what degree/level entities will be trusted, and what types of security policies could be specified within specific context. The influence of context shows the need for defined interface in the domain of pervasive computing. The discussion on context influences raises debatable questions about: how context information would be combined with systems and applications The answers to these questions are application-specific.

2 *Secure Trust Models and Management in Pervasive Computing.* The security matters in pervasive computing are not about a mapping from authentications to access rights or privileges, but it is all about how much we trust users/infrastructure and systems. Trust expresses the level of access to resources that can be granted based on the available information and evidence. Trust information is mandatory for reaching decisions. For trust management to be effective, the contextualization of trust is an important step to build an appropriate level of trust on others. Trust management combines the notion of specifying security policy with the mechanisms for specifying security credential. To achieve that we also need to know the information about trust on the infrastructure and to express how confident we are on the authentication and entity recognition systems. Trust can prevent the loss of privacy by striking the balance between users' trustworthiness and how much could be revealed. This discussion shows clearly how trust, with the combination of context, would adjust/control privacy and security.

3 *Evidence, Authentication and Identity.* The process of authentication (authentication techniques are totally different and varies in pervasive computing) involves collecting evidence about the identity of users. The information of both trust and context are highly considered in the process of authentication, because they give an insight view into user's identity. The concerns of identity in pervasive computing are much bigger than in other applications domains, because in pervasive computing there are huge number of potential users that we may not have enough information about them. Therefore, contextual information, trust information, and evidence form the basis for the evaluation of identity and reasoning about it. An adequate level of available information and evidence will facilitate the process of authorizations. The relationship between evidence, authentication, and identity could be considered as a dependency relationship, in the sense that evidence is highly required for the process of authentication, which in turn provides valid identity.

4 *Social and Technical Approaches to Privacy Protection.* With the advances of technology, privacy solutions has to consider both technical and social approaches. This consideration is important for pervasive computing to be socially acceptable. On the other hand, both the confidentiality and integrity of the information must be controlled. The

information must be protected from unauthorized, unanticipated, or unintentional modification. Authentication is required to ensure that information has not been changed. Besides ensuring confidentiality and availability, it is also important to ensure that resources are used by granted users. To sum up, the proposed solutions should avoid the separation between technical and social factors influencing privacy.

5. Future Research Directions

The motivations for the workshop were not centered on very far-fetched and obscure scenarios. The selected scenarios were drawn from considering the state of the art in technology as well as known developments in academic and industrial research. As can be seen from the contributions to the workshop, the general approach was always to reference what happens in everyday life when humans make socially oriented decisions pertaining to privacy, security and trust, then to draw parallels with technology. The term "context" is discussed in many current theses, as it is understood to be the foundation of deriving meaning, and an understanding of meaning must exist for a decision to be made - at least one that is logically founded. Therefore as technology moves towards more intelligent systems, which can derive meaning from the world through sensing, data processing, ontology and rule execution, these systems should also be capable of making appropriate privacy, security and trust decisions. Thoughts on machine intelligence and automation often bring visions of degrading human control and the rise of the machine as a super power. However, machine intelligence and automation are not (and should not be) intent on taking the human completely out of the position of control. They should rather assist the human in making precise and reliable decisions that do not require excessive, peripheral signals and feedback from computer systems and their environment. The state of the art in privacy, security and trust, alongside the trends in technology, suggests that the balance between control and necessitated system feedback to humans needs to be found. Future research should therefore be more human-centric rather than mechanistic, such that security requirements, policies and enforcement measure can be based on (and in some cases automatically derived from) human-defined processes, situations, entitlements and relationships. Nevertheless, we need to note that pervasive computing is no longer "the future" – it is already part of our today.

References

[1] Ambient Agoras. German Fraunhofer Gesellschaft (FhG), Darmstadt. http://www.ambient-agoras.org/.

[2] Matt Blaze, Joan Feigenbaum, and Jack Lacy. Decentralized trust management. In *Proceedings of the 1996 IEEE Symposium on Security and Privacy,* pages 164–173, Los Alamitos, USA, May 1996. AT&T.

[3] V. Cahill, E. Gray, J.-M. Seigneur, C. Jensen, Y. Chen, B. Shand, N. Dimmock, A. Twigg, J. Bacon, C. English, W. Wagealla, S. Terzis, P. Nixon, G. Serugendo, C. Bryce, M. Carbone, K. Krukow, and M. Nielsen. Using Trust for Secure Collaboration in Uncertain Environments. *Pervasive Computing Magazine,* 2(3):52–61, 2003.

[4] Anind K. Dey. Understanding and Using Context. *Personal and Ubiquitous Computing Journal,* 5(1):4–7, 2001.

[5] H. Ishii and B. Ullmer. Tangible Bits: Towards Seamless Interfaces between People, Bits, and Atoms. In *Computing Systems (CHI),* pages 234–241, 1997.

[6] Lalana Kagal, Jeffrey L Undercoffer, Filip Perich, Anupam Joshi, and Tim Finin. A Security Architecture Based on Trust Management for Pervasive Computing Systems. In *Grace Hopper Celebration of Women in Computing,* October 2002.

[7] David Kirsh. Adaptive Rooms, Virtual Collaboration, and Cognitive Workflow. In *Cooperative Buildings. Integrating Information, Organizations, and Architecture,* number 1670 in LNCS. Springer-Verlag, 1999.

[8] Steve Mann. WearComp.org, WearCam.org, UTWCHI, and Steve Mann's Personal Web Page - research, 2004.

[9] Nexus. University of Stuttgart, Germany. http://nexus.informatik.uni-stuttgart.de/.

[10] Bruce Schechter. Seeing the light: IBM's vision of life beyond the PC, 1999.

[11] Smart Spaces. National Institute of Standards and Technology. http://www.nist.gov/smartspace/.

[12] SECURE Project Official Website. http://secure.dsg.cs.tcd.ie, 2002.

[13] Mark Weiser. The Computer for the 21st Century. *Scientific American,* pages 66–75, September 1991.

I

THE INFLUENCE OF CONTEXT ON PRIVACY, TRUST AND SECURITY

This page intentionally left blank

OVERVIEW

When interactions in pervasive computing are studied, the context in which these interactions are carried out must be taken into consideration. Context contributes to the meaning that a human being assigns to communication. The same data exchanged can mean something completely different in two different contexts. As an example, consider information about the routes to nearby hospitals. A pharmaceutical agent might query for that information sitting in a hotel room, planning a presentation tour for the next day. The same information might be accessed in case of an accident. Even if it is the same person accessing the same set of data, the meaning changes completely with the context in which the query is processed.

This, by all means simplified, example shows that contextual information can be used to alter the behaviour of an application, to adapt it to the current needs of its user. As the application changes, so change its requirements and provisions regarding privacy, trust and security. In the case of the example, privacy is probably not a main issue. It might be desirable to allow anonymous emergency calls, or it might not, for example for discouraging false alarms. However a sales agent would probably like to introduce himself anyway. But since a simple data query for most unclassified information would normally not involve the exposure of sensitive information, privacy concerns are rather low.

Security issues are less obvious. Even though anonymity is not a main issue, the sales agent might be interested in keeping information about his queries confidential. His personal record should be able to show up in the hospital's contact data base, but he would require the hospital's computing service to keep his query activity undisclosed and only revealed for specific purposes, e.g. security audits or billing. The sales agent has to trust the service in having and enforcing a suitable security policy. At the same time, the sales agent has to rely on his own devices to treat application data securely. This is subject to his own security policy, which generally applies to all applications running on his devices.

In the case of an emergency, security requirements are quite different. One point is that an emergency call should not be suppressed by

anybody. Also, a user would almost surely be willing to trigger an emergency call using any device that is at hand. If there are several devices available, he wouldn't even be required to trust any single device to operate correctly, but would likely push the emergency buttons on all of them, just to make sure.

The purpose of this rather elaborate example is to show that a change in context also changes the requirements on privacy, trust and security. These changes are analogous to the changes of the functional behaviour of an application in different contextual settings. But there is another, different aspect to contextual information. If such information becomes available to an adversary, the danger arises that through service usage, user behaviour or confidential communication is disclosed. Information about the context, in which interactions take place, can reveal sensitive information about the parties interacting, about their preferences, their goals, and the relations amongst them.

One of the most significant contextual features is the location of an entity. Lots of information can be inferred if the location of an entity is monitored, especially a human being. Therefore, location information should only be disclosed under certain well-established circumstances. This problem is addressed in the first paper, which is entitled "Survey on Location Privacy in Pervasive Computing". It is observed that location information can be obtained by several means, either through direct communication with the respective entity, or through indirect means such as observation or inference. Several technical and non-technical approaches for maintaining and managing location privacy are discussed. The main problem is that on one hand, location information is useful for enhancing services but on the other hand, it should be disclosed only sparingly. It is often feasible, however, to restrict the use of location information to lower levels of an application only, thereby trying to avoid revealing it unnecessarily, which decreases the requirements on the trust that is put in other entities.

Privacy concerns increase rapidly with respect to the advances in context-aware applications, because individuals fear that their personal information will be known and disclosed to others. The privacy concerns, in context-aware systems, stem from the fact that individuals are not aware of how much personal information is collected and what the content of context data is. In the second paper "Exploring the Relationship between Context and Privacy", consideration is given to the relationship between privacy and context with special attention to inference-based attacks on context information consisting of location, time and identifiers. The authors not only explore the relationship between privacy and context, but also propose a generic framework for modelling pri-

vacy in context-aware systems. The framework describes the essential and interrelated components for privacy in context-aware systems. The components considered in the framework are: a data model, an adversary model, inference rules, and privacy requirements. The authors back the framework up with the formalization of some components using structured data.

The last paper of this session, "Security, Privacy and Trust Issues Raised by the Personal Server Concept", presents the issues raised by an extreme (but not unlikely) incarnation of a ubiquitous computing device: the Personal Server. This is a small device that is able to store large amounts of data and can communicate wirelessly. However, it has no user interface, no means for direct interaction with a human being. It is clear that such a device needs special care if critical data is to be stored on it. The most pressing question is access control: who should be allowed to access the data, and how can the device be sure of the user's identity? But there are other issues as well, depending on the actual function of the device in a certain application scenario, for some of which there are no satisfying solutions yet.

This page intentionally left blank

SURVEY ON LOCATION PRIVACY IN PERVASIVE COMPUTING

Andreas Görlach[1], Andreas Heinemann[2], Wesley W. Terpstra[1,2]

[1] *ITO, TU Darmstadt, Germany*
{goerlach,terpstra}@ito.tu-darmstadt.de

[2] *GKEC, TU Darmstadt, Germany*
{aheine,terpstra}@gkec.tu-darmstadt.de

Abstract The goal of ubiquitous computing research is refine devices to the where their use is transparent. For many applications with mobile devices, transparent operation requires that the device be location-aware. Unfortunately, the location of an individual can be used to infer highly private information. Hence, these devices must be carefully designed, lest they become a ubiquitous surveillance system.

This paper overviews existing location-sensing mobile devices, vectors for a privacy invasion, and proposed solutions. Particular attention is paid to required infrastructure and the accuracy of the location information which can be stolen. Solutions are examined from the perspective of attacks which can be reasonably expected against these systems.

Keywords: Pervasive Computing, Ubiquitous Computing, Privacy, Location Privacy, Tracking, Positioning, Survey

1. Introduction

The proliferation of portable electronic devices into our day-to-day lives introduced many unresolved privacy concerns. The principle concern in this paper is that these devices are being increasingly equipped with communication capabilities and location awareness. While these features present a wide array of new quality-of-life enhancing applications, they also present new threats. We must be careful that the potential quality-of-life lost through the surrender of private information does not overwhelm the benefits.

An important question is how much privacy protection is necessary. Perfect privacy is clearly impossible as long as communication takes place. Therefore, research aims at minimizing the information disclosed. The required level of this protection is not a matter of technology; different people have different privacy needs. Nevertheless, technology should not *force* society to accept less privacy.

The major privacy concern with mobile devices equipped with communications ability is that they can reveal the location of their bearers. This concern is in itself not new; people can recognize each other. What is new is the increased scope of the problem due to automated information gathering and analysis. Poorly designed mobile devices enable anyone to obtain another's location.

If we allow automation to create an effective public record of people's locations, discrimination against minorities will be impossible to control. AIDS patients could be identified by the offices of doctors they visit, Alcoholics Anonymous members by their group meetings, and religious groups by their churches.

This paper will present an overview of the state-of-the-art in location privacy. In Section 2, mobile devices which possess both location awareness and communication ability will be examined. Section 3 lists attacks by which an invader can obtain private location information. Existing countermeasures and safeguards are detailed in Section 4. These include high level schemes such as policies which operate like contracts, and lower-level solutions which reduce information disclosure. Among the latter are anonymous routing algorithms, schemes for hiding within a group, methods to passively determine location, and frequency modulation techniques to hinder triangulation.

2. Location-Aware Communication Devices

Many technologies can determine the location of an individual. This section provides an overview of what technologies are presently deployed and which are coming in the near future.

One of the earliest systems designed for location tracking is the Global Positioning System (GPS) [9]. This system uses satellites to help devices determine their location. The GPS works best outdoors where it has line-of-sight to the satellites and few obstructions. For commercial products, resolution to within 4m is achievable. The GPS is widely deployed and integrated, especially in map applications. Although GPS devices do not transmit, they are being increasing integrated into PDAs and other devices which do.

For indoor use, the Active Badges [23] from AT&T Laboratories Cambridge were developed. These are small devices worn by individuals which actively transmit an identifier via infrared. This information is received by sensors deployed in the environment. This system provides essentially room-level resolution and has problems following individuals due to the infrequency of updates. The environment consolidates this information and can provide the current location of an individual.

A later refinement, the Bat [24], increased the detected resolution. With the increased resolution, the Bat can be used to touch virtual hot spots. Their work reports accuracy as good as 4cm. These refined devices used ultrasonic pings similar to bat sonar. However, once again the environment measures the Bat's location as opposed to real bats which learn about their environment.

The Cricket Location-Support System [18] system takes a similar approach. It uses radio and ultrasonic waves to determine distance and thus location. Like the Cambridge Bat, resolution to within inches is possible. As opposed to the similar Cambridge work, beacons are placed in the environment as opposed to on individuals. The Cricket devices carried by individuals listen to their environment in order to determine their location. In this way, the device knows its location, while the environment does not.

An approach to location sensing which does not require new infrastructure is taken by Carnegie Mellon University [21]. Here, the existing wireless LAN is observed by devices to recognize their location. By passively observing the signal strengths of various base stations, a device can determine it's location. Though there are no requirements for new infrastructure, there is a training overhead. During training a virtual map of signals is created which is used by the devices to determine their location.

Cell phones can be abused to provide location information. Although not originally intended for this purpose, the E-911 [19] requirements in the US forced cell phone providers to determine customer location when they dialed an emergency phone number. Although this practice was clearly beneficial, the technology has since spread. The underlying problem is the omnipresent possibility of performing triangulation (with varying accuracy, though).

In the near future Radio Frequency Identification (RFID) [8] will be found in many consumer goods. Intended as a replacement for barcodes, these tiny devices are placed in products to respond to a wireless query. Unlike barcodes, RFIDs are distinct for every item, even those from the same product line. This allows companies to determine their inventory

by simply walking through the shelves and automatically recording the observed products.

3. Attacks on Location Privacy

In a successful privacy attack, some party obtains unauthorized information. Individuals intend that some information about themselves should be available to others, and that the rest remain private. The means by which the individual's preferences were circumvented is the attack vector.

The main privacy concern with regards to ubiquitous computing is that many new *automated* attack vectors become possible. Loosely categorized, automated digital devices obtain information either through communication, observation, or inference. In this section the attack vectors available in each of these channels will be explored.

3.1 First-Hand Communication

An attacker obtains private information through first-hand communication when an individual unwittingly provides it directly to the attacker. In a world with ubiquitous computing, the threat of disclosure via accident or trickery is significant. All digital devices of a given type, by virtue of being homogeneous, make the same mistakes—and don't learn from them. The designers of the Windows file sharing protocol never intended it to be used to obtain people's names. Nevertheless, Windows laptops will happily reveal their owner's name to anyone who asks it. Due to a bug in bluetooth phones, attackers may often trick the phone into revealing its address book and phone number [16]. By asking a device with known location for owner information, both of these attacks pinpoint the owner's location, among other things. Naturally, these attacks can be built into an automated device.

Many ubiquitous devices also exhibit unwanted behaviour. The Bats and Active Badges broadcast their location information for all to hear. WLAN cards periodically emit traffic which includes their unique MAC ID. Devices providing exact their location information to location based services also seems overly permissive. At the bare minimum, these problems must be addressed.

A unique characteristic of digital devices is their potential for brainwashing. Manufacturers may choose to place secret spyware in their products[1] as a means to recoup financial losses. Furthermore, a vulnerability may allow an attacker to completely assume control of the device, and thus obtain a live location feed. For devices where the loca-

tion information is known to the infrastructure, the threat of a system vulnerability is magnified.

3.2 Second-Hand Communication

Attacks via second-hand communication relay information from one party to another unauthorized party. The primary difference between these attacks and first-hand attacks is that the individual no longer controls the information. Fortunately, in the human scenario, talking about individuals behind their back requires some expenditure of breath. Unfortunately, aggregation and spreading of this information in a digital system is significantly easier.

This behaviour has already been observed in the Internet where Doubleclick regularly sells personal habit and preference information. It seems naïve to assume that the much finer grained information available from ubiquitous devices will not similarly be sold. Services are already available for individuals to locate their friends via the cell phone networks [17].

3.3 Observation

Attackers may also obtain information by configuring devices to observe their environment. The most obvious problem is the deployment of many nearly-invisible cameras in the environment. However, there are other risks which are more feasible to launch with current technology.

One of the more interesting attacks that can be launched against mobile communications-equipped devices is triangulation. By measuring timing delays in a signal, the attacker can determine the location of the device. This is similar to how the Bat operates, only using electromagnetic waves instead of sound waves.

3.4 Inference

One of the fears about automated privacy invasion is the compilation of a profile. After gathering large amounts of information via communication and observation, an automated system combines these facts and draws inferences. Given enough data, the idea is to build a complete picture of the victim's life.

From a more location-centric point of view, location information could be processed to obtain useful information for discrimination. If a person regularly visits the location of a group meeting, she is probably a member of that group. In the consumer arena, the fact that an individual shops at a particular store at regular intervals may be useful information for price discrimination [1].

Tracking an individual's location through time may also enable an attacker to link information to the individual. For example, if an individual's car regularly sends out *totally anonymous* weather requests, it might still be possible for a weather network to track the car by correlating the observed request locations. Later, when the individual buys gas at an affiliate's gas station, the network can link the individual's name and bank account to the tracked car. Now, the network can deduce information such as where the person shops, lives, and works; who the person regularly visits; etc.

4. Solutions

In the literature there exist several approaches to protect the location of a user. Most of them try to prevent disclosure of unnecessary information. Here one explicitly or implicitly controls what information is given to whom, and when. For the purposes of this paper, this information is primarily the identity and the location of an individual. However, other properties of an individual such as interests, behaviour, or communication patterns could lead to the identity and location by inference or statistical analysis.

In some cases giving out information can not be avoided. This can be a threat to personal privacy if an adversary is able to access different sources and link the retrieved data. Unwanted personal profiles may be the result. To prevent this, people request that their information be treated confidentially. For the automated world of databases and data mining, researchers developed policy schemes. These may enable adequate privacy protection, although they similarly rely on laws or goodwill of third parties.

4.1 Policies

In general, all policy based approaches must trust the system. If the systems betrays a user, his privacy might be lost. Here, the suitable counter-measure is a non-technical one. With the help of legislation the privacy policy can be enforced.

All policy based systems have the drawback that a service could simply ignore the individual's privacy preferences and say, "To use this service you have to give up your privacy or go away." This certainly puts the user in a dilemma and he will probably accept these terms as he wants to use the service.

A Privacy Awareness System (pawS) for Ubiquitous Computing Environments. In [14, 15] Langheinrich proposes the *pawS* sys-

tem. *pawS* provides users with a privacy *enabling* technology. This approach is based on the Platform for Privacy Preferences Project (P3P) [4], a framework which enables the encoding of privacy policies into machine-readable XML. Using a trusted device, the user negotiates his privacy preferences with the UbiCom environment.

Framework for Security and Privacy in Automotive Telematics. A framework for security and privacy in automotive telematics, i.e. embedded computing and telecommunication technology for vehicles, is described by Duri *et al.* [5]. The primary goal of their framework is to enable building telematics computing platforms that can be trusted by users and service providers. They do that by installing a *data protection manager* to handle sensitive data. Thus they implement a middleware working with different key concepts which for example influence location data accuracy and enable user defined privacy policies.

Concepts for Personal Location Privacy Policies. Snekkens [22] presents concepts which may be useful when constructing tools to enable individuals to formulate a personal location privacy policy. Snekkens's idea is that the individual should be able to adjust the accuracy of his location, identity, time, and speed and therefore have the power to enforce the need-to-know principle. The accuracy is dependent on the intended use of the data, and the use in turn is encoded within privacy policies.

4.2 Protecting First-Hand Communication

Most approaches address the problem of information disclosure. Many different ideas have been proposed to prevent unnecessary information from becoming known to a third party.

ANODR: ANonymous On Demand Routing. With the scenario of a battlefield in mind, Kong and Hong described in [13] their scheme ANDOR. This is a routing protocol addressing the problems of route anonymity and location privacy.

The intention is that packets in the network can not be traced by any observing adversary. Additionally, their routing scheme provides unlinkability. Prior to one node's ability to send a message to another, a route must be established through route discovery. This route discovery is achieved by broadcasting and forwarding packets. The sender of a message is anonymous because it is impossible to judge whether a node is actually sending a message it generated or is simply forwarding a packet as part of a route.

MIXes in Mobile Communication Systems. It is easy for cellular networks like GSM to track their mobile subscribers. Location information is required in order to route calls appropriately. Avoiding this by simply broadcasting is not an option because of the limited bandwidth in current cellular networks. In [6] this is investigated and the application of MIXes (see also [3]) is proposed.

In their system, the scheme does not keep the identity—telephone number—of the recipient anonymous. Only the location of the recipient is protected. Remarkably, their system remains secure even if *all* of forwarding nodes are observed by an adversary.

Mix Zones. A recent approach which is somewhat similar to mix networks is *mix zones* [2]. In these networks, the infrastructure provides an anonymity service. The infrastructure delays and reorders messages from subscribers within a mix zone to confuse an observer.

A problem with this system is that there must be enough subscribers in the mix zone to provide an acceptable level of anonymity. Beresford and Stajano conducted statistical attacks against these systems and found the afforded security to be quite low. Even large groups using the Active Bat remained vulnerable.

Temporal and Spatial Cloaking. In [10], Gruteser and Grunwald propose a mechanism called *cloaking* that conceals a user within a group of k people. They consider a user as *k-anonymous* if, and only if, they are indistinguishable from at least $k-1$ other users. To achieve this, the accuracy of the disclosed location is reduced. Then any of the people within the disclosed area could have been the user. Similarly, they consider reducing the accuracy of disclosure timestamps. Like Stajano and Beresford they, too, measured anonymity in experimental setups, but unlike them Gruteser and Grundwald identified concrete values which in their view provide certain levels of anonymity.

The Cricket Location-Support System. In order to prevent the potential misuse of personal information, the most convincing solution is to not let out any information at all. This idea is applied directly to the Cricket Location-Support System [18]. As described in section 2, the mobile device never transmits at all; rather, it passively listens to its environment.

This system is ideally suited to an office. The transmitters need not be connected to each other or a network. This not only supports privacy, but also makes things cheaper and more maintainable. However, the use of services sometimes makes it necessary for the device to disclose

its location. For example, using a printer implicitly reveals that the device is near to the printer.

PlaceLab. Place Lab [20] uses a decentralized approach with WLAN hotspots as beacons thus exploiting the same idea as Cricket for outdoor environments. In order to determine its own position each device hosts a previously downloaded database of an access point to location mapping. With their Place Bar component end users are able to adjust their location granularity when revealing location information to third parties.

In addition to that the privacy and security concept of Hong et al. [11] also take access point privacy, network service privacy and web service privacy into account. Access point privacy aims at protecting access point owners by hashing the stored MAC addresses.

The Blocker Tag. A special case among pervasive devices are RFIDs. People carrying objects which contain RFIDs might not even be aware of the existence of these devices because of their size[2] and their passive nature. A second specific property of RFIDs is their unability to do any computation like e.g. encryption. So they require their own measures for privacy protection.

Juels, Rivest and Szydlo examined several possible solutions in [12] ranging from destruction of the tag through less destructive approaches to regulation (i.e. policies). Since the authors see disadvantages in all of the examined solutions they present their own approach which is the development of a special tag: the Blocker Tag. This tag blocks attempts of readers to identify RFIDs. In order not to block desired RFIDs or to temporarily enable the reading of RFIDs the blocking process can be done selectively.

Hindering Triangulation. As mentioned in Section 3.3, data can be gathered by observing a device or person. On the physical layer it is usually possible to locate a sending device by recording signal delays and performing triangulation.

In [7] frequency modulation schemes are discussed to prevent location of mobile devices. The researchers performed an in-depth analysis of direct sequence spread spectrum. Their idea is to make it difficult to distinguish a signal from random background noise. This is done by distributing the data on pseudorandomly chosen channels. By knowing a shared secret the supposed receiver is able to reassamble the messages. The drawback of this solution is that it requires existing infrastructure to be changed and consumes considerably more bandwidth.

5. Conclusions

The solutions we have seen can be categorized into policies and information minimizing at the source. These approaches aim to address threats in the areas of first- and second-hand communication, observation, and inference.

Policies seem to work well wherever consent underlies the transaction. For example, when information is to be provided to a service, an agreement can be reached regarding the further distribution of the information. If no agreement can be reached, then the individual will be unwilling to use the service, but the service will likewise not obtain the information or any associated remunerate. Similarly, the individual can negotiate terms about how his information may be used; this can address attacks based on inference.

There is no consent in observation. This means that policies can not be applied to these attacks since the individual is in no position to negotiate. Here, legal safeguards and countermeasures are required. Unfortunately, there is currently insufficient discourse between technical and legal experts.

Accuracy reduction techniques apply primarily to first-hand communication problems. These schemes aim at reducing the amount of confidential information disclosed to third parties. There are a variety of techniques which obscure the location information, the timestamp of the transaction, and the identity of the individual.

As mentioned in the introduction, privacy issues are fundamentally not technical. As ubiquitous devices permeate the every-day lives of ordinary citizens, our privacy protection measures will have increasing impact on their lives. It is important that research into privacy protection bear in mind what must be protected. This is more the area of social sciences, and thus requires more inter-disciplinary discourse.

6. Acknowledgements

This work was sponsored in part by the Deutsche Forschungsgemeinschaft (DFG) as part of the PhD program “Enabling Technologies for Electronic Commerce”.

Notes

1. For example the Kazaa Media Desktop
2. The smallest RFIDs are currently only of 0.4mm * 0.4mm size.

References

[1] Joseph Bailey. Internet Price Discrimination: Self-Regulation, Public Policy, and Global Electronic Commerce, 1998.

[2] Alastair R. Beresford and Frank Stajano. Location Privacy in Pervasive Computing. *PERVASIVE computing, IEEE CS and IEEE Communications Society,* (1):46–55, 2003.

[3] David Chaum. Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms. *Communications of the ACM,* 24(2):84–88, 1981.

[4] Lorrie Cranor, Marc Langheinrich, Massimo Marchiori, Martin Presler-Marshall, and Joseph Reagle. The Platform for Privacy Preferences 1.0 (P3P1.0) Specification. `http://www.w3.org/TR/P3P/`, seen 2004.

[5] Sastry Duri, Marco Gruteser, Xuan Liu, Paul Moskowitz, Ronald Perez, Moninder Singh, and Jung-Mu Tang. Framework for Security and Privacy in Automotive Telematics. In *International Conference on Mobile Computing and Networking,* pages 25–32. ACM Press, 2002.

[6] Hannes Federrath, Anja Jerichow, and Andreas Pfitzmann. MIXes in Mobile Communication Systems: Location Management with Privacy. In *Information Hiding,* pages 121–135, 1996.

[7] Hannes Federrath and Jürgen Thees. Schutz der Vertraulichkeit des Aufenthaltsorts von Mobilfunkteilnehmern. *Datenschutz und Datensicherung, Verlag Vieweg, Wiesbaden,* 6(6):338–348, 1995.

[8] Klaus Finkenzeller. *RFID-Handbook, 2nd Edition.* Wiley & Sons LTD, 2003.

[9] I. A. Getting. The Global Positioning System. *IEEE Spectrum,* 30(12):36–47, December 1993.

[10] Marco Gruteser and Dirk Grunwald. Anonymous Usage of Location-Based Services Through Spatial and Temporal Cloaking. In *MobiSys,* pages 31–42. USENIX, 2003.

[11] Jason Hong, Gaetano Bordello, James Landay, David McDonald, Bill Schilit, and Doug Tygar. Privacy and Security in the Location-enhanced World Wide Web. In *Proceedings of Ubicomp 2003,* October 2003.

[12] Ari Juels, Ronald L. Rivest, and Michael Szydlo. The Blocker Tag: Selective Blocking of RFID Tags for Consumer Privacy. In *V. Atluri, ed. 8th ACM Conference on Computer and Communications Security,* pages 103–111. ACM Press, 2003.

[13] Jiejun Kong and Xiaoyan Hong. ANODR: ANonymous On Demand Routing with Untraceable Routes for Mobile Ad-hoc Networks. In *Proceedings of the 4th ACM international symposium on Mobile ad hoc networking & computing,* pages 291–302. ACM Press, 2003.

[14] Marc Langheinrich. Privacy by Design – Principles of Privacy-Aware Ubiquitous Systems. In Gregory D. Abowd, Barry Brumitt, and Steven A. Shafer, editors, *Ubicomp,* volume 2201 of *Lecture Notes in Computer Science,* pages 273–291. Springer, 2001.

[15] Marc Langheinrich. A Privacy Awareness System for Ubiquitous Computing Environments. In Gaetano Borriello and Lars Erik Holmquist, editors, *Ubicomp,* volume 2498 of *Lecture Notes in Computer Science,* pages 237–245. Springer, 2002.

[16] Adam Laurie. Serious Flaws in Bluetooth Security Lead to Disclosure of Personal Data, http://www.bluestumbler.org, 2003.

[17] Mobiloco - Location Based Services for Mobile Communities. http://www.mobiloco.de/.

[18] Nissanka B. Priyantha, Anit Chakraborty, and Hari Balakrishnan. The Cricket Location-Support System. In *Mobile Computing and Networking,* pages 32–43, 2000.

[19] Jeffrey H. Reed, Kevin J. Krizman, Brian D. Woerner, and Theodore S. Rappaport. An Overview of the Challenges and Progress in Meeting the E–911 Requirement for Location Service. *IEEE Communications Magazine,* 5(3):30–37, April 1998.

[20] Bill Schilit, Anthony LaMarca, Gaetano Borriello, William Griswold, David McDonald, Edward Lazowska, Anand Balachandran, Jason Hong, and Vaughn Iverson. Challenge: Ubiquitous Location-Aware Computing and the Place Lab Initiative. In *Proceedings of The First ACM International Workshop on Wireless Mobile Applications and Services on WLAN (WMASH 2003).* ACM Press, September 2003.

[21] Asim Smailagic, Daniel P. Siewiorek, Joshua Anhalt, David Kogan, and Yang Wang. Location Sensing and Privacy in a Context Aware Computing Environment. In *Pervasive Computing,* 2001.

[22] Einar Snekkenes. Concepts for Personal Location Privacy Policies. In *Proceedings of the 3rd ACM Conference on Electronic Commerce,* pages 48–57. ACM Press, 2001.

[23] Roy Want, Andy Hopper, Veronica Falcão, and Jonathan Gibbons. The Active Badge Location System. *ACM Transactions on Information Systems,* 10(1):91–102, 1992.

[24] Andy Ward, Alan Jones, and Andy Hopper. A New Location Technique for the Active Office. *IEEE Personal Communication,* 4(5):42–47, 1997.

EXPLORING THE RELATIONSHIP BETWEEN CONTEXT AND PRIVACY

Timo Heiber, Pedro José Marrón
University of Stuttgart
Institute for Parallel and Distributed Systems (IPVS)
{timo.heiber,pedro.marron}@informatik.uni-stuttgart.de

Abstract Privacy is an important consideration for context-aware systems, because an individual's context contains a large amount of personal information. In this article, we describe a generic framework to model privacy in context-aware systems. We also present an example instance of the framework to demonstrate its practical application.

Keywords: Pervasive computing, security, privacy, context-awareness, inference control

1. Introduction

Most people agree that privacy protection is an important aspect of networked and distributed applications, especially in the fields of mobile and pervasive computing. However, it is hard to agree on a common definition of privacy, for two main reasons: First, the definition depends on the highly variable preferences of individuals and socio-cultural groups. Secondly, in contrast to the related security goal of data confidentiality, privacy is not an all-or-nothing notion. It is often acceptable to divulge a limited amount of personal data, whereas it may be unacceptable if large amounts of the same type of data become known.

The problem becomes even harder when one considers the question of personal privacy with respect to context-aware applications, i.e. applications that take the context of entities into account. In pervasive computing, the most important entities are individuals. According to [3], context is information that describes the situation of an individual, which means that the question of personal privacy arises naturally: The amount of context information that is personal (such as the location of a user) or related to personal information (such as the location of a user's mobile device) could conceivably grow quite large. Additionally,

someone interested in obtaining personal information (hereafter termed "adversary") would have a multitude of opportunities. Moreover, the semantics of context information can be leveraged to infer context information that is not explicitly stated in the available pieces of context information. Consider, for instance, the point that the location of a certain user's mobile device can be used to infer information about the location of that user.

Assuming a global view of the problem, there are three main questions that influence a user's degree of privacy in context-aware environments:

1 How much personal context data can be collected by an adversary?

2 What is the content of that context data?

3 How successful is the adversary in attributing that data to a particular person?

In this paper, we present a generic framework for privacy in context-aware computing systems. We focus on how to model inference-based attacks on context information within this framework. We demonstrate the feasibility of our approach by showing how to model inferences based on the context information of location, time and identifier (sometimes called primary context [3, 8]).

This paper is structured as follows: In Section 2, we present an example scenario from which we derive a generic framework for privacy in context-aware systems in Section 3. We then discuss the formalization of this generic framework in Section 4, with a concrete instance provided in Section 5. We review the related work in Section 6. In Section 7, we summarize our approach and discuss directions for future work.

2. Motivation

Consider a scenario with an abundant supply of context-based information systems: Location-based services track the locations of their customers and supply information relevant at their current location (e.g. route planning, public transport timetables etc.) while "infostations" supply information to anyone in their transmission range. User Alice uses her personal devices to communicate with such information systems. Location tracking and communication with the location-based service is done via a mobile phone network that can provide high location resolution (e.g. UMTS). Access to the infostations is gained through her WLAN-equipped PDA.

We assume that Alice needs to authenticate herself to the location-based information system (LBS) for billing purposes. A a consequence,

she is known to the LBS under the identifier *Alice-LBS* (which might be a pseudonym). The PDA uses a wireless LAN adapter with the constant device ID (MAC address) *Alice_PDA*.

Now consider adversary Eve that has gained access to the information generated by the transmissions of Alice's devices (for example, a UMTS service provider that also monitors WLAN traffic at some locations). Eve could then collect two types of location traces for all users. With respect to Alice, she would obtain location traces under the ID *Alice-LBS* and also other location traces under the ID *Alice_PDA,* using the location information that comes implicitly with the WLAN transmissions.

Eve's next step would be to correlate both types of location traces in order to see whether a pair of location traces from different sources matches. That way, two different identifiers *(Alice_LBS* and *Alice-PDA)* could be linked to the same person (Alice). Furthermore, only one success in this regard will be enough to link Alice's identifiers *Alice-LBS* and *Alice_PDA* from this point on.

The important point of this scenario is that with increasing amounts of context information, attempts to penetrate an individual's privacy will also be increasingly successful, because the adversary will be able to leverage the semantics of context data items to infer additional information that is not explicitly stated in the context information. Even if data is only stored in pseudonymous form, adversaries will often be able to link items of context data based on their content. Moreover, the problem discussed here is not restricted to a specific application scenario, but remains valid for any form of constant identifier, for instance RFID tags that can be interrogated remotely.

Note that the amount of data used by Eve in the example above is comparatively small and restricted to identifiers and spatio-temporal coordinates. This is an indication that the privacy problems will become even worse when context-aware computing is used ubiquitously and other forms of context data are taken into account.

3. A Generic Framework for Modeling Privacy in Context-Aware Systems

In this section, we describe a common framework for privacy that reflects the main factors that are relevant for any model of privacy in context-aware systems. It consists of several interrelated components that model the essential parameters for privacy in context-aware systems. These are the contents of the data, the capabilities of the adversary to obtain data, possible inferences and the actual privacy requirements, as shown in Figure 1.

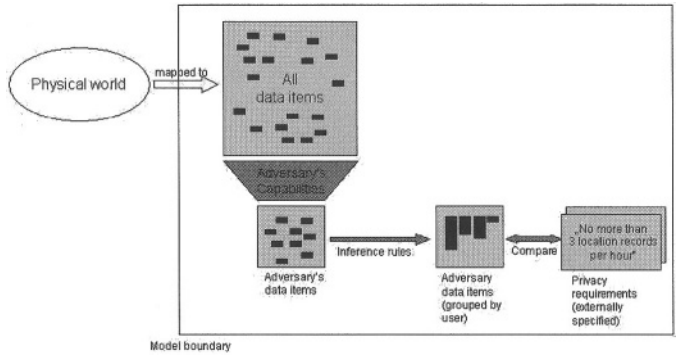

Figure 1. Generic framework for privacy

The *physical world* is external to our model boundaries and represents the events occurring in the physical world: people walking around, people using computers to access information or send messages etc.

Some of these events leave an "electronic trace", e.g. cause records of information to be stored on a computer system. We refer to each of these discrete records as a *data item.* Conceptually, we consider the *set of all data items* to be one database to which an adversary trying to violate people's privacy would like to gain access. A data item may, for instance, be created when a person sends an e-mail message, or when his or her whereabouts are recorded by a location tracking system. In the example of Section 2, data items containing Alice's location are generated at the location-based service and also due to the communication of the PDA.

A subset of all existing data items is available to an adversary *(adversary's data items).* The size and exact composition of this subset depends on the *adversary's capabilities,* which are a function of her abilities and the access control mechanisms employed to restrict access to the data items. In the "Alice" example, we assumed Eve to have access to the data stored at both the public service and the two non-public services.

The next step of the adversary is to apply *inference rules* to organize the data available to her. Organizing the data items refers to grouping the data items by user. This grouping is done by examining the contents of the data items to determine which data items have been created through the activities of the same user.

Whether the *privacy requirements* of a certain user have been violated is determined by how much and what type of information the adversary has gained access to.

Based on the description above, a model for privacy in context-aware system needs four components:

1 A *data model* that describes what kind of data items are created.

2 An *adversary model* that describes what data items the adversary can gain access to.

3 The *inference rules* that can be applied to the data by the adversary.

4 And, finally, a characterization of the *privacy requirements* for the system.

A formal characterization of these four components makes it possible to derive the knowledge that can possibly be gained by an adversary and evaluate it with respect to previously stated privacy requirements for the context-aware system.

In the following section, we provide a formalization of the first three components using structured data and predicates on this data. That way, it is possible to derive the knowledge that can possibly be gained by an adversary. This suffices to evaluate whether simple privacy requirements like "no more than n location records within a time interval of length t" hold. A system for formally stating complex privacy requirements within this model is beyond the scope of this paper and will be considered in future work.

4. Formalization of the Model

The generic formalization of the model is based on predicates that describe the capabilities and inferences of the attacker.

4.1 Generic Data Model

With respect to the data model, two questions need to be answered:

1 What information do we need to represent?

2 How can we design a flexible and extensible data model?

Let us first consider what information we need to model with respect to user privacy. Here, the most important pieces of information are identity, location and time. We refer to these as *primary context.* All other context information is referred to as *secondary context.* This includes

specific constant properties of an entity (such as user preferences) and its current state and activity. Our definition of context is a variation of existing definitions [3, 8].

In order to achieve flexibility and extensibility, we use the following generic representation for data items: A data item is a four-tuple *(ID, Location, Time, Secondary Context)* containing the following information:

ID The identifier under which this data item was created. In the example, this field would contain Alice's customer ID or the MAC address for her PDA.

Location Spatial information about a data item. This field would for example contain Alice's location when it is stored by the location-based service.

Time Temporal information about the data item. In the example, this refers to the time of a communication or the time of an update of location data.

Secondary Context Any other information in the data item. An example for *secondary context* would be the content of Alice's communication.

We think of data items as generic records that can be further structured into subfields, depending on the actual data created in an application scenario. Using dot-notation to refer to subfields, we would, for example, model the different types of IDs by substructuring the ID field into *ID.MAC_Address* and *ID.Customer_ID*. Such substructuring can also be used to introduce name spaces in order to avoid clashes between the customer IDs of several service providers. In this case, we would introduce the field *ID.Provider* to name the service provider explicitly. A complete instance of the generic data model can be found in Section 5.1.

4.2 Generic Adversary Model

The adversary model is, in effect, a filter applied to the set of all data items. We represent the capabilities of an adversary with a generic predicate *visible_to_adversary*. This predicate is defined on data items and evaluates to true if the adversary can learn this data item. A concrete instance of this predicate can be found in the example in Section 5.2.

4.3 Generic Inference Rules

Context data items are a-priori independent of each other. However, primary context contains sufficient information relating to users that can be exploited to learn whether data items were caused by the same user. In effect, the adversary infers, based on the fields of a data item, that certain data items relate to the same person. As a result, the adversary collects of sets of data items, where all data items in the same set can be attributed to the same person. In this section, we describe the generic structure of an inference system based on primary context.

Referring back to our example again, we saw how inference allowed the adversary to link data items based on their content (in this case, user IDs). Also, correlation of spatio-temporal coordinates made it possible to link unrelated identifiers for Alice and increase the amount of knowledge about her. This means that there are two types of inference rules: Linking based on user IDs, which only requires examining the content of two single data items, and correlation of coordinates, which needs sufficient overlap in whole location traces of a user. That is, in the second case, two whole *sets* containing already linked data items must be examined in order to obtain a match.

We represent these linking strategies by two generic inference rules, one that deals with linking data items, thereby aggregating them into sets of linked data items and one that deals with linking sets of already linked data items, thereby producing even larger sets of data items. The privacy of a person degrades directly with the size of the set of data items attributable to him or her.

Formally, the generic inference model provides two inference rules, one that works on pairs of single data items and one that works on pairs of sets of data items. These rules are based on two predicates, which are instantiated according to the data model and the inference possibilities of the application scenario:

linkable The predicate *linkable* is defined on data items: Two data items are *linkable* if their respective contents warrant the conclusion that they relate to the same person. For example, two data items that contain the same unique identifier for a person could be considered to be linkable. The predicate is transitive and induces an equivalence relation on data items.

matching The predicate *matching* is defined on sets of data items: It represents those cases where correlation of two sets of linked data items leads to the conclusion that both sets relate to the same person. For example, two sets of data items are matching if they

contain large numbers of matching location records from highly different locations, *and* there is no pair of data items from the different sets that record different locations for the same point in time. Note that the negation implies that this predicate is not necessarily transitive.

Using these predicates, an adversary can execute Algorithm 1 and, after that, Algorithm 2. Each of the sets obtained in this way is a representations of an actual person. This means that, for the adversary, a person is defined as a set of the data items created by that person. Section 5.3 will provide a concrete instance of these predicates.

Algorithm 1 Generic algorithm for collecting linkable data items

function collect_items (D : **set of** data_item) : **set of** (**set of** data_item)
begin
if $D = \{\}$ **then**
  return $\{\}$
end if
{Otherwise, build a set of data items that are (transitively) linkable}
select any $d \in D$
$D := D - \{d\}$
$R := \{d\}$
for all $d' \in D$ **do**
  if $\exists d'' \in R : linkable(d', d'')$ **then**
    $D := D - \{d'\}$
    $R := R \cup \{d'\}$
  end if
end for
return $\{R\} \cup collect\_items$(D)
end

5. Modeling the Example Scenario

In this section, we formalize the data model, adversary model and inference rules used in the example scenario of Section 2.

5.1 Data Model

For the example, we need to model two types of data items, one for the WLAN communication and one for the location-based service.

Algorithm 2 Generic algorithm for building sets attributable to the same person

```
function collect_sets (S : set of (set of data_item))
                       : set of (set of data_item)
begin
if ∃s1, s2 ∈ S : matching(s1, s2) then
  S := S − {s1, s2}
  S := S ∪ {s1 ∪ s2}
  return collect_sets(S)
end if
return S
end
```

WLAN Communication. Each transmission of a WLAN-equipped device creates a data item. In order to represent this, we use data items of the following format:

ID This field has the following subfields:

ID.MAC_Address The MAC address of the device.

ID.Technology The technology used to make transmissions. All data items caused by 802.11 wireless LAN devices will have the constant value "IEEE 802.11 MAC".

Location The location at which the transmission occurred. This location is the area served by a certain WLAN access point.

Time The time at which the transmission occurred. If an adversary can perceive a transmission, this information will be fairly exact, since the delay between physical transmission and reception will be negligible.

Location-Based Service. The location records for the location-based service have the following form:

ID Again, we make use of subfields:

ID.Customer_ID The customer for which this record is created.

ID.Provider The name of the service provider.

Location The location at which the transmission occurred as determined by the location system in use.

Time The time at which this record was created. Again, it should be possible to determine this information in a fairly exact way.

For two locations l_1 and l_2, we write $l_1 \sim l_2$ if and only if l_1 and l_2 are less than 50 meters apart. We also define a comparison operator $\approx$ for times. For two times t_1 and t_2, $t_1 \approx t_2$ if and only if t_1 and t_2 are less than one minute apart. For the purpose of this discussion, we assume that location and temporal information can be determined with a good enough accuracy for theses operators.

5.2 Adversary Model

In the context of our framework, the amount of data items the adversary Eve can actually perceive will depend on her capabilities. For the sake of the example, we assume that adversary Eve is capable of overhearing wireless LAN transmissions and has full access to location-based service X:

A data item d is *visible_to_adversary* if

$$\begin{aligned} d.ID.Technology &= \text{“IEEE 802.11 MAC”} \vee \\ d.ID.Provider &= \text{“Location-based Service X”} \end{aligned}$$

5.3 Inference Rules

The predicates *linkable* and *matching* are defined as follows:

Linkable. Two data items, d_1 and d_2 are *linkable* if

$$\begin{aligned} (d_1.ID.Technology &= d_2.ID.Technology \wedge \\ d_1.ID.MAC\_Address &= d_2.ID.MAC\_Address) \vee \\ (d_1.ID.Provider &= d_2.ID.Provider \wedge \\ d_1.ID.Customer\_ID &= d_2.ID.Customer\_ID). \end{aligned}$$

This definition assumes constant identifiers and defines linkability by identity of the identifiers. *ID.MAC_Address* and *Provider* are used to provide name spaces for the identifiers.

Matching. For a parameter k, which is dependent on the accuracy with which location and time information can be captured and compared, two sets of data items D_1 and D_2 are *matching* if for some $k' \geq k$

$$\begin{aligned} &(\exists D_1' = \{d_{11}, \ldots, d_{1k'}\} \subseteq D_1, D_2' = \{d_{21}, \ldots, d_{2k'}\} \subseteq D_2 : \\ &\quad \forall 1 < i < k' : d_{1i}.Time \sim d_{2i}.Time \wedge d_{1i}.Location \approx d_{2i}.Location) \wedge \\ &(\neg \exists d_1 \in D_1, d_2 \in D_2 : \\ &\quad d_1.Time \sim d_2.Time \wedge d_1.Location \not\approx d_2.Location) \end{aligned}$$

where ~ and ≈ are defined as in Section 5.1.

The predicate *matching* is defined by requiring two sets of data items to contain a sufficient number of items that place the user at the same location at the same time. The match fails if the two sets contain data items that have the user at different locations at the same time.

5.4 Remarks

A set of data items derived through application of this definition describes a person in terms of the device he or she used and the locations at which that occurred. It is noteworthy that sets derived in this way do not contain directly identifying information. However, sufficiently detailed location information would make identification of any person easy (e.g. because most people spend most of their time at home). Also, after a person has been identified once, his or her name can always be linked to the use of his or her personal device.

6. Related Work

Pervasive Computing scenarios [7, 12] are full of privacy issues. However, much of the current work in this field has, with some exceptions, not yet progressed much beyond the conceptual stage [10, 11].

The Platform for Privacy Preferences Project (P3P) [1] aims at developing machine-readable privacy policies and preferences. This approach is somewhat related to our model component for privacy conditions. An interesting issue that comes up in both P3P and our work and that is worth further investigation is how preferences can be described in an easy to understand and human-readable form and then transformed into a more formal representation. Marc Langheinrich, one of the authors of P3P has also extended the P3P concept to ubiquitous computing scenarios [6].

The Freiburg Privacy Diamond [13] is more closely related to our approach. The authors model privacy in mobile computing environments using relations between the sets of users, devices, actions and locations. The only inference rule in their model is transitive closure. As a result, the expressiveness of the model is limited. The authors also discuss the possibility of including probabilities and time in their model, although it remains unclear where the probabilities come from and the concept of time is only mentioned briefly in the paper.

Snekkenes [9] discusses access control policies for personal information. He proposes a lattice model to describe the accuracy of information (e.g. location, time or identifying information). The way we represent identifying information as sets of data items relating to the same person

is comparable (but not identical) to his approach. The high-level view of the privacy problem presented in our work does not yet consider the accuracy of other types of information. We plan to consider the question of accuracy of information in our further work.

Hengartner and Steenkiste [4, 5] consider access control with respect to personal information in a pervasive computing context. Their second work [4] mentions the need to model the relationships between different "pieces of information", although the paper does not yet give any details about their approach. The generic data model and inference system presented here is an attempt to provide such a model.

Inference Control [2] is the common term for approaches to limit the inference capabilities of an adversary. Our framework provides a method to model the semantic inference capabilities of the adversary, in contrast to the more common syntactic approaches to inference control.

7. Conclusion and Further Work

The contributions of this paper are twofold: First, we presented a generic framework for discussing the privacy problem in context-aware systems. Secondly, we introduced the inference problem for context data by providing a generic model for the representation of context data and concepts for the modeling of inferences based on the primary context of location, time and ID.

We are presently in the process of formalizing the inference model for primary context more rigorously, using a restricted form of First-Order Logic. We are also exploring ways to represent inexact information and uncertain inference within our model.

Future work will extend our model to inferences that take secondary context into account. Note that the inclusion of the context information of *Activity* alone will open up further inference problems (not the least of them being the fact that activities will often provide implicit location information).

Additionally, we are working on methods to evaluate the accuracy of models based on our framework and to derive access control rules for personal information derived from such models.

Appendix: Discussion Results

This appendix briefly summarizes the results of the discussion following the presentation of the preceding work at the Workshop on Security and Privacy in Pervasive Computing.

A point was made that one could always assume that the adversary can see all data items generated and forego the adversary model altogether. In the case of *pervasive* computing, however, this does not seem appropriate, since a potentially very large

amount of data items will be generated and many classes of adversaries will have access to only a small subset of them. In pervasive computing, the trade-off between privacy and functionality needs to be considered explicitly. If we erroneously assume highly powerful attackers, and limit the generation of data items (and, consequently, possible functionality) based on this assumption, we might needlessly restrict functionality.

Additionally, two major lines were identified for future work:

1 Usage of the model. Two usage scenarios for a privacy model are possible: The first one aims at evaluating a given pervasive computing scenario and determining whether a certain set of privacy requirements can be satisfied. The second one is based on observing the history of a user's interactions and advising whether a certain action of that user would have harmful effects on his or her privacy.

2 Incorporating probability into the model. Participants at the workshop also favored the early inclusion of probability and probabilistic inference into the framework.

The authors would like to thank all participants for the lively discussion and the excellent feedback given.

References

[1] Lorrie Cranor, Marc Langheinrich, Massimo Marchiori, Martin Presler-Marshall, and Joseph Reagle. The Platform for Privacy Preferences 1.0 specification. W3C Recommendation, The World Wide Web Consortium, April 2002.

[2] Dorothy E. Denning. *Cryptography and Data Security.* Addison-Wesley, 1982.

[3] Anind K. Dey. Understanding and using context. *Personal Ubiquitous Comput.,* 5(1):4–7, 2001.

[4] Urs Hengartner and Peter Steenkiste. Access control to information in pervasive computing environments. In *Proceedings of the 9th Workshop on Hot Topics in Operating Systems (HotOS IX),* Lihue, Hawaii, May 2003.

[5] Urs Hengartner and Peter Steenkiste. Protecting access to people location information. In *Proceedings of the First International Conference on Security in Pervasive Computing (SPC 2003),* Lecture Notes in Computer Science, Boppard, Germany, March 2003. Springer-Verlag.

[6] Marc Langheinrich. A privacy awareness system for ubiquitous computing environments. In Gaetano Borriello and Lars Erik Holmquist, editors, *UbiComp 2002: Ubiquitous Computing, 4th International Conference,* volume 2498 of *Lecture Notes in Computer Science,* page 237ff, Göteborg, Sweden, September 29 - October 1 2002. Springer-Verlag.

[7] Friedemann Mattern. The vision and technical foundations of Ubiquitous Computing. *Upgrade,* 2(5):75–84, October 2001.

[8] Kurt Rothermel, Martin Bauer, and Christian Becker. Digitale Weltmodelle – Grundlage kontextbezogener Systeme. In Friedemann Mattern, editor, *Total Vernetzt?!* Springer-Verlag, 2003.

[9] Einar Snekkenes. Concepts for personal location privacy policies. In *Proceedings of the 3rd ACM conference on Electronic Commerce.* ACM Press, 2001.

[10] Ubicomp 2002. Socially-informed design of privacy-enhancing solutions in ubiquitous computing: Workshop at Ubicomp'2002, sep 2002.

[11] Ubicomp 2003. Ubicomp communities: Privacy as boundary negotiation: Workshop at Ubicomp'2003, oct 2003.

[12] Mark Weiser. Some computer science issues in ubiquitous computing. *Communications of the ACM,* 36(7):75–84, July 1993.

[13] Alf Zugenmaier, Michael Kreutzer, and Günter Müller. The Freiburg Privacy Diamond: An attacker model for a mobile computing environment. In K. Irmscher and K.-P. Fähnrich, editors, *Proceedings KIVS 2003.* VDE-Verlag, February 2003.

PRIVACY, SECURITY AND TRUST ISSUES RAISED BY THE PERSONAL SERVER CONCEPT

John Light, Trevor Pering, Murali Sundar, Roy Want
Intel Research

Abstract This paper is a survey of user risks associated with the Personal Server concept. The Personal Server concept is used as a surrogate for future mobile devices. The risks are threats involving security, privacy and trust. An overview of the concept is provided, followed by descriptions of three usage models: mobile storage, application server, and beacon receiver. Each usage model description includes a discussion of risks that result from that usage. No solutions are provided.

Keywords: Privacy, Security, Trust, Ubiquitous Computing, Pervasive Computing, Mobile Computing, Personal Server

1. The Personal Server Concept

Among other ends, Pervasive Computing deconstructs the User Interface which has dominated computing for the last two decades. Since the Personal Computer (from Apple and IBM) arrived in the early '80s, User Interface has consisted of a human sitting upright in front of a vertical display surface wielding a keyboard and pointing device on a horizontal surface. This paradigm is unchallenged for "real" computers, but the advent of Personal Digital Assistants and especially cell phones has challenged it in the larger arena.

The Personal Server project [6, 5] explores an extreme alternative approach to this paradigm by asking "What if your computer had no standard user interface?" How would that change what our computers consists of and how we use them? How would the world have to change in order to accommodate us? How would that change how we feel about computing? How would it change the impact computers have on our lives?

To explore these questions we created a mobile device with considerable processing power, storage, battery capacity and communication

capability but no display or input device. It is a fully capable computer without an inherent user interface. We don't expect to see a product built this way, but we hope that what we learn can be applied to building better mobile computing devices of all sorts.

The Personal Server prototype consists of an Intel PX255 processor, which includes Intel XScale®technology, two Compact Flash slots for memory expansion, a Zeevo Bluetooth radio, and a battery capable of running the device for about a day. The prototype is being manufactured and sold by Crossbow Technologies for the benefit of researchers in many disciplines who want a compact, highly capable mobile computing platform. An open source Linux distribution is available on SourceForge to support it. Compact Flash cards with capacities of up to 4 gigabytes are currently being sold, and larger ones have been announced.

The Personal Server is analogous to a personal version of the back-end servers that provide file, web, database and application services to desktop computers. Just as the Personal Computer took the mainframe computer out of the back room two decades ago, and the notebook PC took the Personal Computer out of the office, and the PDA took the PC onto the street, the Personal Server takes the back-end server out of the back room and puts it in the pocket or purse. An important implication of this analogy is that while PCs of all sorts are often turned on and off, servers tend to be "always on", providing services even when the user is not directly engaged. The Personal Server is designed to run all day in the user's pocket, and this is a characteristic it shares with the cell phone.

Capabilities of a Personal Server may eventually be included in some other form of mobile device, such as a Personal Digital Assistant or cell phone, since its physical components are very similar to both.

The immediate questions posed by the Personal Server concept are:

- What computing needs can such a device satisfy?
- What personal needs can such a device fulfill?
- How does one interact with such a device?
- Can interaction with such a device be effective and satisfying?
- Can interaction with such a device be safe?

This paper explores the issues related to last question using our learnings from the other questions.

2. Summary of Issues

Because the Personal Server explores an extreme computing model, it raises unique issues of security, privacy and trust in addition to those present in any mobile device. We expect aspects of the Personal Server to make their way into mainstream products in the future, and the Personal Server project provides a relatively clear view of what those issues may be.

Any mobile device raises concerns about security ("Can someone modify or destroy my data?"), privacy ("Can someone read my data?"), and trust ("Can I count on my data being available when I need it?"). The way these issues manifest themselves depends on the nature of the device, the nature of its use, and the expectations of its user.

The Personal Server concept expands on those issues because of its lack of display and dependence on a wireless connection to the world. For any computer system, the most severe threats involve external communication, and **all** of the Personal Server's operations involve interaction with external sources. Moreover, the Personal Server concept proposes new primary modes of external interaction such as annexing external User Interaction devices and listening to Information Beacons. Annexation raises new questions for secure authentication, and listening to beacons raises new issues of privacy.

This paper summarizes the security, privacy and trust issues uncovered by the Personal Server project. We will not explore issues that are common to all mobile devices, concentrating on those that are unique to Personal Server concept. We hope that this exposition of issues will add to the overall picture [4] of what we need to do to make the Pervasive Computing environment safe.

3. Generic Risks

Any mobile device carries risks involving security, privacy and trust. Solutions to eliminating or mitigating such risks are an on-going effort by the mobile computing community. The Personal Server project assumes that those efforts will be successful and expects to benefit from them. We will survey them quickly to provide a more complete picture of the issues.

At one extreme of the mobile device playing field are the smart card and USB Flash storage device, sometimes called a USB dongle. Both have a primary purpose of carrying information safely from one place to another. Both are implemented with storage and a processor sufficient to interface them to other computing devices, and that is their primary purpose. In one case the storage and device size are very small (smart

card), and in the other case (USB dongle) the storage capacity can be quite large in a package not much bigger. The biggest difference is that the smart card is designed to only talk with trusted readers while a USB dongle can connect with nearly any computer.

At another extreme is the notebook computer. Some are barely mobile, and they typically include large amounts of storage. Most have many I/O mechanisms, but I/O other than the keyboard, display and pointer is usually of secondary importance. The primary purpose of most notebook computers is as a more or less complete, self-contained computing environment. A notebook computer may be just as vulnerable to risks of security, privacy and trust, but many of those risks can be mitigated by working without connection to the external world until a safe venue is attained.

Most mobile devices fit within those extreme, but they all share some common concerns.

- How likely is the device to be stolen?
- How likely is the device to be lost?
- If it is lost or stolen, what is the likelihood that its contents will be stolen?
- If it is lost or stolen, how quickly and easily can it be replaced?
- If it is lost or stolen, how much information and work will I lose?
- Can its contents be stolen during normal usage?
- How susceptible is my interaction with the device to being observed?
- Can someone introduce a malign agent into the device?

Some of these concerns are bigger problems for some devices than others. The likelihood of being stolen is a complex function of perceived value versus perceived risk on the part of a potential thief. A device that is often put down on surfaces is more likely to be stolen or lost. Moreover, the availability of effective (and used) security and privacy technologies can make the loss of data less of a problem. The availability (and use) of backup or synchronization services can mitigate the replacement problem.

The Personal Server and other devices with Personal Server capabilities are vulnerable to these same risks. We expect products with these capabilities will use the best known practices to deal with these and

other generic risks. The rest of the paper discusses risks that are introduced or emphasized by the Personal Server concept, which we will refer to as *incremental risks.*

4. Mobile Storage Issues

The earliest usage models explored on the Personal Server were its use as a file and web server. These are traditional uses of a traditional backroom server, and the Personal Server's wireless file server capability is an obvious extension of the current popularity of USB dongles. Portable storage devices have always held an important place in personal computing, and USB dongles have largely inherited the place once occupied by floppy disks.

The obvious difference in this use with devices such as the USB dongle is the wireless connection provided by the Bluetooth radio. Instead of reaching into your pocket for a USB dongle, fumbling with your computer to plug it in, and trying to remember to take it with you when you leave, you can use the Personal Server while it stays untouched in your pocket. This simplicity of use comes at the cost of some implementation complexity and incremental risks.

One class of incremental risks for the Personal Server involves the nature of wireless connections. When your USB dongle connects to a computer, it is typically obvious what connection has been made: the physicality of the connector ensures the integrity of the connection. A wireless connection, on the other hand, can be ambiguous. How do I know what connection I've made, and how do I know there is not a "man-in-the-middle"? There are no natural physical artifacts to answer those questions.

Any storage device must be able to reliably hold data. A mobile storage device must deal with physical threats to the device, e.g., theft, dropping, losing, etc., which are normally dealt with by some form of synchronization or backup. Furthermore, the normal usage of such a device exposes it to hosts outside of the user's direct control, e.g., a friend's or customer's notebook computer, etc., which exposes it to intentional or unintentional data loss. Some storage devices include a physical switch to write-protect the contents, but such switches are hard to use, so small that few people even know they are there, and unlikely to be used at critical times. They also provide only binary control: if anything is to be written, then all protection goes away.

A mobile storage device should be able to hold data securely. Hard drives typically depend on the physical security of their location to provide data security, but a mobile storage device is more likely to fall into

the hands of someone who wants to steal the contents, through either theft or loss of the device. Furthermore, the normal usage of the device exposes its contents to theft whenever it is connected with a host not directly controlled by the storage device owner. This is true whether the host is operated by the user (a rented computer) or not (a customer computer). Most current devices expose all their contents whenever they are plugged in, and the few with authentication methods expose all their contents after authentication succeeds. Ideally, only the data relevant to a transaction would be accessible at any one time.

A mobile storage device must provide reasonable access to its held data. The word "reasonable" refers to a tradeoff between the user's risk and effort. Security often deals with such tradeoffs, but the need to include untrusted hosts in the security equation makes solutions more difficult. For example, common security methods such as typed passwords are less effective in the common usage model since they expose the passwords themselves to theft. This can lead to more complicated security measures, which may discourage using either the device or the security measures. It is not sufficient to prove that a procedure is secure unless you can also prove that people will use it. This problem encourages the development of alternative authentication methods.

A mobile storage device can act as a vector for worms, viruses and other forms of malware. Because it promiscuously connects to multiple devices and connects quite directly (typically as a mapped file system), it is an ideal vector for malware. All such devices are currently vulnerable to existing viruses, and we expect malware to be written specifically for mobile storage devices as the use of such devices proliferates. Since the current crop of mobile storage devices are seen as big floppy disks, this problem is being treated as a host issue, but it is not practical to scan all the contents of a multi-gigabyte storage device every time it is plugged into a host. The device itself must be involved in supporting the protection process, and the host must be able to trust that involvement.

The Personal Server project has explored solutions for some of these problems, using the device's processing power to counter its vulnerability. For example, we have considered structured availability of data, new forms of authentication [2], and access journaling. The Personal Server can also present its contents in the form of a Web site, which reduces some threats to the Personal Server but not the host. Discussion of these solutions is not within the scope of this paper.

5. Application Server Issues

The processor of the Personal Server allows it to act as an application server. In this case the data for an application is stored on the Personal Server, and a program that implements the application runs there as well. In some cases the application can run with little or no user interface, but in others a user interface is needed. If the Personal Server capability is embedded in a device with a display screen, that screen might be used for the application.

Some applications require a bigger screen than a mobile device can reasonable provide, and some applications involve collaborative use with colocated individuals. In those cases, an external screen might be used with a mobile device. Desktop computer users have had remote access to their machines for years, and we believe this capability may become common with mobile devices as well. Thus, this problem is not limited to the Personal Server model.

The model here is that someone with a mobile device (e.g., a Personal Server) would walk up to a public display, take some action on that public display, and create an interaction session on that display with an application running in the mobile device. For the duration of the session, the user would use the affordances of the public display to interact with the application and the results would be shown on the public display.

Known as *annexation* [3], this use of an external interaction device can provide a larger or shared screen when needed. Several relevant problems arise from annexation.

- How do you know which display you are annexing? This may seem obvious, but if you annex an interaction device that someone else controls, they might steal or destroy your information before you even know there is a problem.

- How do you know the interaction device isn't recording your session? There are lots of nefarious uses for a session recording.

- How do you authenticate yourself to your mobile device without exposing passwords? This problem is common with the previous section on mobile storage devices.

- How do you know that your interaction session is controlling your mobile device? An observer might be able to simulate your typical session with another device (after observing a previous session) well enough for you to be fooled into typing sensitive information into it.

- How do you know there is not a man-in-the-middle passing your interaction through until you have authenticated yourself? The man-in-the- middle may then either steal information or take control.

- How do look at information on a public display without displaying more than you want?

The last question is really a whole class of questions about how we deal with information in settings that are not entirely private. The advent of Pervasive Computing and the transformation of the office are combining to make our work places more communal or public and less private. Our databases and web sites are often not organized according to sensitivity of information so accessing one piece of information often exposes other pieces that shouldn't be exposed. In the privacy of an office this is usually acceptable, but in many other places we would like to work, it is not.

Since the Personal Server is "always on", it can run applications that might operate independently of user involvement. Such software *agents* can recognize context, respond to events, monitor activity, and notify the user, according to the expressed preferences of the device owner. The agent may operate based on external events, and the veracity of those events may be doubtful if the device is under attack. Such agents should be designed to deal with uncertain and false events. More importantly, an agent may be empowered to act externally to the device on the user's behalf, and these actions may need to be performed without user involvement with untrusted external devices. This creates new security challenges.

6. Information Beacon Issues

The wireless capability and "always on" behavior of the Personal Server allows it to act as a receiver for wireless *information beacons.* Information beacons are small wireless transmitters with a relatively small (~10 meter) broadcast radius. They are inexpensive (<US$25), so anyone (store owner, individual, government, etc.) can place them wherever people walk by carrying appropriate receivers. A short repetitive message (~10K bytes) can be received by any receiver as it passes a beacon.

The combination of information beacons and receivers create a new form of location-aware computing, previously described in a workshop at UbiComp 2003 [1]. It requires no central authority for registration, location mapping, or content handling. Instead, the information passes

directly from its source (who owns the information beacon) to its destination (who owns the receiver). The Personal Server can run software agents that process the incoming beacon messages and act on or archive them without direct user involvement.

Any form of location-aware computing raises issues of privacy and trust. We believe the use of information beacons raises fewer such issues than other forms of location-aware computing since it doesn't involve third parties such as cellular vendors or location-database web sites and it doesn't require traceable radio activity on an ongoing basis. Comparing the use of information beacons with other forms of location-aware computing is not in the scope of this paper. We will summarize the privacy and trust issues of this new approach.

Information beacons offer information and services to passing receivers. The information might be as simple as a store description, or it might include a full menu for a restaurant or a coupon for a clothing store. It could offer to sell something to the user, and the transaction might be able to take place immediately. Previous forms of location-aware computing have concentrated on immediate notification of "interesting" events because of the high cost of maintaining and processing significant state in a centralized resource for each user. We believe the cost of handling state can be much lower in a distributed approach. The new approach concentrates on building a personalized location database for the user, providing a useful source of context and state computations and reducing the need for interruptions commonly seen in other approaches. An agent running in the receiver might interrupt the user, but it would be based on considerably more context than is available to some other approaches.

Three classes of privacy and trust issues arise with the new approach. One class involves external tracking. If a user is communicating continuously with a series of information sources as she passes through an environment, software with a global view of the information sources could track her location and path. This is similar to the concern that the cellular network can track you while you carry a cell phone. This problem can be mitigated by avoiding use of a traceable identifier in communications with the information beacons. The problem can be eliminated entirely if the transmissions are entirely unidirectional. That is, if the receiver doesn't have to send any radio message in order to receive the beacon information, then there is essentially no way for the receiver to be tracked.

Another class of issues involves self tracking. As the receiver collects information from beacons, it likely creates a time stamped record of locations in its persistent storage. This record can be a major source of

value for this approach to location-aware computing, but it can also be a risk in the case that a receiver is lost, stolen or subpoenaed. To mitigate this risk, the user should have full and nuanced control over both the collection and retention of such data. By "nuanced" we mean that the user should be able to have detailed control over various aspects of the data collection and retention, not just the ability to enable and disable.

The third class of issues involves user preferences. The agent that responds to beacon messages must be configured to behave as the user wishes. These preferences form a personal database that may be quite sensitive, depending on its contents. A user may want a mobile agent to work in the more personal parts of her life, and the preferences expressed to that agent may be especially sensitive. The point is that metadata may create as much of an incremental risk as data.

The use of information beacons is an exemplar of the class of applications that can be built on an "always on" platform. Any such program that interacts with the outside world via radio, infrared, RFID, etc., is likely to have similar issues with privacy and trust. As with location-aware computing there are often multiple approaches to architecting the system. The architecture that is easiest or most obvious (or appears to have the most revenue potential) may not be the one that offers privacy and trust.

7. Summary

Because the Personal Server defines a new computing model and new usage models, it exposes new risks to security, privacy and trust. Whether the Personal Server as presented here ever becomes a product is not important, but it is clear to us that various capabilities of the concept will become part of other mobile devices. The Personal Server project provides an opportunity for us to identify these risks at an early stage and provide solutions before they are needed. This paper describes what has been learned so far about risks facing any mobile device that incorporates aspects of the Personal Server concept.

References

[1] J. Light, E. Pattison, T. Pering, M. Sundar, and R. Want. Fully Distributed Location-Aware Computing. UbiComp 2003 workshop, 2003.

[2] T. Pering, M. Sundar, J. Light, and R. Want. Photographic Authentication through Untrusted Terminals. *IEEE Pervasive Computing,* 2(1):30–36, 2003.

[3] J. S. Pierce, H. E. Mahaney, and G. D. Abowd. Opportunistic Annexing for Handheld Devices: Opportunities and Challenges. In *Proceedings of the Human-Computer Interaction Consortium,* 2004.

[4] B. Schneier. *Secrets and Lies: Digital Security in a Networked World.* Wiley, 2000.

[5] R. Want and T. Pering. New Horizons for Mobile Computing. In *Proceedings of PerCom,* 2003.

[6] R. Want, T. Pering, G. Danneels, M. Kumar, M. Sundar, and J. Light. The Personal Server: Changing the Way We Think about Ubiquitous Computing. In *UbiComp 2002; Ubiquitous Computing : 4th International Conference,* volume 2498 of *LNCS,* pages 194–209. Springer-Verlag, 2002.

This page intentionally left blank

II

SECURE TRUST MODELS AND MANAGEMENT IN PERVASIVE COMPUTING

This page intentionally left blank

OVERVIEW

Entities in pervasive computing will be required to operate while disconnected from their home network, thus having no controlling authority and no connection to their certification hierarchies through which to determine their identities. Therefore, there is no specific security infrastructure that can be relied on. In many cases, conventional security mechanisms seem inappropriate for handling the dynamic situations arising in pervasive computing. If there is no security infrastructure that can be referred to, it is hard to imagine on what grounds a security relevant decision should be based. In order to provide a framework for reasoning about such decisions in highly uncertain environments, trust management systems have been proposed, which give flexibility in modeling situations in which there is not enough information about the entities. It is for these reasons that researches felt that the approach of trust management fits in nicely with pervasive computing applications.

Although the field of trust management is relatively new, the proposed solutions so far are potentially promising. Thus, the aim is to analyze the current trust management systems and envisage their application in pervasive computing. Any proposed trust management model for pervasive computing should be capable of addressing the confidentiality of information, and to ensure privacy by protecting personal information from malicious users.

To facilitate the exploitation of trust management models in pervasive computing and to visualize their solutions, we encouraged authors to submit papers around this research topic. One of the workshop sessions was devoted to topics on trust models and management in pervasive computing. The research papers included in this section highlight various activities in different academic institutions in the domain of trust research.

The first paper, "The Role of Identity in Pervasive Computational Trust" was presented by Jean-Marc Seigneur of TCD. It shows how identity can be managed in a trust-based security framework. In the proposed framework, the authors mainly consider the issues of identity in pervasive computing, where there is no central authority legitimate for all entities. The paper then goes on discussing how trust-based deci-

sions are taken into consideration with respect to resource accessibility in pervasive environments. Security decisions are taken based on trust determination with a focus on context and identity information of these environments. The paper nicely links trust to context (awareness) and identities in pervasive environments.

The second paper of the session was "Towards a Next-Generation Trust Management Infrastructure for Open Computing Systems", and presented by Yücel Karabulut from SAP. The paper highlights the design and the requirements for the next-generation trust management infrastructure by exploiting other approaches (SPKI, SDSI). The paper also discusses the boundaries and interfaces between security, privacy and trust.

The last paper of this session was "Research Directions for Trust and Security in Human-Centric Computing", presented by Irfan Zakiuddin. The paper investigates some issues and approaches to achieving trust and security in computing that retains a human-centric property. The argument of the authors is that they typically have to trust systems to offer their services without understanding their trustworthiness. In their proposed framework that investigates security issues in pervasive computing, they center their argument on three levels: user, service and infrastructure level.

The presented work, in terms of the above mentioned papers, opened the discussion on their research statements and other related research agenda in trust management within the context of pervasive computing. The participants seconded the attempts of fulfilling the security concerns and requirements in pervasive computing. The papers posed many questions about the implementation side of the models. This motivated the audience to focus on questions regarding specific pervasive applications and they looked for answers and solutions for them.

The main point that came out manifestly from the session was how to advance the developed model. Without a proper evaluation and assessment, there is a lack of evidence about their realistic application into pervasive computing. In this aspect, the next research objective is to answer questions about how pragmatic is to adopt these models. There are also some open questions about the computational costs and how they respond to security attacks. Any attempts to consider these issues would greatly enhance the trust management models.

THE ROLE OF IDENTITY IN PERVASIVE COMPUTATIONAL TRUST

Jean-Marc Seigneur*[1], Christian Damsgaard Jensen[2]

[1] *Trinity College Dublin*
Ireland
Jean-Marc.Seigneur@trustcomp.org

[2] *Technical University of Denmark*
Denmark
Christian.Jensen@imm.dtu.dk

Abstract A central element in the human notion of trust is to identify whom or what is under consideration. In the digital world, this is harder to achieve due to more or less trustworthy technical infrastructure between interacting parties. However, we argue that uncertain identification may enhance privacy protection. We present the role of identity and how identity can be managed in a trust-based security framework, in order to balance these concerns, and present a discussion of our design and implementation choices.

Keywords: Trust, identity, pervasive computing

1. Introduction

Weiser's vision of ubiquitous/pervasive computing [28] will only become true when computing capabilities are woven into the fabric of every day life, indistinguishable from it. The goal is to enhance the environment and help people in their daily activities. However, the current state of the art in pervasive computing does not properly address security and privacy [2]. For example, illegitimate monitoring, can arise in such an environment due to the proliferation of sensor technology. The ability of computing systems to identify and adapt to their environmental context

* This work is sponsored by the European Union through the Information Society Technologies (IST) programme, which funds the IST-2001-32486 SECURE project and the IST-2001-34910 iTrust Working Group.

is called context-awareness [7]. Privacy can be seen as a fundamental human right "to be left alone" [3] or a basic need (according to Maslow's hierarchy of needs [20]) for a private sphere protected against others. Regardless of the definition, different mechanisms have been proposed to protect the privacy of people due to information technology. The most common mechanisms are either legislative or technological, depending on whether privacy is seen as a right which should be protected by law or a need which should be supported by the devices that are used to access the online world. We do not consider the general privacy threat of pervasive sensors but focus on the technological aspects of privacy protection in trust/risk-based security frameworks (TSF), especially techniques to control the dissemination of personal information at the level of identity. It is important that these frameworks maintain a trade-off between privacy and trust [25]. We use TSF in its broad sense: any TSF can be used (even though the TSF being developed in the SECURE [23] project is an example of an advanced TSF). In the human world, trust exists between two interacting entities and is very useful when there is uncertainty about the outcome of the interaction. Trust can be seen as a complex predictor of the entity's future behaviour based on past evidence. Others have shown how trust can be formalized as a computational concept [14, 19]. The aim of the SECURE project is an advanced TSF formally grounded and (re)usable. The basic components of a TSF (depicted in Figure 1) should expose a decision-making component that is called when a requested entity has to decide what action should be taken due to a request made by another entity, the requesting entity.

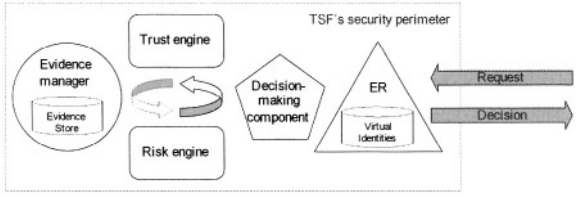

Figure 1. High-level View of a TSF

In order to take this decision, two sub-components are used:

- a trust engine that can dynamically assess the trustworthiness of the requesting entity based on pieces of evidence (e.g., observation or recommendation [27])
- a risk engine that can dynamically evaluate the risk involved in the interaction and choose the action that would maintain the appropriate cost/benefit

In the background, another component is in charge of gathering evidence (e.g., recommendations, comparisons between expected outcomes of the chosen actions and real outcomes...) This evidence is used to update risk and trust information. Thus, trust and risk follow a managed life-cycle. The Entity Recognition (ER [24]) module deals with digital identities and is in charge of recognizing them. We especially put emphasis on ER in the remainder of the paper. The next section contrasts digital and real-world trust, with an emphasis on a key element which is identity. A discussion on the advantages and disadvantages of alternative formats for trust values of entities is given in Section 3. The engineering of identity in SECURE and feedback on the design and implementation choices made is presented in Section 4. Section 5 surveys related work and we draw conclusions.

2. Contrasting Digital and Real-World Trust

In the real-world, rich context is available for trust-mediated decisions. For social scientists [21], there are three main types of trust: interpersonal trust, system trust and dispositional trust. Dispositional trust is said to be independent of any party or context. Interpersonal trust is requesting entity and context specific. So, trust partly depends on context. In computing systems, sources of context are fewer and less certain due to more or less trustworthy technical infrastructure between interacting parties. Dey defines context as "any information that can be used to characterize situation" [5] and emphasizes that not all types of context are equally important. The most important types are: location, identity, time and activity. Time is supposedly the easiest type to get (if there is no misconfiguration or timing-attack). Location is rather new but pervasive computing will provide it. Even though the notion of identity is part of legacy security mechanisms, identification is more or less certain depending on resources spent for security. Capturing the real external activity of the user is still challenging for pervasive context-aware computing [5]. Since we argue for Dey's view on context (i.e., identity is part of context, indeed an important part), we say that the level of trust is computed based on context. This is slightly different from the alternative of computing trust based on identities and then context. More has to be said about the notion of identities in computing systems. Traditionally, users to be enrolled in the administered computing infrastructure are known and what they do electronically is bound to their real-world identity. This allows for the possibility of bringing the faulty user to court. In an open environment (with no unique authority) like the Internet, it is not uncommon to be able to create

as many virtual identities as wanted (e.g., email addresses) with weak links to the real-world identity. Public Key Infrastructures (PKI) with central authorities have not shown their feasibility to legally bind any human with a cryptographic key yet (mainly due to management issues). On one hand, initiatives are needed to solve the problems of managing these multiple and dependable identities [6]. One of the main issues for TSFs in pervasive computing, where no central authority is legitimate, is the fact that it is hard to verify that a sole person has created many identities who blindly recommend one of these entities in order to fool the TSF. The level of trust in the latter entity eventually increases and passes above a threshold which grants the asset. This type of attack is called the Sybil attack [8]. On the other hand, these different virtual identities can be used as pseudonyms, which are privacy enhancing techniques due to their level of indirection between the real-world identity and the electronic data. Trust, as with privacy, is dynamic and evolving interaction after interaction. The intrinsic property of trust to evolve autonomously improves the capability to auto-configure [24]. Privacy is a constant interaction where information flows between parties [13]. Privacy expectations vary [1, 13] and depend on context [15]. So, privacy policies based on context [9, 12, 17, 18] and trust [25] can be made closer to the real-world privacy expectations. However, recalling the process of trust formation makes apparent the fact that privacy is at stake in trust-based systems. In order to be able to trust another entity, the first step is to establish the level of trust in that entity[1], which is the result of an analysis of the existing knowledge and evidence. Thus, trust relies on profiling, where more information is better, because it allows the likely behaviour of the other entity to be more accurately predicted. Any link with the real-world identity of the user changes this information into sensitive personally identifiable information (PII). This is aggravated than in real-life because information is easily stored and retrieved for a long period of time. In Section 4, we present how we engineered identities for pervasive computational trust in order to mitigate these issues. The next section discusses some of the advantages and disadvantages of trust values format, which is a significant difference between real-world trust and computational trust. In the real-world, there is no such well-defined format, which is essential for computing systems to communicate.

3. Trust Values Format: Interoperability, Privacy, Scalability

In the literature, there is no real consensus regarding the digital representation of trust, e.g., the format of trust values of an entity, and

pieces of evidence exchanged between interacting parties. In this section, we look at the format from the point of view of privacy protection, interoperability, ease and accuracy of trust calculation, performance and scalability. Trust values can be more or less expressive (i.e., they contain more or less information): the level of expressiveness seems to depend on the application. However, due to the broad range of applications that can be found in pervasive computing, there should be a context mapping mechanism to adjust trust values calculated in one application to different applications or more generally different contexts. Such a mechanism increases interoperability. More expressive representation is likely to help this mapping. A trust value may be the aggregation of trust values in specific contexts: this helps to exclude trust irrelevant to the context of interest. For example, if there are two applications: one for allowing the requesting entity to drive a car and another one to ride a motorcycle, the trust value is the aggregation of a trust value for cars and a trust value for motorcycles. It makes sense that the trust value for motorcycles can be extrapolated from the trust value for cars, because the same traffic laws apply and the ability to position yourself in traffic is similar for cars and motorcycles. A trust value may simply consist of the inexpressive result of trust calculation due to privacy reasons but it is harder for mapping. In our example, knowing that the trust value for car is 0.6 (which can be the result of many pieces of evidence) is less useful that knowing that the trust value contains the success rate of the driving exam questions also found in the motorcycle exam. A trust value may include these pieces of evidence to facilitate mapping but this may violate privacy (see the latter example). At the other extreme, a trust value may only consist of the pieces of evidence without the trust calculation result, because the trust value calculation can reveal more than the value (e.g., how trust is calculated). Another reason may be that the observers or recommenders are willing to provide objective evidence without wanting to disclose the subjective feeling represented by some trust value calculation. From a performance and scalability point of view, the more trust contexts are aggregated and pieces of evidence are stored in the trust value, the larger the trust value becomes. In fact, performance and scalability are of great concern in pervasive computing where severely resource constrained devices may be found. Large trust values mean that fewer can be stored. Past history of one specific entity may be longer with trust values with more evidence though.

4. SECURE: Feedback on Choices Made Regarding Identity

The current official format of trust values of an entity in SECURE is used within the Trust Information Structure [27]. There are three layers: the bottom layer with the list of pieces of evidence; a middle layer with two types of trust values (trust value due to observations and trust value due to recommendations) to avoid issues related to the use of second-hand evidence; the top layer with combined trust values which are used as the local trust values for the requesting entities. An outstanding choice related to the format of trust values has still to be made. In SECURE, it is possible to query another entity to obtain the trust value of a third requesting entity. This trust value is used as it is provided. This process is called a reference. If the trust value contains an aggregation of trust values related to different contexts/applications, the requesting entity doing the reference can choose to ask either for the full trust value or the part of the trust value of interest. For example, if the request for driving a car is made, the part of the trust value related to driving a motorcycle is not sent in the reference trust value. Again, there is a privacy issue. Requesting for the full trust value is a bigger privacy threat for the entity sending the reference than sending a specific part of the trust value. It is less privacy risky for the entity asking for the reference because it discloses less about what the requesting entity has asked for than if only a specific part of the trust value is requested. An advantage of getting the full trust value is to allow for the best context mapping possible without several exchanges between entities involved in the reference. Due to the notion of reciprocity in privacy concerns, the final choice seems to be in favour of asking for parts of trust values. In doing so, the sending entity knows more about what the requesting entity asked the requested entity for (to compensate the disclosure of its part of trust value) and the requested entity is still able to carry out the decision making. The most appropriate way of referencing may depend on the type of application though. Since the beginning of the SECURE project, the viability of suing any real-world identity has been considered marginal. Our expectation is that entities are in general virtually anonymous to the extent that identity conveys little information about likely behaviour. What is important as a prerequisite is not really "Who exactly does this entity represent?" but "Do I recognize this entity as a trustworthy collaborator?" As there is no a priori information concerning likely behaviour; identity therefore does not imply privilege. Before retrieving trust from the TSF, interacting entities must be recognized. It has been observed that authentication in pervasive computing

systems is not necessarily enough to ensure security, because identity conveys no a priori information about the likely behaviour of the other entity [4, 24]. We have proposed Entity Recognition (ER) [24] as a more general replacement for authentication that does not necessarily bind an identity to the recognised entity (i.e., authentication is a special case of recognition that binds an externally visible identity to the recognised entity). We conjecture that the ability to recognise another entity, possibly using any of its observable attributes, is sufficient to establish trust in that entity based on past experience. Our end-to-end trust model starts with recognition [24], which is a more general concept than authentication, i.e., entity recognition encompasses authentication. To allow for dynamic enrollment of strangers and unknown entities, we have proposed the entity recognition (ER) process, which consists of the following four steps.

1 Triggering of the recognition mechanism

2 Detective Work to recognize the entity using the available recognition scheme(s)

3 Discriminative Retention of information relevant for possible recall or recognition

4 Upper-level Action based on the outcome of recognition, which includes a level of confidence in recognition

From a privacy point of view, this use of virtual identities - pseudonyms (mapping to principals in SECURE) - is a first technological line of defence. In a TSF, the minimum requirement is a local reference for the formation of trust, which is in turn managed by other components in the TSF. According to the privacy protection principle of "collection limitation" [17], data collection should be strictly restricted to mandatory required data for the purpose of the collection. Our requirement is to establish the trustworthiness of entities and not their real-world identity. This is why pseudonymity, the level of indirection between trust and the real-world entity, is necessary. Transaction pseudonyms [15] (i.e., a pseudonym used for only one transaction) and anonymity cannot be effectively used because they do not allow linkability between transactions as required when building trust. There is an inherent conflict between trust and privacy because both depend on knowledge about an entity but in the opposite ways. Although trust allows us to accept risk and engage in actions with a potential harmful outcome, a computational TSF must take into account that humans need (or have the right to) privacy. However, depending on what benefits can be reaped through

trustworthiness, people may be willing to trade part of their privacy for increased trustworthiness: hence, contextual privacy/trust trade is needed. We have proposed [25] a model for privacy/trust trade based on linkability of pieces of evidence. If insufficient evidence is available under the chosen pseudonym, more evidence may be linked to this pseudonym in order to improve trustworthiness and grant the request. Some thresholds should be set concerning the acceptable evidence that should be disclosed. This is why we have introduced the link selection engagement (liseng) algorithm to ensure that the Minimal Linkability principle[2] [25] is taken into account. During a trade process, the following three levels must be balanced: the level of privacy asset of the evidence envisaged to be disclosed; the trustworthiness assessment impact of the evidence to be disclosed; and the utility of the requested action. We have emphasized that care should be taken when linked evidence on multiple virtual identities is assessed. The most important requirement is to avoid counting the same evidence twice when it is presented as part of two different pseudonyms or overcounting overlapping evidence. We found [25] that in some cases, passing recommendations in the form of a simple trust value, instead of all supporting information, does not fulfil the latter requirement. Assessing evidence may require analysis and comparison of each piece of evidence to other pieces of evidence. This is in favour of a trust value format including as fine-grained pieces of evidence as possible. Our initial investigations have shown [25] that combining levels of trust in entities is not uncommon. For example, the outcome of ER can be a set of n principals p (i.e., virtual entity or pseudonym) associated with a level of confidence in recognition *lcr*: $OutcomeOfRecognition = \sum_{i=1}^{n} (lcr_i, p_i)$

When we apply the APER [24] scheme (message-based recognition using cryptographic keys, hashes of previous messages and challenge/responses) to recognise the sender of an email, we may combine the level of trust of principals who were using emails with a text email address and upgrade to emails as APER messages. A tool kit, called the Claim Tool Kit (CTK) [24], has been developed to facilitate the development of message-based recognition. The second scheme, called VER [24], we have been implementing is based on vision recognition: once again principals recognised with different recognition techniques must have their pieces of evidence linked and assessed. To cope with scalability, we have proposed to forget about entities, that the entity has not collaborated with after a certain time or more generally based on context [24, 25].

5. Related Work

One of the main issues for the management of multiple dependable identities is the support of trust levels [6]. We indeed demonstrate in this paper that the SECURE project addresses this issue. Wagealla et al. [27] use trustworthiness of an information receiver to make the decision on whether private information should be disclosed or not, which is another way to envisage the relation between trust and privacy. Kosba and Schreck [15] highlighted the fact that reputation systems do not mandatory require explicit link with real world identities. We added that too much evidence can lead to the disclosure of the implicit link [25]. Others [10, 11, 15] have presented how pseudonyms can be used for privacy protection and shown that different levels of pseudonymity and configurations exist. Their work is valuable to choose the right type of configuration and pseudonymity. Previous work on identity management in ubicomp environments [12, 18] demonstrates that the model of switching identities according to context is appealing and meaningful for users. Our own prototype [24], where pseudonyms are disclosed based on location, confirms the usefulness of context. Different TSFs have been used for sharing personal information in ubicomp environments [9, 26]. However, these TSFs do not use pseudonyms and their focus is not on identity matters. Another related work, although this one only focuses on recommendation, is the OpenPrivacy platform [16]. The user can create many pseudonyms linked with specific information. Langheinrich's work [17] is valuable to understand privacy in context-aware pervasive computing. Robinson and Beigl [22] investigate one of the first real trust/context-aware spaces based on the Smart-Its context sensing, computation and communication platform, which could also be used for an ER scheme based on context.

6. Conclusion

Identity is a central element of computational trust. In pervasive computing, where there is no central authority legitimate for all entities, more or less trustworthy technical infrastructure between parties facilitates attacks (e.g., the Sybil attack) on trust/risk-based security frameworks. However, this weakness can be used for privacy protection. Different alternatives are possible for the implementation of identity in a TSF. There is a trade-off between the aimed level of trust, privacy, interoperability and scalability. We argue for a solution that explicitly takes into account these different levels and so can be used in a diversity of applications (as it can be expected in pervasive computing). We propose the following generic mechanisms to engineer this solution. The

potential weakness of the technical infrastructure is taken into account in our ER process thanks to levels of confidence in recognition. Our privacy/trust trade model includes means to link pieces of evidence of different pseudonymous virtual identities whilst respecting the Minimal Linkability principle. In addition to the fact that identity is a part of context, context-awareness is promising for auto-configuration, privacy protection, interoperability and scalability.

Notes

1. In this paper, we use the following terms as synonyms: level of trust and trustworthiness. In a TSF, they are represented as a trust value. This is different than trust, which is the concept.

2. "No more evidence than needed should be linked."

References

[1] B. D. Brunk, *Understanding the Privacy Space,* in First Monday, vol. 7, no. 10, Library of the University of Illinois, Chicago, 2002.

[2] R. Campbell, J. Al-Muhtadi, P. Naldurg, G. Sampermane, and M. D. Mickunas, *Towards Security and Privacy for Pervasive Computing,* in Proceedings of the International Symposium on Software Security, 2002.

[3] T. M. Cooley, *A Treatise on the Law of Torts,* Callaghan, Chicago, 1888.

[4] S. Creese, M. Goldsmith, B. Roscoe, and I. Zakiuddin, *Authentication for Pervasive Computing,* in Proceedings of Security in Pervasive Computing, LNCS, Springer, 2003.

[5] J. L. Crowley, J. Coutaz, G. Rey, and P. Reignier, *Perceptual Components for Context Aware Computing,* in Proceedings of Ubicomp, LNCS, Springer, 2002.

[6] E. Damiani, S. D. C. d. Vimercati, and P. Samarati, *Managing Multiple and Dependable Identities,* in 7(6), pp. 29-37, IEEE Internet Computing, 2003.

[7] A. K. Dey, *Understanding and Using Context,* in Personal and Ubiquitous Computing Journal, vol. 5 (1), pp. 4-7, 2001, http://www.cc.gatech.edu/fce/ctk/pubs/PeTe5-1.pdf.

[8] J. R. Douceur, *The Sybil Attack,* in Proceedings of the 1st International Workshop on Peer-to-Peer Systems, 2002, http://research.microsoft.com/sn/farsite/IPTPS2002.pdf.

[9] J. Goecks and E. Mynatt, *Enabling Privacy Management in Ubiquitous Computing Environments through Trust and Reputation Systems,* in Proceedings of the Conference on Computer Supported Cooperative Work, ACM, 2002.

[10] I. Goldberg, *A Pseudonymous Communications Infrastructure for the Internet,* PhD Thesis, University of California, 2000, http://www.isaac.cs.berkeley.edu/ iang/thesis-final.pdf.

[11] R. Hes and J. Borking, *Privacy Enhancing Technologies: The Path to Anonymity,* ISBN 90 74087 12 4, 2000, http://www.cbpweb.nl/downloads_av/AV11.PDF.

[12] U. Jendricke, M. Kreutzer, and A. Zugenmaier, *Pervasive Privacy with Identity Management,* in Proceedings of the Workshop on Security in Ubiquitous Computing, 2002, http://citeseer.nj.nec.com/544380.html.

[13] X. Jiang, J. I. Hong, and J. A. Landay, *Approximate Information Flows: Socially Based Modeling of Privacy in Ubiquitous Computing,* in Proceedings of the 4th International Conference on Ubiquitous Computing, LNCS 2498, pp. 176-193, Springer-Verlag, 2002.

[14] A. Jøsang, *A Logic for Uncertain Probabilities,* in Fuzziness and Knowledge-Based Systems, vol. 9(3), 2001.

[15] A. Kobsa and J. Schreck, *Privacy through Pseudonymity in User-Adaptive Systems,* in Transactions on Internet Technology, vol. 3 (2), pp. 149-183, ACM, 2003.

[16] F. Labalme and K. Burton, *Enhancing the Internet with Reputations,* 2001, www.openprivacy.org/papers/200103-white.html.

[17] M. Langheinrich, *A Privacy Awareness System for Ubiquitous Computing Environments,* in Proceedings of the Ubicomp conference, 2002, http://citeseer.nj.nec.com/517334.html.

[18] S. Lederer, C. Beckmann, A. K. Dey, and J. Mankoff, *Managing Personal Information Disclosure in Ubiquitous Computing Environments,* Intel Research, IRB-TR-03-015, 2003.

[19] S. Marsh, *Formalising Trust as a Computational Concept,* PhD Thesis, Department of Mathematics and Computer Science, University of Stirling, 1994, http://citeseer.nj.nec.com/marsh94formalising.html.

[20] A. H. Maslow, *Motivation and Personality,* Harper, 1954.

[21] D. McKnight and N. L. Chervany, *The Meanings of Trust,* MISRC 96-04, University of Minnesota, Management Informations Systems Research Center, 1996.

[22] P. Robinson and M. Beigl, *Trust Context Spaces,* in Proceedings of Security in Pervasive Computing First International Conference, LNCS 2802, Springer-Verlag, 2003.

[23] The SECURE project, Website, http://secure.dsg.cs.tcd.ie.

[24] J.-M. Seigneur and C. D. Jensen, *The Claim Tool Kit for Ad-hoc Recognition of Peer Entities,* in Journal of Science of Computer Programming, Elsevier, 2004.

[25] J.-M. Seigneur and C. D. Jensen, *Trading Privacy for Trust,* in Proceedings of iTrust'04 the Second International Conference on Trust Management, LNCS 2995, Springer, 2004.

[26] B. Shand, N. Dimmock, and J. Bacon, *Trust for Ubiquitous, Transparent Collaboration,* in Proceedings of the 1st Percom conference, IEEE, 2003.

[27] W. Wagealla, S. Terzis, and C. English, *Trust-Based Model for Privacy Control in Context-Aware Systems,* in Proceedings of the 2nd Workshop on Security in Ubiquitous Computing, 2003, http://www.vs.inf.ethz.ch/events/ubicomp2003sec/papers/secubi03_p03.pdf.

[28] M. Weiser, *The Computer for the 21st Century,* Scientific American, 1991, http://www.ubiq.com/hypertext/weiser/SciAmDraft3.html.

This page intentionally left blank

TOWARDS A NEXT-GENERATION TRUST MANAGEMENT INFRASTRUCTURE FOR OPEN COMPUTING SYSTEMS

Yücel Karabulut
SAP Research, CEC Karlsruhe, Vincenz-Priessnitz-Str. 1, 76131 Karlsruhe

Abstract Basically, there are two intertwined kinds of security mechanisms: monitoring including access control and cryptographic protocols. The purpose of an access control system is to enforce security policies by gating access to, and execution of, processes and services within a computing system. Specification and enforcement of permissions can be based on asymmetric cryptography. In order to employ asymmetric cryptography in open computing environments we need appropriate trust management infrastructures that enable entities to establish mutual trust. Management of trust is organized within a public key infrastructure, PKI for short. Credentials assert a binding between a principal, represented by a public key, and some property. Current proposals investigating the definition of PKI and the application of credential-based access control treat existing PKI models (e.g. X.509) and trust management approaches (e.g. SPKI/SDSI) as competing technologies. We take a different position. We argue here that a trust management infrastructure for open computing environments has to use and to link existing approaches. We explain which requirements a next-generation trust management approach has to fulfill. After presenting an application scenario, we finally outline the design of a next-generation trust management approach that we believe really would appear to be worthwhile for a broad spectrum of applications.

Keywords: PKI, certificates, credentials, trust management, access control, X.509, SPKI/SDSI

1. Introduction

The proper administration of computing systems requires to specify which clients are allowed to access which services, and to effectively and efficiently enforce such specifications. In a local computing system, a specification can be represented by traditional access rights granted to

known identified individuals and thereby to the processes under their control. The enforcement is mostly based on identification and authentication of requesting individuals over a trusted physical path and on keeping track of the processes they are controlling. In the Internet most interactions including business transactions occur between strangers, due to billions of spontaneous users and the fact that most of them do not share a common security domain. Thus, Internet constitutes a global computing infrastructure in which entities need to reason about the trustworthiness of other entities in order to make autonomous security decisions. In the modern computing environments [11] emerging from these trends, some basic assumptions of traditional access control approaches are not longer valid. Traditional access control mechanisms operate under a closed world assumption, in which all of the entities are registered and locally known. When the server and the client are unknown to one another and when resources are to be shared across administrative boundaries, the conventional authorization scheme fails. Thus, we cannot reasonably assume anything like a trusted physical path between remote agents. In order to overcome these and related difficulties a diversity of proposals has arisen. While all proposals exploit cryptography, some of them use symmetric cryptographic mechanisms, like Kerberos [12], and others rely on asymmetric cryptography, like X.509 [10] and SPKI/SDSI [9]. Accordingly, we can specify and enforce permissions of clients on remote servers by employing modern access control approaches which are based on asymmetric cryptography. In order to employ asymmetric cryptography in open computing environments we need appropriate trust management infrastructures that enable entities to establish mutual trust. Management of trust is organized within a public key infrastructure, PKI for short. Credentials are digital and digitally signed documents that assert a binding between a principal, represented by a public key, and some property. Current literature treat existing PKI models and trust management approaches as competing technologies even as dueling theologies [4]. We take a different position. We argue that a trust management infrastructure for an open and dynamic computing environment has to use and to link existing PKI models. Accordingly, we designed a hybrid PKI model to be used for specifying and enforcing permission in open computing systems. The hybrid PKI model, as reported in [2, 3, 8], unifies and extends previous PKI approaches [9, 10]. The sole purpose of this position paper is to stimulate discussion in a workshop on security and privacy in pervasive computing. In particular, it is not our goal here to put forth new results and proposals. All of the technical material alluded to here has been developed in previous work [2, 3, 5–8].

2. Thoughts on a Next-Generation Trust Management Infrastructure

2.1 An Application

A typical scenario exploiting the use of credentials for access control runs as follows. A client is represented by (one of) his public key(s) and characterized by the assigned properties. A resource owner follows a confidentiality policy that is expressed in terms of characterizing properties. An agent as resource owner receives a signed request together with a set of credentials stemming from the pertinent client. The agent firstly ensures the authenticity with respect to the bound public keys and with respect to the actual holder of the corresponding private key by applying appropriate challenge-response protocols and secondly evaluates his trust in the signing issuer. Then the agent decides on the permission of the request by evaluating the properties extracted from submitted credentials with respect to his confidentiality policy. Depending on the application and the underlying trust relationships between the involved entities, such scenarios can be realized by employing different PKI models and trust management approaches. We see arguments of the style this-model-is-better-than-another-model. PKI trust relationships must be built on real-world trust relationships. In many real-world scenarios, trust relationships consist of hierarchies, trust networks, and combinations of two. Therefore, we argue that a trust management infrastructure, as required by dynamic computing environments, has to use and to link both kinds of PKI models. More concretely, we consider the following scenario. In [1], we proposed a secure information integrating mediation approach (i-mediation for short) considering the dynamics and conflicting interests of mediation participants. In mediated information systems, a client seeking information and various autonomous sources holding potentially useful data, are brought together by a third kind of independent components, called mediators. Data sources in i-mediation, following property-based security policies, aim at supporting a wide range of potential clients, which are in general unknown in advance and may belong to heterogeneous and autonomous security domains. This raises the challenge how remote and autonomous entities can agree on a common understanding of certified properties, and other issues related to these properties (e.g. encoding formats). In such situations the sources wish to be assisted to determine potentially eligible clients. To reach potentially eligible clients, which might belong to remote security domains, the sources will need to trusted mediating agents having the required domain expertise as well as the relationships with the potential clients. As a concrete solution, we proposed

an additional mediation functionality, called entity finding mediation, f-mediation for short. F-mediation employs our hybrid PKI model (see Section 2.3 and [2]).

2.2 Outline of the Infrastructure

In [2], we classified previous PKI approaches as based on trusted authorities with licensing and dealing with free properties (characterizing attributes including identities) and the corresponding certificates , e.g. X.509, or based on owners with delegation dealing with bound properties (including capabilities) and the corresponding credentials , e.g. SPKI/SDSI. We extended and integrated these approaches into a hybrid PKI model which uses protocols to convert free properties into bound properties. Furthermore, we unified licensing and delegation by introducing administrative properties. An instance of the full hybrid PKI model consists of overlapping components of three kinds: a) trusted authorities (also called trustees) and licensees for and a holder of a free property together with a verifier of this free property, b) an owner and delegatees for and a grantee of a bound property, and c) a holder of free properties and a grantor of a bound property. The grantor follows a property conversion policy that maps free properties on bound properties, where the property conversion policy is a part of grantor's whole security policy. More precisely, the property conversion policy specifies which set of free properties an entity has to enjoy in order to obtain a bound property assignment. A typical interaction for a property conversion process runs as follows: A holder of free properties requests a promise for a permission, i.e., a bound property. For this purpose, the holder shows her certified free properties and applies for a bound property from the grantor who is acting as an authorizer on behalf of and in explicit delegation of a resource owner. The grantor, after verifying the submitted free property-certificates with the supporting licences, applies his conversion policy on the free properties extracted from the submitted certificates, and finally, if all checks have been successfully completed, grants a bound property-credential where the subject (grantee) is the same as in the submitted free property-certificates. Our hybrid PKI model brings together different PKI models and trust management approaches. The business advantage of such a model is clear. By employing a unifying PKI model, which provides a seamless interoperation between heterogeneous and autonomous security domains, organizations can broaden their potential customer base and collaborators base.

2.3 Required Features

As a basis for emerging distributed applications which aim to follow credential-based access control policies, we would like to see the following features supported by a next-generation trust management infrastructure that enables interoperability between heterogeneous security domains:

- support for free properties (e.g. personal data, a skill, group membership)
- support for bound properties (e.g. a ticket, a capability, a role)
- conversion of free properties into bound properties
- the model of trusted authorities with licensing (e.g. X.509)
- the model of owners with delegation (e.g. SPKI/SDSI)
- support for administrative properties (e.g. trustee, licensee, delegatee)
- recursive trust evaluation (e.g. path validation, chain reduction)
- expressive certificates or credentials
- expressive authorization policies supporting role-based access control
- authorization decision engines
- credential management components (e.g. issuing, revocation)

In addition to these features, the anonymity need of the clients has to be considered. While requesting accesses to the resources, clients may be unwilling to reveal their identities for private reasons and thus prefer to remain anonymous. Additionally for a resource owner, it may be necessary to see evidences of a client's eligibility rather than to know who they are. Thus, the trust management infrastructure should support concepts (e.g. pseudonyms) to support anonymity of the clients.

Acknowledgments

It is a pleasure to thank Joachim Biskup with whom I've had extensive discussions about trust management, PKI models and secure mediation.

References

[1] C. Altenschmidt, J. Biskup, U. Flegel and Y. Karabulut: Secure Mediation: Requirements, Design and Architecture. Journal of Computer Security, 11(3):365-398, 2003.

[2] J. Biskup and Y. Karabulut: A Hybrid PKI Model with an Application for Secure Mediation. In 16th Annual IFIP WG 11.3 Working Conference on Data and Application Security, pages 271-282, Cambridge, England, July 2002. Kluwer Academic Press.

[3] J. Biskup and Y. Karabulut: Mediating Between Strangers: A Trust Management Based Approach. In 2nd Annual PKI Research Workshop, pages 80-95, Gaithersburg, Maryland, USA, April 2003.

[4] B. Chinowsky: Summary of the panel discussions Dueling Theologies. In 1st Annual PKI Workshop, Gaithersburg, Maryland, USA, Apr. 2002.

[5] Y. Karabulut: Investigating Trust Management Approaches for Trustworthy Business Processing for Dynamic Virtual Organizations. Special Session on Security and Privacy in E-Commerce within the 7th International Conference on Electronic Commerce Research, INFOMART, Dallas, June 2004.

[6] Y. Karabulut: Developing a Trust Management Based Secure Interoperable Information System. Special Session on Security and Privacy in E-Commerce within the 6th International Conference on Electronic Commerce Research, INFOMART, Dallas, October 2003.

[7] Y. Karabulut: Implementation of an Agent-Oriented Trust Management Infrastructure Based on a Hybrid PKI Model. In 1st International Conference on Trust Management, LNCS 2692, pages 318-331, Crete, Greece, May 2003.

[8] Y. Karabulut: Secure Mediation Between Strangers in Cyberspace, Ph.D. Thesis, University of Dortmund, 2002.

[9] SPKI/SDSI. http://theworld.com/ cme/html/spki.html.

[10] X.509. http://www.ietf.org/html.charters/pkix-charter.html.

[11] A. S. Tanenbaum and M. van Steen. Distributed Systems: Principles and Paradigms. Prentice Hall, Upper Saddle River, NJ, Sept. 2002.

[12] B. Neuman and T. Ts'o: Kerberos: An Authentication Service for Computer Networks. IEEE Communications, 32(9):33-38, Sept. 1994.

RESEARCH DIRECTIONS FOR TRUST AND SECURITY IN HUMAN-CENTRIC COMPUTING*

Sadie Creese[1], Michael Goldsmith[2], Bill Roscoe[3], Iran Zakiuddin[4]

[1] *Systems Assurance Group,*
QinetiQ, Malvern Technology Centre, UK
S.Creese@eris.QinetiQ.com

[2] *Formal Systems (Europe) Ltd.*
www.fsel.com
michael@fsel.com

[3] *Oxford University Computing Laboratory*
Bill.Roscoe@comlab.ox.ac.uk

[4] *Distributed Technology Group,*
QinetiQ, Malvern Technology Centre, UK.
I.Zakiuddin@signal.QinetiQ.com

Abstract Pervasive networks foresee communicating and computing devices embedded throughout our environment. This will cause huge increases in the complexity of network infrastructures and the information services available over them. The challenge of managing information services, while maintaining security and privacy will be great. It is not clear that current security paradigms will map readily into such future environments. This paper outlines the authors' current position regarding the technical challenges which will need to be addressed in order to make secure pervasive computing environments a reality.

*This research is being conducted as part of the FORWARD project which is supported by the U.K. Department of Trade and Industry via the Next Wave Technologies and Markets programme. www.forward-project.org.uk

1. Introduction

The ubiquitous paradigm foresees devices capable of communication and computation embedded in every aspect of our lives and throughout our environment. This will increase both the complexity of information infrastructures and the networks which support them. New forms of interaction are envisaged, which will aim to push the technology into the background making the information services human-centric in delivery. Computing devices will be less and less noticeable, creating a feeling of being surrounded by "ambient intelligence".

As these pervasive computing technologies become deeply intertwined in our lives we will become increasingly dependent on them, implicitly trusting them to offer their services without necessarily understanding their trustworthiness. Undoubtedly the timely provision of bespoke services will require personal or valuable data to be digitally stored and made available. The increased digitalisation of our assets, coupled with the increasingly intangible way that networks use information, will make it difficult to ensure that trusted services are indeed trustworthy. Will users have to decide how to interact with systems without understanding the associated risks?

This paper presents our thoughts on a particularly important, often critical, property that will be required of such systems, namely *Information Security.* We consider both the technical requirements for secure pervasive computing environments and the human centric properties users are likely to demand of such systems. We highlight the issues we feel require addressing by the research community. The thoughts that we present in this short article are guided by our previous work on pervasive computing security: [2] and [3].

2. Challenges to Information Security

The concept of *authorised access* is enormously important to security, underpinning most principal security properties:

- **Confidentiality.** Information is only made available to those who are authorised to have it.
- **Integrity.** Only authorised users may manipulate information.
- **Availability.** Information services must be accessible to those authorised.

Underpinning the notion of authorisation is that of *authentication,* which concerns proving the validity of an authorising claim. Traditional notions of authentication concentrate on the notion of proving the claim of

an identity (if identity can be proved, then this is a basis for authorisation). In [2] we provided a critique of traditional identity authentication, arguing its unsuitability for pervasive networks because:

- Interaction would be between devices and it does not seem plausible that the identity of an arbitrary device, in an arbitrary environment, can be reliably determined. Furthermore in some applications mass-produced devices might not have unique identities.
- The value of authenticating an identity depends on the trustworthiness of the owner of the identity. If we do not know, either beforehand or by other means, that the owner of the identity is trustworthy, then little is gained by authenticating that identity. Thus, simply proving the identity of a device would be of limited value, since it provides little assurance that the device will behave in a trustworthy manner.

There were subsidiary reasons for doubting the value of identity authentication, such as the viability of certification infrastructures to support authenticating the identities of the huge numbers of devices that are likely to exist.

After presenting the above deconstruction we proposed that authentication for pervasive computing is revised to mean *attribute authentication.* Any device will have a range of attributes, such as its location, its name, its manufacturer, aspects of its state, its service history, and so forth. In a given situation some attributes will need authenticating and the attributes should be chosen to achieve assurance about *which* devices are the subject of interaction, and *what* those devices will *do.*

Protocols for authentication and authenticated key exchange have been the subject of intense study [1]. Moreover, the subject of verifying such protocols has achieved significant advances [5]. For analysis and formal verification it is vital to be precise about the threat model which a given protocol must resist. The standard model of the attacker is due to Dolev and Yao [4], and it underpins a large portion of the research community's efforts. However, the Dolev-Yao threat model (as it is referred to) significantly predates the promulgation and widespread acceptance of the pervasive computing vision. In [3] we proposed that such a threat model was too simplistic and unable to capture the authenticated key agreement protocols that might be required for pervasive networks. The principal amendment was to propose a "two-channel" threat model, as follows:

1 An *E*-channel which captures human or other "external" participation in bootstrapping an authenticated link. On the one hand,

compared to the Dolev-Yao model, the attacker's capabilities on the *E*-channel were significantly limited. But on the other hand the bandwidth for communication on the *E*-channel is assumed to be small.

2 An *N*-channel which captured the main medium for devices to create an perform secure electronic communications. The attacker would have similar capabilities to the Dolev-Yao attacker on the *N*-channel, but the bandwidth for communication is much greater.

A successful protocol for initialising a secure link in pervasive networks depends on sound use and interaction of the two channels. Our understanding of the literature to date has led us to believe that the two-channel threat model is powerful abstraction capable of formalising a wide range of protocols.

Pervasive computing frequently makes the *E*-channel available thanks to the locality and context-dependent nature of authentication. And it may be *necessary* to use the *E*-channel due to the potential lack of ubiquitous PKI services and useful device identities.

Our two papers indicate how fundamental security parameters will change, as information services become pervasive. Both point to an increase in the breadth and heterogeneity of the problem space. Instead of authenticating identities, we may be obliged to authenticate any of a very wide range of attributes; and instead of the standard Dolev-Yao threat model, we have a matrix of threat models. This broadening of the problem space clearly indicates that ubiquitous, human-centric computing will make the problem of achieving trusted and trustworthy information services harder.

To structure our understanding of the broader problem space and to help organise discussion, we propose that the subject is factored into three sub-domains:

- **User Level.** This includes all the involvement of human users in achieving, violating or enabling the violation of security. It also includes the design of user interfaces. The user interfaces will themselves connect this level to the service level.
- **Service Level.** This level encompasses all applications, though our interest is primarily in security applications. The service level will make use of information resources and computing and processing capabilities offered by the infrastructure level.
- **Infrastructure Level.** This level contains the hardware present in the pervasive networks, the information resources, the communi-

cations architectures, the middle-ware and the software processing architectures.

Commonly, when such a layered factorisation is proposed, there is much debate and argument about the number of layers, the contents of the layers and so forth. In this case such debate would miss the point: we do not prescribe this layered decomposition as canonical. It is merely a conceptual tool inspired by the fact that achieving trust in human-centric computing will not be possible without a careful consideration of the humans' role – thus the User layer to make this explicit. Furthermore, the delivery of ambient intelligence services will require a range of resources, communications and computing capabilities that will be globally standard and locally available – thus the need for an infrastructure layer. It should also be noted that the two-channel threat model, that we summarised above, implies two layers. Given that successful abstraction, we hope a layered decomposition of the problem will be a fruitful way to proceed.

3. Future Research Challenges

This final section contains an outline of some of the important issues that we feel need to be debated and understood, arranged according to the loose layering that we mentioned in the previous section.

3.1 User Level

1 The human's role in achieving trust needs to be clearly understood. The security requirements in the examples in [2], and the new modelling paradigm, in [3], derived from what assurances a human with a sound knowledge of information security would seek. In implementing them we made use of things a human would be willing and able to do to achieve these. The important question here is this: to implement and achieve trustworthy interaction should the broad strategy be to minimise the human's role, or should it be assumed that humans can and should retain significant ownership of protecting their assets. Arguments for retaining the human's role include:

- the fact that people do care about their assets (and will continue to do so as they are digitised);
- people want to retain ownership of whatever they regard as precious; and
- the fact that people increasingly use electronic security mechanisms, especially PIN numbers.

Arguments for minimising the human's role are:

- the difficulty of designing trustworthy and effective human-computer interfaces;
- the general fact that most security violations involve irresponsible use or management by people;
- the fact that PIN numbers are frequently poorly managed and stolen; and
- the desirability of relieving the human user of tasks which might become very frequent and burdensome, or be necessary when the human is not in a position to do them.

This is clearly a fundamental question, but it may not be necessary to understand it as an exclusive choice.

2 With regard to the problem of enabling users to retain control of who and what they trust, this seems to define a whole service category of decision support tools. The tool will inevitably have some measure of control over the decisions that its owner makes. How well understood is the science of making such tools trustworthy? Clearly, for such tools to be effective their interface to the user must itself be effective. How well understood is the science of making the interface to such tools trustworthy?

3 Conversely, if the aim is to minimise the user's role in implementing security, then it should first be noted that this may make the problems of responsibility and liability harder.

4 The work in [2] and [3] lays the groundwork for understanding what a human security expert might require, and what is needed to establish authentication in pervasive environments. When using the concept of weakened Dolev-Yao channels as suggested in the second of these papers, it is important to investigate ways in which these channels can be realised both with, and more importantly without, human participation.

3.2 Service Level

1 An interesting issue arises whether tools and technologies for trust and security enable their users, or act on behalf of their human owners. This is whether to allow certain actions to go ahead in the presence of incomplete information. Of course the greater the importance of security, the less we will be inclined to do this.

2 Decision making about trust and security might be enabled, if it were possible to "quantify trust". This is hard, and any scheme will be prone to criticism, but the definition of trust as an "acceptable level of risk" might provide one basis for a way ahead. The notion of attribute authentication might provide an appropriate setting for trying to quantify trust and make decisions about acceptable risk. In any instance it is unlikely that it will ever be possible to attain *complete* assurance about all the relevant attributes of the devices involved. Whether the user has control or not, a decision, based on incomplete information, will have to be taken about an acceptable level of risk.

3 However well we define our interface, we may still need to provide an underlying service which supports an appropriate authentication policy, depending on the context of use. Such a service or application would have to be able to tolerate heterogeneous user interaction, and still provide reliable security. However, would such a tool be considered trustworthy by users, given that it might be able to change the users command if it felt the user were mistaken. In addition, what data needs to be provided to such applications in order that they can provide the user with appropriate decision support regarding authentication policies?

4 Essential to pervasive computing will be the ease with which users can traverse distinct networks. This will require unparalleled levels of interoperability on the application level, heterogeneous devices and users will need to interact with a range of trust and security mechanisms. How can we enable such interoperability? Should we be subscribing to the top down approach of generating one standard, or ontology, to which all services subscribe, such as that being developed in the SWAD-Europe project[1]? Is there a real alternative?

5 Where users and devices fail to authenticate should we provide services for broadcasting that fact, equivalent to revocation lists? If authentication means attribute authentication, then what would be the form and content of such "attribute revocation lists"?

3.3 Infrastructure Level

1 We have discussed, above, the two-channel abstraction [3], where the *E*-channel involves physical interaction and is critical to bootstrapping authentication. Typically the implementations of such channels will require things like physical contact, line-of-sight in-

teraction, human intervention, and so forth. In any particular case the reliability of this channel will be crucial, and should be the subject of debate.

2 It is likely that any channel which assumes a weakened Dolev-Yao threat model will rely on (relative) contextual information about the processes using it. Therefore, we might regard such channels as a nice abstraction of the idea of context sensitivity.

3 Most authentication mechanisms currently rely on asymmetric encryption, which is computationally expensive, and requires larger keys – thus consuming more bandwidth. For pervasive computing, where many devices will be relatively weak in their computational and communication capabilities, it is highly desirable to find authentication mechanisms based on symmetric encryption, or one-way functions. Furthermore, the domination of asymmetric cryptography has, in part, been spurred by the need to implement identity authentication. Can attribute authentication provide the impetus for developing and deploying cheaper encryption techniques for authentication?

4 Will the trusted computing initiative[2] bring about solutions for supporting the authentication of device behaviours (these being some of the key attributes that will need authenticating)? If devices are reconfigurable in the field, then they are not necessarily the same as when they left the factory. What is the impact of this? Can we achieve "biometrics" for devices on which we could base our authentication of behaviours.

5 How can major global technology initiative, such as Grid computing[3] and Semantic Web[4] provide the information, computing and communication resources, to enable solutions to trust and trustworthiness in human-centric computing?

6 Finally, what can we do without infrastructure? Or, more precisely, what do we do fundamentally need, and what can we "create" spontaneously, on an on-demand basis?

Notes

1. www.w3.org/2001/sw/Europe/
2. www.trustedcomputing.org
3. www.gridforum.org
4. www.semanticweb.org

References

[1] Colin Boyd and Anish Mathuria. *Protocols for Authentication and Key Establishment.* Springer-Verlag, 2003.

[2] S. Creese, M. H. Goldsmith, Bill Roscoe, and Irfan Zakiuddin. Authentication in Pervasive Computing. In D. Hutter, G. Müller, W. Stephan, and M. Ullmann, editors, *First International Conference on Security in Pervasive Computing,* volume 2802 of *LNCS.* Springer-Verlag, 2003.

[3] S. Creese, M. H. Goldsmith, Bill Roscoe, and Irfan Zakiuddin. The Attacker in Ubiquitous Computing Environments: Formalising the Threat Model. In *Formal Aspects of Security and Trust,* Pisa, 2003. Springer-Verlag.

[4] D. Dolev and A. C. Yao. On the Security of Public Key Protocols. *IEEE Transactions on Information Theory,* 29(2), 1983.

[5] P. Y. A. Ryan, M. H. Goldsmith, S. A. Schneider, G. Lowe, and A. W. Roscoe. *The Modelling and Analysis of Security Protocols: the CSP Approach.* Addison-Wesley, 2001.

This page intentionally left blank

III

EVIDENCE, AUTHENTICATION, AND IDENTITY

This page intentionally left blank

OVERVIEW

The creation, storage, verification and destruction of evidence are all considered very sensitive matters in business and law. As Pervasive Computing bridges technical and social spaces, the convergence of concerns is becoming more apparent. Evidence is created both explicitly and implicitly; a user may request a receipt or ticket after payment in order to later claim the usage of a service or entry to a particular domain. Receipts and tickets are therefore explicitly created forms of evidence, whereas other forms of evidence that are created based on observation, analysis or incidental recording, and can be however reconstructed to prove the transpiring of some event, are considered implicitly created. These implicitly created forms of evidence have become more common in pervasive computing scenarios. These can be a hazard for intrusion of privacy yet they may enhance usability of applications or rightfully help to identify and incriminate subjects acting contra to policies of a particular domain. Therefore, evidence is initially created to record that an event has occurred but at some point in the process it may become equivalent to an identity - given that there is some context-bounded function that renders it unique.

An identity is the fundamental requirement for a subject to gain special access to a service or protected object and subsequently right to use this object. The identity of a subject may range from “Jack Jackson” to “jack@jackson.corp” to “the leader of Jackson Corp” to “owner of black Mercedes”. The actors in the application will therefore need some means of interpreting the evidence made available in order to formulate some assertion of identity. The verification of this identity is also tied to the stakeholders in the application and their contacts. For example, Jack Jackson may prove to be “owner of black Mercedes” by showing his car key or in a more formal scenario, showing his insurance and registration certificates as proof of ownership. However, he may not want to reveal these certificates or driver’s history to the local carwash. In everyday social situations, humans are very capable of adapting their identity and personal data disclosure; however, this is more of a challenge for digitally stored data and extensive internetworking.

Pervasive computing suggests more automated, less obtrusive, flexible and dynamic means of humans interacting with even more distributed and invisible computational services, such that entering a password or retrieving the relevant certificate from a mass of attribute data stored in a local file system tends not to scale. Certainly these conventional means of authenticating identity and attributes will not become defunct overnight; rather we will continue to see network domains federate through mediators brokers or even as peers, in order to reduce the amount of sensitive data that needs to be transferred between domains but also the amount of passwords, credentials and authentication requests that a user will have to remember, store and respond to respectively. This was the initial, publicized motivation behind Microsoft Passport, which was not entirely accepted because of the inherent need for users to trust the secure management of their personal data to passport-providing services typically owned by Microsoft. However, a more open set of standards are emerging and being supported by the Liberty Alliance, which is a consortium of large Telcos, system integrators, application and security service providers. The keyword "Federated Identities" suggests that users may "link" elements of their identity between accounts without centrally storing all of their personal information (see http://www.projectliberty.org/).

This chapter does not seek to present a comprehensive overview or comparison of standardization activities in the Federated Identities or Enterprise Applications domain, but uses this as a means of understanding the core issues when security infrastructures are to be as distributed as the people, applications and resources which comprise pervasive environments. The title of "Evidence, Authentication and Identity" suggests that the parameters submitted to and exchanged between authentication and authorization servers will be more and more generalized and complex types. Evidence can be virtually anything that a particular application context deems as relevant to asserting the identity of a subject, and hence allowing them to be authenticated.

Hoffmann emphasizes this need for a balance between disclosure of user-related information and service provisioning by describing a user-centric Identity Management Architecture, which enhances the privacy of the user by making informed decisions of which element of personal identity should be disclosed for a particular service. In his presentation during the workshop he used a business trip scenario as his motivation. His claim was that a personal software agent, which only requires the minimal input from the user, could eliminate the complexities of planning the travel itinerary. This minimal input would be a calendar entry for example. Many questions arose in the discussion including the local-

ization of the agent, specification and management of privacy policies, control of the agent, and automated checking of identity validity.

Spahić and colleagues from Freiburg went on to discuss why configuration-less authentication methods are required in Pervasive computing. Their solution surrounds the notion of "Pre-authentication", which is based on wireless technology and the use of context information. The basic concept is that there may be some evidence of familiarity between principals, which precedes full authentication. For example, the fact that two people are on the same bus provides them with a situation-constrained channel to start an interaction. Spahić presented a medical scenario as their motivating application but this raised some discussion - for example, how useful is the trust-security tradeoff for an emergency medical scenario? There are therefore some applications where more relaxed security is desirable yet others where the preference of usability cannot prevail.

Robinson followed this context-based security trail of dissertation by addressing the issue of access control based on the evidence that a situation is controllable. The fundamental idea is that there are situations where a subject will give up temporary/transient control of a subset of its resources, for the duration of a particular situation. That is, the perception of security requirements sometimes gives way to the need for a service. Discussion went along the lines of how "zones of control" could be realized, requirements for authorization to such a zone of control and the issues concerning overlapping of virtual and physical security.

This page intentionally left blank

USER-CENTRIC IDENTITY MANAGEMENT IN OPEN MOBILE ENVIRONMENTS

Mario Hoffmann
Fraunhofer Institute Institute for Secure Telecooperation, Germany

Abstract Two levels of identity management can be determined. The first level considers Enterprise Identity Management which is currently on the roadmap of most companies dealing with huge knowledge bases of employees and/or customers. At this level identity management means (1) to provide employees with role-based access to documents and resources and (2) to consolidate and concatenate partial customer identities for simplifications in customer administration.

Almost at the same time the second level of identity management occurred. Personalised context-aware services have begun to enter, particularly, the mobile communication market and, obviously, detailed user profiles are essential to provide reasonable personalised services. These services are based on the user's current location, his environment, and personal preferences. Here, identity management becomes a key technology in order to keep those additional information under control. However, this pursuit of control finally leads to severe implications due to privacy violation.

Hence, a third level of identity management has to be introduced: User-centric Identity Management. User-centric identity management allows the user to keep at least some control over his personal data where several different approaches in this paper have to be discussed. Specifically, a framework will be described which adds user-centric identity management to a context-aware mobile services platform. This platform has been already designed to support and dynamically combine services especially of small- and medium-sized independent service providers.

Keywords: Privacy Protection, Multilateral Security, Identity Management, Context-aware Services, Location Based Services, Web Services, Multi-Agent-Systems

1. Introduction

With the roll-out of UMTS and public WiFi hotspots in several European countries the usability and acceptance of Location Based Services

(LBS) will finally succeed; the higher bandwidth promises Quality of Service (QoS) for video and audio streaming. Moreover, after the development and evaluation of LBS since the hype in 2000 the market is expecting a high penetration of so-called context-aware services in the next step. Whereas, context comprises not only the user's location but also the current time, the user's environment (in terms of additional sensor information based for example on RFID [7], the corresponding preferences, and the user's service history [6].

On the one hand, in order to provide specific personalised value-added services the collection, the analysis, and the management of user related information is mandatory. The more service providers know from their customers the more precise they fulfil and predict the user's needs. Basically, identity management is assumed to be the key technology to bring together, consolidate, and analyse available information and partial identities of users and costumers automatically [1].

On the other hand, especially, in "old Europe" protecting one's privacy is not only a discontinued model, although Sep 11th also has left after-effects in keeping telecommunications and the Internet under specific surveillance. However, keeping the balance between security requirements and privacy constraints of all involved parties is still one of the key concepts and paradigms in our Institute's research work. Thus, so-called multilateral security [5] serves as the basic security concept of a framework for context-aware mobile services. Especially, this framework enables corresponding platforms to realise user-centric identity management [2].

2. A Privacy Enhancing Framework

In principal, the overall architecture of the framework can be divided into three parts representing three different areas of trust that will be discussed in detail in the next paragraphs.

2.1 The User's Mobile Device

The first area comprises the user's mobile device which contains small databases of frequently used contacts and addresses, and a qualified calendar view instead of the hole history of dates and appointments like in current implementations. So-called privacy enhancing services take care of the user's authentication and authorisation in case of accessing personal settings, preferences, and the *homebase.* Any kind of personal information is securely maintained by taking advantage of Smartcard technology and biometry. Each class of mobile devices is supported by a resource dependent user interface, i.e. taking into account different dis-

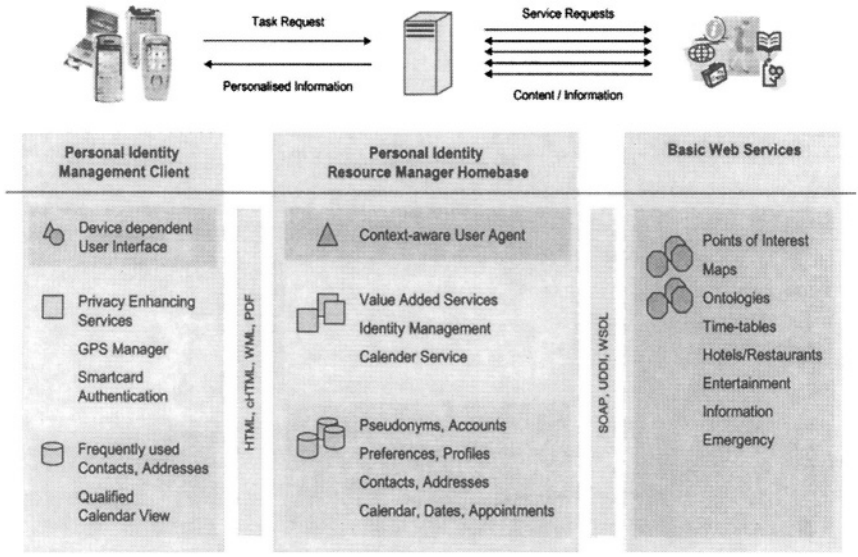

Figure 1. User-centric Identity Management Framework Architecture

play resolutions, browser capabilities, and user preferences. The mobile device with an enabled *Personal Identity Management Client* is considered as providing enhanced control mechanisms.

2.2 The Personal Homebase

In order to get access to context aware services the user triggers the corresponding user agent residing at the *Personal Identity Resource Manager Homebase,* the second area of trust. The homebase is, first, characterised by databases containing the hole set of the user's pseudonyms, preferences, contacts, and calendar entries. Moreover, value-added services consolidate incoming content from *Basic Web Services,* an identity manager balances grant and rejection to personal identity information, and a calendar service manages dates, appointments, and their dependencies. Finally, the homebase is featuring a context-aware user agent which coordinates the activities between the different modules and prepares all results and service information depending on the user's current mobile equipment.

Control mechanisms and, thus, the level of trust highly depends on the location of the homebase. There is several opportunities: First, the homebase might be managed by the user himself on a personal server, e.g. the PC at home. However, this solution assumes a highly skilled and

experienced user. Second, the user might take advantage of a personal homepage administrated by his company - many companies offer such a personal service for their employees. Third, Internet Service Providers (ISPs) specialised on managing and maintaining user identities could offer an appropriate service. Fourth, according to ISPs, Mobile Network Operators (MNOs) might take advantage of their huge subscriber community and could provide enhanced identity management services. Two well known approaches already introduced are Passport by Microsoft and the concept of federated identities by the Liberty Alliance Project [3, 4].

2.3 The World Around Us

However, nowadays most users are at least aware of threats concerning the Internet such as viruses, worms, and trojan horses. In addition, recent security analyses of web-portals of Mobile Network Operators (MNOs) have shown that even those portals can be easily misused in order to spy out the user's preferences and personal configurations, to order mobile services and additional mobile phones for free, and even to enter mobile devices with malicious code. Nevertheless, especially mobile devices are still be considered as save and secure, although, there is neither integrated reliable encryption of application data or communication channels nor trustworthy service authentication nor additional user authorisation.

In contrast, in the proposed framework value-added services are only performed at the homebase. The homebase receives the necessary information by particular Basic Web Services based on the users' preferences. Those specialised services provide a combination of useful low-level information such as a city map, emergency services, points of interest, and for example the schedule of the public transport system.

In general, this approach has several advantages. On the one hand, the user's privacy is warranted by consolidating and analysing basic information at the homebase where value-added services such as the planning of a business trip can be provided. At that place under the user's control, finally, the current location based on a passive positioning system such as GPS will be added. Therefore, the exact position of the user never leaves the user's area of trust comprising the device and the homebase. On the other hand, small- and medium-sized specialised and independent service providers – for example hotels and restaurants – are no longer restricted to only one single-sign-on-portal (often considered as single-point-of-failure) where they offer their services together with hundreds of others. They simply describe and provide their services and

information in a standardised form based on Web Service technology and can be finally found at so-called UDDI repositories.

3. Conclusion & Outlook

At that point, the question of the business case can only partially be answered. The business model for Basic Service Providers is as simple as efficient. Only standardised information have to be offered and could be charged depending on their preparation cost. Much more difficult is the business case regarding ISPs and MNOs, although they obviously scent the big market by managing identities and offering value-added services as described above. However, chained identity information, customer loyalty, and targeted marketing are the other side of the story.

Currently, Open Source seems to have a head start considering a feasible and reasonable model to enhance the different components of the user agent's homebase. Although, applications developed under the paradigm of Open Source, indeed, are not free of bugs. At least, the fact that everybody can participate in the development, testing, and debugging process is a promising approach to provide a service platform in respect of the user's security requirements and privacy constraints.

Both simplifying service provisioning and establishing privacy protection are only two advantages identified. Others still have to be elaborated and realised. For example, digital rights management (DRM) offers adequate mechanisms not only for protecting digital content such as video and audio files but could also be applied to identity management; users could associate their partial identities with specific purposes and an expiration date. Furthermore, mobile devices might be enhanced with the Trusted Computing Platform (TCP) in order to protect the access to Smartcards or biometry sensors. However, this is subject of forthcoming research work.

References

[1] Johann Bizer, Dirk Fox, and Helmut Reimer, editors. *DuD - Datenschutz und Datensicherheit. Schwerpunkt: Identitätsmanagement.* Vieweg Verlag 09/2003.

[2] Mario Hoffmann, Jan Peters, and Ulrich Pinsdorf. Multilateral Security in Mobile Applications and Location Based Services. In *ISSE - Information Security Solutions Europe,* Paris, France, October 2002.

[3] Birgit Pfitzmann. Privacy in enterprise identity federation - policies for Liberty 2 single signon. *Elsevier Information Security Technical Report (ISTR),* 9(1):45–58, 2004.

[4] Birgit Pfitzmann and Michael Waidner. Federated Identity-Management Protocols -Where User Authentication Protocols May Go -. In *11th Cambridge International Workshop on Security Protocols.* Springer-Verlag, 2003.

[5] Kai Rannenberg. Multilateral Security - A Concept and Examples for Balanced Security. In *Proceedings of the 2000 Workshop on New Security Paradigms,* pages 151–162, Ballycotton, Ireland, 2000. ACM Press.

[6] Jeroen van Bemmel, Mario Hoffmann, and Harold Teunissen. Privacy and 4G Services: Who do you trust? 10th Meeting - Wireless World Research Forum, New York, NY, USA, Oct 27-28, 2003.

[7] Stephen A. Weis, Sanjay Sarma, Ronald L. Rivest, and Daniel W. Engels. Security and Privacy Aspects of Low-Cost Radio Frequency Identification Systems. In *Proc. of First International Conference on Security in Pervasive Computing (SPC 2003),* volume 2802 of *LNCS,* Boppard, Germany, March 2003. Springer-Verlag.

PRE-AUTHENTICATION USING INFRARED

Amir Spahić [1], Michael Kreutzer[2], Martin Kähmer[2], Sumith Chandratilleke[2]

[1] *Faculty of Organization and Informatics Varaždin*
University of Zagreb
Pavlinska 2, 42000 Varaždin, Croatia
aspahic@foi.hr

[2] *Institute of Computer Science and Social Studies*
Dept. of Telematics
University of Freiburg
Friedrichstraße 50, D-79098 Freiburg, Germany
{kreutzer, kaehmer, sumith}@iig.uni-freiburg.de

Abstract Using complex authentication and verification methods is not always feasible in application fields with time and resource restrictions. However, fast and configuration-less authentication methods are required in many pervasive computing applications using wireless connectivity. In this paper we present an authentication mechanism which uses context information for its first phase, the so called pre-authentication phase. During this phase a connection between two devices is established to generate a common secret as a prerequisite for the subsequent authentication. We present an implementation of a special device called "magic wand", using optical communication for the pre-authentication phase. With the help of this device it is also possible to quickly authenticate devices for subsequent use in service discovery.

Keywords: Ad-hoc authentication, pre-authentication, infrared, transient security association, wireless connectivity.

1. Introduction

Using complex authentication and verification methods is not always feasible in application fields with time and resource restrictions. However, fast and configuration-less authentication methods are required in many pervasive computing applications using wireless connectivity. For

instance when entering a room that is enhanced with pervasive computing technology (cf. [8]) the user might need to securely associate his personal device to the (wireless accessible) computing environment, despite there is no pre-existing mutual knowledge of the device and the "roomware". Not in all cases physical contact can be used to run a secure authentication, e.g. in the case where sensors are built into the ceiling.

In this paper we present an authentication mechanism which uses context information for its first phase, the so called pre-authentication phase. During this phase a connection between two devices is established to generate a common secret as a prerequisite for the subsequent authentication. During this time it must disable intentional or unintentional involvement of a third party. Besides solving the security problem the mechanism should be fast, cheap, simple and easy to use.

2. Attacker Model

The endpoints of the communication, i.e. both devices in consideration, are assumed to be trustworthy. Using exclusively wireless technology, the focus of the attacker model lies in the air interface. According to the application fields in infrastructure-less environments and the dynamics of an ad-hoc basis of usage we assume as intentional attack eavesdropping (originated by a man in the middle also capable to effectuate a subsequent replay attack) and as unintentional attack the identification of the false device (misdirection). Denial of service attacks are not regarded in this paper.

3. Related Work

Even if the wireless technology gets more and more important, two devices that are in the range of each other, should not in each case "talk" to each other: this imposes not only scalability problems but also security problems, especially related to authentication (cf. [9]). However, as in [3] suggested, an authentication mechanism is needed to explicitly "marry" two formerly mutually unknown devices, i.e. two devices which haven't any (even partial) knowledge about the existence of one other. Such an authentication mechanism has been proposed by [1] and has been called "ad-hoc authentication" by [2].

As the focus of [1] lies in asymmetric cryptography with PKIs, its mechanisms even protect against active attacks like impersonation during authentication establishment. However, it is questionable whether this attacker model is realistic for the majority of the application scenarios and whether there are lightweight mechanisms to deploy and main-

tain pre-shared secrets. Being secure against passive attacks is often sufficient in infrastructure-less environments. Furthermore asymmetric cryptography with PKIs can be only performed by computationally strong devices and pre-shared secrets. For the verification of the validity information of public keys online access to black lists is needed.

Mechanisms to establish shared secrets using context information have also been proposed by [7], [6], [4], and [5].

In this paper we present a concrete implementation and evaluation of the ideas and concepts presented in [2].

4. Four Phases of Ad-hoc Authentication

According to [2] the four phases of ad-hoc authentication are (in [1] these are almost the same, however it lacks the last phase):

1 Pre-authentication phase: Secure establishment of a shared secret or mutual knowledge of identifying data about the other device (for example a public key). This may be done not only by direct communication, but also with the help of, or even exclusively by using context (in [1] the latter is also called "demonstrative identification"). Context may not only be sensed but also can be explicitly created, cf. the acceleration events of smart-its friends [4].

2 Authentication phase: Verification of the identity using the shared secret.

3 Use of authentication phase: In most cases authentication is the basis for subsequent security mechanisms like access control, encryption, integrity, etc.

4 Releasing the security association phase: This means "forgetting" the data collected in 1 and 2, i.e. explicitly deleting any information relating to the (former) partner device. This is done to prevent replay attacks.

Phases 1 and 2 are separated, because they fulfill different tasks: phase 1 uses context information to securely select devices and to establish a "context"-key. Authentication in phase 2 typically is done using undirected radio technology like Bluetooth or WiFi. Phase 2 checks that the correct devices are interconnected using the commonly generated key. There are many protocols that fulfill this task, challenge-response for instance. The key of phase 1 can also be used in phase 2 to establish a (stronger) session key, but authentication is not restricted to this mechanism. If both devices would share a secret or knowledge about each other, their designated authentication mechanism should be used only

after the secret of phase 1 has been verified as this procedure protects against denial of service attacks that take place on the radio channel (this saves energy as well: dependent on the scenario, the radio link only needs to be activated after a successful run of phase 1).

5. Pre-Authentication Mechanism

5.1 Design decisions

Beside the desired security properties our guiding design criteria for the mechanism are: fast, cheap, simple, lightweight, and no pre-existing mutual knowledge of the devices.

As context is used for pre-authentication, a location-limited channel must be taken (cf. [1]). When using communication technologies, they should have physical limitations in their transmissions, for example the necessity of line of sight and limited range, like "the PDAs are directed to each other and have a distance less than 20 cm". The reasons for the need for a location-limited channel are twofold:

- This kind of channel guarantees authenticity. This means that it is impossible or at least difficult for an attacker to transmit in the location-limited channel. This property is sufficient to ensure that information exchanged over the location-limited channel will allow the parties involved to securely authenticate each other (even in the presence of potential attacker).

- This kind of channel prevents unintentional false identification of the partner device (misdirection).

We argue that all undirected radio technologies are inappropriate for a secure implementation of phase 1 (this also holds for Near Field Communication[1]) as these mechanisms do not guarantee the necessary location-limited channel. An attacker can use devices with high energy radio to fool the selection of devices and to become man-in-the-middle. To make a reliable statement on this question further research is necessary.

We chose IrDA (standard according to the Infrared Data Association based on optical communication via infrared) for the pre-authentication phase. The connection is limited to a one-meter distance and the beam widening is only 30°. Infrared beams have defined orientation and we can use them to transfer data (the so called "point and shoot" principle).

There are also some other possible solutions concerning phase 1 like physical contact (for instance key distribution device interface, also called "fill gun"), common acoustical experience, shaking two devices

(common acceleration experience, cf. [4]), etc. We suppose usability to be scenario-dependent. The question of usability and fitness for different application fields must be evaluated by making a usability study with implementations of all these approaches.

We establish a Diffie-Hellman secret in the pre-authentication phase. An eavesdropper cannot calculate the resulting common key of both parties even if the attacker is able to intercept all messages.

Phases 2, 3, and 4 can be done based on a radio link (we are planning to use Bluetooth).

We are using PDAs for our sample implementation and evaluation, as they are widespread general purpose devices.

5.2 Pre-authentication as a three-step mechanism

Now we can define pre-authentication as a three-step mechanism:

1 Step: Establish an infrared connection.

2 Step: Use the Diffie-Hellman algorithm to create a key.

3 Step: Terminate the infrared and establish the radio (Bluetooth) connection.

These three steps are the basis for the subsequent phases.

6. Implementation

The used PDAs have some limitations like small processor power and restricted energy resources. In the following we will call our implementation of pre-authentication phase IrEx, which is implemented as a client-server model. At least one of the partners must have a self-standing application running called the irexserver. The client part is also a self-standing application, called the irexclient. The device which starts the irexclient is the client which initiates the exchange.

First we developed IrEx primarily for the Pocket PC platform but now we have a solution for Palm devices too. This means we can use Pocket PC - Pocket PC, Palm - Palm and Pocket PC - Palm connections.

6.1 Basic idea

To create a Diffie-Hellman key we need the following parameters:

1 P and Q - large prime numbers, where (P ¡ Q) and ((Q-1) / 2) is also a prime number

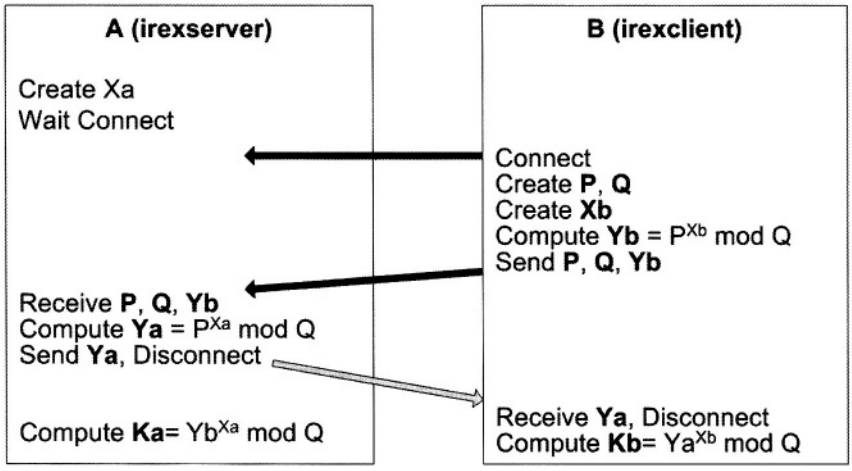

Figure 1. Protocol of the pre-authentication phase

2 Xa and Xb - secret random numbers (each side has its own X)

3 Ya and Yb - public numbers

P and Q are common for both sides so we give the right to the initiator to choose them. In our case these numbers are determined by the irexclient application. The random number generators on both sides create their appropriate X numbers. Those numbers are kept secret and they are not exchanged.

As it can be seen in figure 1, the server waits for an incoming request. When the client is started it tries to connect itself to the server. If the client finds the server, first it sends and then receives the necessary parameters for the creation of the common secret key K. After exchange of three messages on both sides, both of them can compute K and terminate the infrared connection.

6.2 Implementation details

To test the IrEx application, HP iPaq 2210 Pocket PCs were used. It is recommended to specify buttons in order to start the applications irexserver and irexclient. The irexserver starts with the right button and the irexclient with the left button. Pressing the right button starts the server program. It listens for incoming requests or shuts down if the client program is called.

When the irexserver is started on both sides, nothing will happen until one side initiates an exchange and thus takes the client role. When the

right button is pressed the irexclient (and with it the secret key creation) is started.

When an incoming request is noted, the server closes all other server ports until the exchange procedure is done. Pressing the left button (on the second device) starts the client program and the exchange procedure is started. This procedure consists of the following steps:

- The server opens a socket and waits for the client
- The client opens the socket and sends the request (sending the generated prime numbers and the public key in the same message)
- The server gets the prime numbers (P, Q), computes its public key and sends it to the client; at the same time it sends a message to close the socket_client.

If no message arrives the client knows after a waiting time of 5 seconds, that something went wrong, and it gives a double beep alert. In the other case, it gives a single beep and an LED signal. The server has no timeout.

The procedure will not start automatically; it demands explicit user action which increases the security.

7. Conclusion and Future Work

We have made a conceptual design of the pre-authentication phase of ad-hoc authentication. Furthermore, we implemented it using the infrared technology. The performance results of the first tests are promising; however, to gain reliable results we need further systematic testing. We plan to make a comparison between the times for connection establishment, for message exchange, and for crypto-calculations.

Our application does not have the full functionality of ad-hoc authentication. When using Bluetooth for phase 2 (authentication) and possibly 3 (use of authentication) we will take the key from the IrEx application as a basis to ensure the desired security property between the partners. The Bluetooth part will probably implement the OBEX protocol too. At this stage of the project we will make performance measurements and user tests.

Our research still leaves the question how to set some kind of "reincarnation policy" and authorization schemes.

Our prototype inspired us to a new user interface: the initiating PDA can be used as a "magic wand" to select one device from a group of devices. This is useful to support device discovery: in many cases it is more natural to select a device in the vicinity by pointing to it instead

of selecting it from a list of service alternatives on a small screen. It is even thinkable to implement such a "magic wand" on a smaller device, to enhance it with an RFID reader, thus building new security bridges between the real world and the virtual world.

8. Acknowledgments

We thank Kerry McGawley for her helpful comments on readability and Prof. Dr. Günter Müller and Prof. Dr. Blaženka Divjak for encouraging support. This research has been supported by the Kolleg "Living in a smart environment" of the Gottlieb Daimler- and Karl Benz-Stiftung and by the TEMPUS project of the European Commission.

Notes

1. http://www.nfc-forum.org/

References

[1] D. Balfanz, D. Smetters, P. Stewart, and H. Wong. Talking to strangers: Authentication in adhoc wireless networks. In *Symposium on Network and Distributed Systems Security (NDSS '02), San Diego, California,* 2002.

[2] S. Chandratilleke and M. Kreutzer. Credential-basierte ad-hoc-authentifikation (engl.: Credential-based ad-hoc-authentication). In *netzwoche Netzguide E-Security, Netzmedien AG, Basel,* 2003.

[3] L. M. Feeney, B. Ahlgren, and A. Westerlund. Spontaneous networking: an application-oriented approach to ad hoc networking. In *IEEE Communications Magazine,* 2001.

[4] L. E. Holmquist, F. Mattern, B. Schiele, P. Alahuhta, M. Beigl, and H. W. Gellersen. Smart-its friends: A technique for users to easily establish connections between smart artefacts. In *Proc. of UBICOMP 2001, Atlanta, GA, USA,* 2001.

[5] T. Kindberg and K. Zhang. Secure spontaneous device association. In *Proc. of UbiComp 2003, Seattle, Washington,* 2003.

[6] J. Light. Security, privacy and trust issues raised by the personal server concept. In *this book.,* 2004.

[7] J. Rekimoto, T. Miyaki, and M. Kohno. Proxnet: Secure dynamicwireless connection by proximity sensing. In *Proc. of Pervasive 2004, Linz/Vienna,* 2004.

[8] P. Robinson. Architecture and protocol for authorized transient control. In *this book.,* 2004.

[9] F. Stajano and R. J. Anderson. The resurrecting duckling: Security issues for adhoc wireless networks. In *Lecture Notes in Computer Science, Vol. 1796, Springer,* pages 172–194, 2000.

ARCHITECTURE AND PROTOCOL FOR AUTHORIZED TRANSIENT CONTROL

Philip Robinson
Teco, University of Karlsruhe & SAP Corporate Research. Vincenz-Prießnitz-Str. 1, 76131 Karlsruhe, Germany

Abstract Having considered adaptation and usability as two major challenges for security in pervasive computing, this paper presents an architecture and protocol to address the associated challenges. The major challenges specified are resource modeling, context-awareness, adaptive control and dynamic interaction. With an operation-centric approach, a dual-controller architecture is presented and a coordination protocol specified.

Keywords: Access Control, Adaptation, Authorization, Controller, Control-Situation, Perception, Policy, Resource, Transience

1. Introduction

A system with static interrelationships and purely atomic interactions is simple to manage. However, this is not a practical assumption of real world systems, where items in an environment have multiple relationships with users, including shared ownership, and hence multiple operating modes. Corner & Noble make a similar observation in their work on transient authentication, where they discuss the fallacy of infrequent and persistent authentication between people and their devices [8]. Stajano also addresses this theme by defining techniques and mechanisms for asserting and ending the transient association between people and devices in his Resurrecting Duckling protocol [15]. Therefore, it is often argued that a major requirement for security in pervasive computing is dynamic adaptation as opposed to rigid prescription of system controls [4]. This however requires a richer model of interaction between the security management system and the real world environment of the resources it monitors and controls. Advances in context awareness and senor networking facilitate this form of interaction even if the management system and resources are distributed. The protection goals of a system seem to however act against the goals of awareness, usability and

ubiquity of information and resources, which are the foundation of pervasive computing [19]. Consider a scenario where a homeowner becomes ill, is at home alone in bed but has all entry points such as doors, and windows, as well as access to the home's utilities, locked away and under her sole control. This may serve to work against the homeowner if a friend or medical personnel were to arrive on the scene and try to come to their aid. The home's security system would serve the homeowner better, if it were capable of changing its controls only for the situation that an emergency is detected. Covington used such a scenario as motivation for his work on "Environment Roles" and "Parameterized Authentication" [2], where the signals sensed from the environment were incorporated into both authorization and authentication decisions. Zhang, with the DRBAC (Dynamic Role Based Access Control) model, also notes the necessity to change a user's access privileges when the user's context changes as well as to change the access permissions of a resource when its system information changes [20]. There are also other research contributions that focus on the design of context-based security systems [7] and including context parameters in authorization constraints [6, 11]. From a review of existing and ongoing work, an effective management system in such environments should address the following issues

1 Resource modeling: comprehensive "world/domain" model of resources and their interactions

2 Context awareness: interpretation of environmental and system situations

3 Adaptive control: enforcement of situation-based policies

4 Dynamic interaction: coordinated, dynamic update of user-resource interfaces and modes

This paper addresses the above issues by combining them under the heading "authorized transient control". The intent was to emphasize operational matters as opposed to design and implementation issues that appear to be already well addressed by existing work. The analysis of the problem therefore commenced with a consideration of operational roles as opposed to system components, as is the approach in the area of control systems [17, 13, 9, 10].

Authorized transient control suggests that a user of a system is granted provisional access to a resource, given that certain conditions currently hold. The user is therefore referred to as an "authorized transient

controller (ATC)" in that through interaction with resources in the environment, it is desired that a particular system state be reached, which satisfies the user's goals or asserts task completion. In addition, each environment also has one or more administrators, whose responsibility is to ensure that the system's authorization policies are maintained. This role is referred to as a "fulltime controller (FTC)" in this paper, A control problem arises where two controllers need to balance the controllability of the same target resource, each having different control goals. A discussion of these terms and issues is included in section 2, the problem analysis and approach. Section 3 presents the architecture of the management system described, while section 4 gives an explanation of how the architecture components interact in the form of a message-based, state-modeled protocol, for the coordination state between the two controllers. The conclusion summarizes the contributions of the paper and ongoing work.

2. Analysis and Approach

The term "authorized transient control" is meant to describe an aspect of how entities in a pervasive environment interact with each other and with resources. The difference from traditional computing is the degree of spontaneity and dynamics of interaction afforded in pervasive computing [4, 5]. This results in entities and resources being transiently related. However, there are still security and usability goals to be considered when building applications in such environments, withstanding the great flexibility promised by pervasive computing. The security goal considered in this paper is that of authorization, while effective, reliable coordination of resource access and interaction control underpins a system meeting its usability goals [19]. An *authorized* subject is entitled to access and use a target resource for performing a set of operations provided a set of constraints hold. Authorization can therefore be represented by the template $\langle A, B, P, C \rangle$, where A is a subject or subject-role, B is a target resource, P is a set of permissions (operations to which A is entitled) and C is a set of constraints or conditions that apply to the permissions granted [14]. Secondly, the term *transient* applies to what is actually short in its duration or stay, as opposed to having preconceived intentions and natural tendencies to be long-term or permanent [15, 8]. If an authorization is considered transient, this implies that its constraints are modified by a situation S, where S could be a time-range or other sensed properties of the resource and its environment. Nevertheless, note that the term "transient" has an established meaning in the field of reconfigurable control systems, referring to a phenomenon that arises when

a system switches from one operational mode to a next [13]. This second notion of transient is not discussed in this paper, but is marked as an issue that should be addressed as adaptive security does entail switching operating modes of a target system. Thirdly, if a subject A *controls* a target resource B, A monitors a set of control properties Rn that refer to B and its operational environment, compares them to a set of control reference properties R_0, and generates an action O that counteracts the comparative error between R_n and R_0. This definition of control is derived from Powers' work on "Perceptual Control Theory" [10], which forms a part of the approach discussed later in this section. Other useful descriptions of the term "control" come from Petersen, who states that the role of a human operator [controller] is to bring about desired state changes (or non-changes) in a controlled system [9]. Therefore if A is an authorized transient controller of B, A is permitted to perform an action resultant from comparing the properties R_B and R_A, in order to control the operational state of B to bring about R_{AB} for the validity of a situation S.

Considering the above definitions, resources in pervasive environments can be said to have multiple controllers with different references or operational goals. However, only two types of controllers are considered for the purposes of this paper - the "Interaction Controller" and the "Access Controller". The Access Controller carries out control operations on behalf of a fulltime controller or administrator, while the Interaction Controller acts on behalf of a transient controller or user of a resource. The two controllers therefore have different perceptions of the target resource, its situation and that of its environmental signals. Figure 1 depicts how these two different controllers are seen to operate on the same target resource.

Figure 1 has introduced new terms that may lack intuitive meaning for readers unfamiliar with perceptual control theory (PCT) [10]. PCT is based on the premise that dynamic systems do not plan and process repeatable actions; rather they plan and process perceptions (or desired views of a system), and hence producing repeatable results with varied conditions. The principles of PCT adopted in the controller model in figure 1 are defined below:

[Access/Interaction] Perception: this is the relevant view that a human controller has of a resource dependent on its operational state. The operator need not know every detail of the resource's operational state but sufficient detail for the support of effective control decision-making. The human controller may receive this directly from a resource, but in the model used here, there is an intermediate controller module or agent

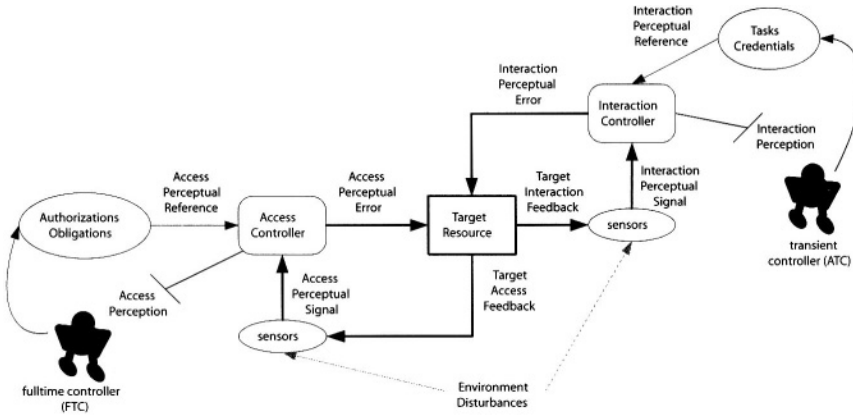

Figure 1. Depicts multi-controller interaction with a target resource by an interaction controller and an access controller

that automatically adjusts the perception in order that in the best cases the human operator constantly receives an "ideal view" of the resource.

[Access/Interaction] Perceptual Reference: this represents the "ideal view" that the controller wishes to receive from the resource. In the case of the fulltime controller (FTC) and Access Controller (AC), the source of the perceptual reference is authorization and obligation policies. These policies are specified by the FTC and enforced by the AC. In the case of the transient controller (ATC) and interaction controller (IC), the source of the perceptual reference is the tasks the ATC wishes to carry out as well as the credentials that certify some set of rights.

[Access/Interaction] Perceptual Signal: this is the input that a controller receives from a sensor system, which represents the control state of the target resource, with respect to its observable properties, as well as that of its environment.

[Access/Interaction] Perceptual Error: this is the calculated comparative error between a perceptual signal and a perceptual reference. That is, this is the controller's calculation of how much the actual perception of the target resource deviates from the ideal perception as defined by the perceptual reference.

Environment Disturbance & Feedback: these are both property sets sensed by a sensor system. "Feedback" is the actualized value of explicitly monitored properties of the target resource, while the "environment disturbance" is monitored properties of the environment. The environment disturbance may have either an indirect or direct effect on the target resource's control state and hence perceptual signals.

From the above model it is observed that feedback from the target resource simultaneously results in two classes of perceptual signals, and that the resource may also simultaneously receive two forms of perceptual errors and control commands. Breemen and Vries discuss and reference a number of problems that arise in systems with multiple controllers [17], which also apply to the interpretation of multi-controller used here. Three of these multi-controller problems addressed by the architecture and protocol are *conflicts, deadlocks* and *coordination* of switching between controllers. Conflicts may arise as a result of contrary perceptual references or if the controllers attempt to simultaneously enforce a control on the target resource. In the context of the access and interaction controllers, a conflict arises if the authorizations and obligations specified at the AC do not support the tasks and credentials of the IC or if the AC tries to perform an access control at the same time the IC performs an interaction control (and vice versa). Deadlocks refer to exceptional control situations which none of the controllers are prepared to handle. There could therefore be a case where an irresolvable perceptual error occurs at both the interaction and access controllers - e.g. hardware or software failure - which may render the target resource as unavailable. The coordination of switching between controllers means that rules have to be defined for when and how control is to be exchanged. Although the AC typically has a higher controller priority than the IC, there may be situations, such as the emergency response scenario, where this priority should be overridden to allow the IC to work more efficiently. This means that the AC will in this case need to adapt its perceptual reference to accept the new controllability of the target resource. The architecture provides more details on the design of a management system to computationally support the controller model, giving consideration to the issues discussed.

3. Theoretical Models and Architecture

The architecture was purposefully designed as four modules, in order to address the four challenges identified for the management system. There are also interfaces between these models, which are discussed further in section 4 - the protocol - but depicted as relationships between classes in the architecture diagram. The challenges of adaptation, context-awareness and resource modeling are not unique to pervasive computing, as there exist well-researched theoretical models for each of these problems. The task was therefore to filter what was relevant to the problem domain discussed in the paper, as well as what showed promise of being computationally implemented. Figure 2 presents the high level

UML diagram of the four conceptual models for the management system and their integration.

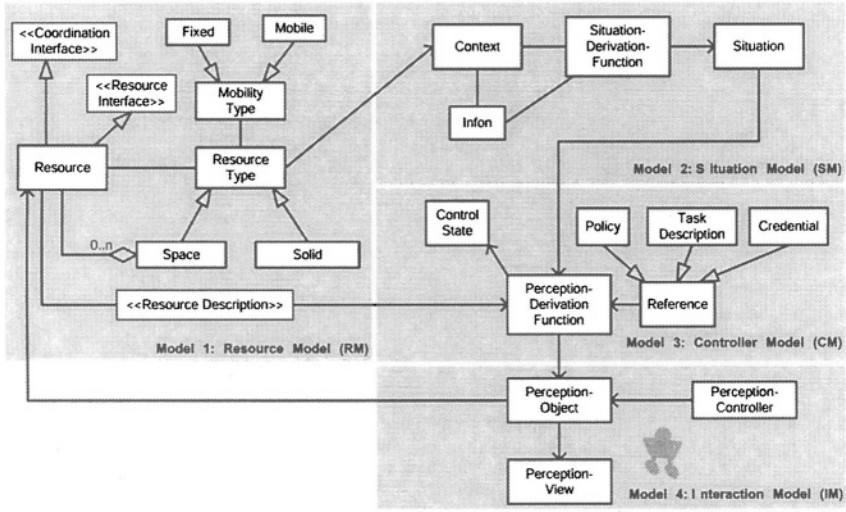

Figure 2. UML Diagram for Authorized Transient Control Management System, showing four component models and their interfaces

Resource Model. The goal of this model is to facilitate dynamic registration of resources in the control environment. Registration implies that a resource receives some identification attributes from the controllers, and each instance of a registered resource in the system providing a "Resource Description", a "Coordination Interface" and a "Resource Interface". The Resource Description is resource metadata, including context, type and mobility, while the Resource Interface describes its methods and parameters offered to authorized users. The Coordination Interface is the interface that the controllers use in order to check the status of the resource and exchange coordination messages pertaining to the system state. This model is therefore the foundation for deriving control situations based on state changes. Petersen refers to this sort of model as a "work domain model", which supplies a management system with different levels of relevant behavior [9]. Petersen refers to a control situation as comprising of actualities, possibilities and norms, which when changing may influence the control situation. This model therefore supports the specification of "norms" - prescription of appropriate state changes in the controlled system [9], while the actual-

ities and possibilities are modeled in the situation and controller models respectively. The class design principles were nevertheless inspired work from Scott et al., where they describe a spatial policy framework for mobile agents based on the ambient calculus [12]. One of the goals was to model the "world" inclusive of both physical entities and virtual agent entities. They began by defining a set of entity sorts/types and then define a calculus for defining their behavior, relationships and interaction rules. The top-level entities in their model are however immediately specialized (e.g. workstation and laptop are two different top-level entities), as opposed to defining object-oriented inheritance relationships between the sorts - this could have served to enhance the semantics used in the calculus. The approach in this paper is therefore to define more high-level resource types; each resource is either of type "Space" or "Solid", where a Space may contain 0 to n resources of either resource type. Secondly, the mobility of a resource also has an influence on its behavior and controllability. Each resource type therefore has a "Mobility" type of either "Mobile" or "Fixed". A cargo container could be classified as either a "mobile space" or a "mobile solid" dependent on the perspective and allowed detail of a controller. That is, a controller who is not allowed to view the contents of the cargo container would perceive it as a solid. Therefore, changes in the structure, location and policy of resources are represented by reassignment of resource types and transfer to different spaces. The norms of the environment are the semantic relationships of the resource types. For example, "a 'fixed' resource cannot be moved" is a norm processed at the lowest level of behavior monitoring. The model also applies to information resources, where e.g. a "namespace" would be considered a "Fixed Space", a folder in the namespace a "Moveable Space" and the electronic documents in folders would be "Moveable Solids". However, a folder may be represented as a solid to a user if he has no rights to read its contents and can only be aware of its existence.

Situation Model. The Situation Model is used to detect and interpret states of resources and effectively the overall environment. The meaning of the term Context tends to vary dependent on the system or domain of usage. It has linguistic, philosophical and computational designations, but at least there is a consistent notion of it being the foundation for attributing meaning to information [3–1, 16, 18]. The Situation Model has its foundations in the work of Surav and Akman, where they have pursued the creation of a computational model for context with situations, based on Barwise's Situation Theory [1]. The Situation Theory has the fundamental components of *infons, situations*

and *constraints,* where each is considered to be a "first-class object". An infon is an atomic, most basic unit of information, such that it can refer to pure facts as opposed to information that requires further inference by the receiver. This would therefore refer to simple sensor readings. A situation is consequently a purposeful collection of infons that form a higher-level statement about a subject, while constraints are the dependencies between situations. A Context relates "infon types" with resources in the environment, and provides a "Situation Derivation Function" that gathers actual "infons" or "actualities", in order to produce a first-class Situation object. The difference between a Context and a Sitaution is compared to that of the difference between a Class and an Object in OO Theory. A Situation is a committed statement or instance of truth or falsehood of a particular Context, including the time of occurrence, source, certainty, and calculated rate of change. It is by listening for certain situations that a controller can determine when "Control Situations" arise.

Controller Model. The Controller Model is where the major decision and coordination logic of the system is performed. This is therefore the mechanism for processing "possibilities", based on derived control situations produced by the Situation Model. The central function or "primary controller logic" is referred to as the "Perception Derivation Function". This function takes resource descriptions, control situations and references as its inputs, and produces an object known as a Perception Object (specified as part of the Interaction Model). The resource description provides the state of the properties of the device including its mobility and spatial type, while the reference is a higher-level system policy, task description or set of authorized user credentials. Authorization policies state what operations a transient controller is allowed to perform on a resource, while an obligation specifies what the controller must do as a control duty in the environment [14]. A task description may be represented as single, atomic task items, or as a stateful workflow. The term "credential" is used to represent identifying properties of a transient controller including alias-password pairs, public keys and certificates, and sensor-obtained properties. The other important function of the controller model is the maintenance of the overall system state. Recall that as this is proposed as a multi controller model this maintenance entails some deal of coordination between the controllers. The Perception Object is therefore functionally related to the Control State of the overall system. The control state is discussed further in the protocol section, as the controller model plays a significant role in executing the protocol.

Interaction Model. The Interaction Model is the simplest of the four models. Its main purpose is to broker the Perception Objects between the Controller Model and a human operator, but it also provides the possibility for "secondary controller logic" to be defined in the "Perception Controller". Primary controller logic refers to making control decisions based on situations, while secondary controller logic is automating control decisions based on a perception. If there is no automation/secondary control logic implemented, all environmental disturbances that shift the error in the system go to a human operator through the Perception View. If the secondary controller logic is sufficient to respond to the control situation, the human controller receives a stable perception view. Readers familiar with the MVC (Model View Controller) software engineering pattern will recognize the design benefits of separating the perception controller from the object in this way.

The specification and interpretation of each model is and can be more detailed, but for the purposes of this paper it is important for the reader to be able to see the relationship between the controllers, the architecture and the protocol rather than the detailed specification of each model.

4. Protocol

The protocol description is in two parts; firstly, it describes the ordering and processing of messages passed between the architecture models. The models are realized as four controller components but the interfaces are actually between subcomponents of the models, as represented in the architecture diagram (Figure 2). Secondly, the protocol is defined as a state model for authorized transient control, which corresponds to controller states as different control situations and events occur. Each controller is composed of two threads - one for "listening" and one for "controlling". The listening thread is interfaced with the resource and situation component models, such that it is responsible for coordinating the reception of situation data and coordinating the output of control commands to the resource. It is hence the interface between the resource and the controller. The "control thread" is interfaced with the controller and interaction models, and is therefore responsible for coordinating forwarding of control situations to the controller logic and receiving subsequent controller commands.

The labels a.0 - a. 10 are particular message types. These are listed below and their flow described in the subsequent paragraph:

0: Initialization / **1**: Infon / **2**: Situation / **3**: Awareness / **4**: Control Situation / **5**: Perceptual Signal / **6**: Perceptual Error / **7**: Control

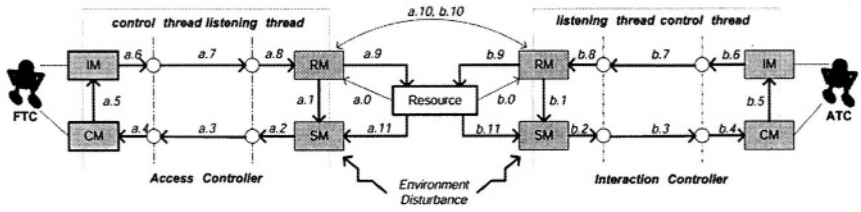

Figure 3. The operational model for controllers using the four architecture component models. The figure shows messages being exchanged between the component interfaces. The diagram is repeated for both the access and interaction controllers, but the message types are essentially the same. Messages "a.n" refer to those of the Access Controller, while "b.n" are the Interaction Controller's messages

Command / **8**: Validated Control Command / **9**: Resource Control Command / **10**: Coordination / **11**: Feedback

The first message (a.0, b.0) passed between a resource and a controller is referred to as the "Initialization Message". This provides the Resource Model (RM) with its Resource Description, which includes the Resource Type and hence set of Context instances that it supports - these are collectively referred to as the "Resource Context". The Resource Description is also forwarded to the Controller Module (CM) as part of the initialization process, but this is not shown in the diagram. If a new resource is added to the environment and registered with the RM, then this event is encoded as an "Infon Message" (a.1, b.1) and sent to the Situation Model (SM). The SM must also account for "Environment Disturbance" when deriving Situation Messages, such that the signals from the environment are converted to Infon Messages. After deriving the "situation" using the Situation Derivation Function, SM creates a "Situation Message" (a.2, b.2) and interrupts the listening thread. The listening thread essentially performs filtering on the messages, based on the "certainty" of the situation message. The listening thread typically discards situation messages with a certainty below a defined threshold. Accepted messages are then marked as "Awareness Messages" (a.3, b.3) and forwarded to the control thread. The control thread also performs a "relevance filter" on the messages in order to determine if it can be classified as a "Control Situation" or if it should be discarded. This is determined by the "freshness" of the message or other parameters. "Control Situation Messages" (a.4, b.4) are those that are accepted as relevant by the control thread and forwarded to the Controller Module (CM). The CM then derives a perception of the resource based on a function taking the control situation and a reference as input. The output is

a "Perceptual Signal Message" (a.5, b.5) to the Interaction Model (IM). Based on an evaluation of the controller logic or direct input from a human operator, the IM passes a "Perceptual Error Message" (a.6, b.6) to the control thread, which then makes a control decision and creates a "Control Command Message" (a.7, b.7) and forwards to the listening thread. The listening thread always plays validating role and hence validates the authenticity (may need to certify some proof that the message is from the correct controller) and relevance (consider that the situation may no longer be valid) for the command message. The listening thread then forwards a "Validated Control Command" (a.8, b.8) to the RM, which can then parse the command into the correct "Resource Control Command" (a.9, b.9) message format, as well as send a "Coordination Message" (a.10, b.10) to other controllers registered with the particular Resource Context. Coordination Messages are used to ensure the cooperative behavior of multiple controllers on the resource, in that the overall management system has a consistent state (see figure 4 below). This is known as a local or reactive coordination mechanism in multi agent systems [17]. Finally, the resource reflects issues a response to the command in the form of a "Feedback Infon Message" (a.11, b.11), closing the control loop. The inputs to the Situation Derivation Function at this stage would therefore include - controller coordination, resource state, and environment disturbance infons.

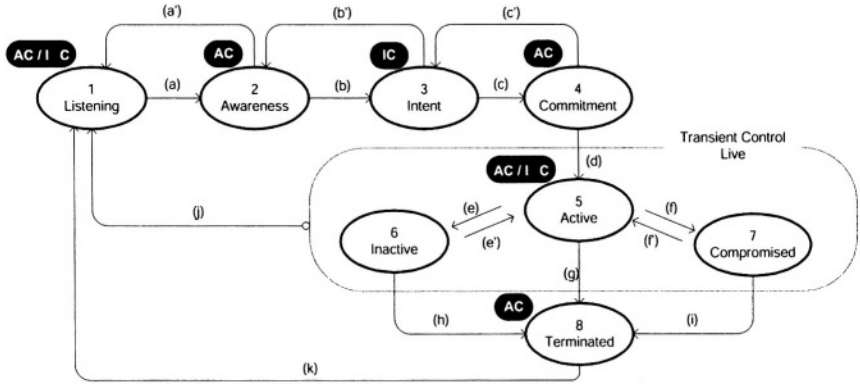

Figure 4. The local state model of the overall control system as coordination and interaction messages are passed between the controllers and their threads respectively. Each state is also labeled with its designated coordinator - AC: Access Controller; IC: Interaction Controller

The listening and control threads of the controllers constantly go through the states of "listen", "receive", "process" and "forward". How-

ever, there is a more complex state model (see figure 4) for the overall management system, which forms the basis for coordinating the controllers. In the model presented, the current assumption is that the Access and Interaction controllers are in the same control domain and therefore their behavior can be mutually trusted. In this case trust relates to their cooperation based on local coordination messages. Each controller is therefore said to be "autonomously reactive", matching the goals of "adaptive security". As the overall system state changes, each controller makes a decision based on its perceptual reference. While the access and interaction control activities are "closed-loop", the coordination activities are treated as "opened-loop", in that there is no imposition that either controller should provide explicit feedback. Nevertheless, if it is recognized that the state models are out-of-sync this is considered a "vulnerability window" or "imbalance" in the system [11], which the controllers must reconcile. Coordination is therefore the task of asserting that a particular state either exists or is transitioning. The controller model (CM) is responsible for maintaining the state machine of each controller; however, one of the controllers plays the role of coordinator per state. The states are described below by defining the controller that coordinates the relevant state, as well as the events that trigger reverse transitions and forward transitions respectively. The applicability of the model to a traditional PC workstation environment is also used as a practical example.

- **Listening:** the "Listening" state is also referred to as a "passive" phase of the system. It implies that there is no control situation detected by either of the CM's. Both the access controller (AC) and the interaction controller (IC) coordinate this phase until the AC is "aware" of a control situation and transitions to the "awareness" state via (a). Consider when a PC is in sleep mode and the monitor runs on low power.

- **Awareness:** the AC coordinates the "Awareness" state. Although the IC may also receive infons from the environment or the resource, it awaits the AC's assertion that a control situation really exists. Awareness could be based on the registration of a new resource, the change in state of a resource, or a disturbance signal from the resource's environment. Note that if a control situation arises and the IC does not respond, given a specified time interval or validity condition, the AC resumes the listening state via (a') otherwise it goes to "Intent" via (b). Consider a user pressing the "Ctrl-alt-del" key on a PC in order to awake it.

- **Intent:** the IC coordinates the "Intent" state. While the "Awareness" state is based on the input of implicit infon signals, the "Intent" state requires some explicit interaction by an operator. This means that unless an operator explicitly issues a command through the interface generated by the IC, the system state resumes "Awareness" via (b'). The IC is responsible for collecting credentials from the user, in order that the system state can be transitioned to "Commitment" via (c). Consider a user entering their username-password when a login screen is presented.

- **Commitment:** the AC coordinates the "Commitment" state, as this is when the authorization decision occurs. The IC therefore forwards credentials of a current user to the AC, which then gathers relevant attributes based on the control situation. If the user can be authorized, the AC asserts that the system state may be moved to "Active" via (d). If the credentials cannot be verified or the attributes of the control situation are or become invalid, the system state resumes "Intent" via (c'). Consider a user entering the incorrect username or password and being re-issued with the login screen.

- **Active:** both controllers AC and IC constantly exchange coordination messages during the "Active" state. The controllers therefore constantly assert that their perceptual references do not conflict with the perceptual signals of the current control situation. That is, the AC does not detect any activity that conflicts with its authorization and obligation policies, while the IC does not detect hindrance or disqualification with regards to user tasks and credentials respectively. However, the "Listening" state may be resumed via (j) if both the IC and AC agree that the extent of inactivity (infon generation) falls below a particular threshold, yet the control situation still holds. Consider an active login session with a PC, and possible timeouts during inactivity.

  - **Inactive:** the system state is moved to "Inactive" via (e) if for some reason there are no coordination messages exchanged between the controllers. Either of the controllers may attempt to re-engage the "Active" state via (e'), and await a coordination acknowledgement. When all attempts at re-engagement fail, the system resumes the "listening" state via (j), or, in the event that the control situation is expired, transitions to "terminated" via (h).

  - **Compromised:** the system state is moved to "Compromised" via (f) if either of the controllers calculates a discrepancy when comparing the control situation with their perceptual reference. If the controller logic cannot reach a point of agreement, and the control situation remains the same, the system state resumes "listening" via (j) to allow the human controllers to do offline negotiations. If the control situation expires, the system state is transitioned to "terminated" via (i)

- **Terminated:** the AC coordinates the "terminated" state, as it manages the control situation policies and transient constraints. Termination is based on the expiry of control situation, either by the AC sensing that the control situation is passed or that the tasks registered at the IC are complete. After the system is cleaned up, it resumes "listening" via (k). Consider a "log-off" event from a PC.

In order to clearly draw the connection between the state model and the earlier discussions about perception - each phase indicates a new perceptual view or class of perceptions, such that human operators can perceive the state of the system based on their interests.

5. Conclusion

This paper has addressed two important themes in security for pervasive computing. Firstly, the issues of coordinating complexity with regards to adaptive security and access controls, and, secondly, the issue of conflicting usability and security goals. The approach was to design the system as a dual-controller system, where one controller was responsible for mediating interactions and the other for authorizations and access controls. The system architecture of each controller was designed to explicitly address the four challenges stated in the introduction of resource modeling, context awareness, adaptive control and dynamic interaction. The protocol considers that the dual controllers need to be coordinated, and also includes how the coordination between these two controllers is managed. The properties of authorization, transience and control are therefore captured by the interaction of controllers and their human operators.

Although quite some work related to transient and adaptive security has been identified throughout the text, it was found that the focus was either on design aspects, a particular requirement (policies, context awareness, resource modeling, user modeling etc), or described a specific

mechanism. Contributions to the operational aspects (message protocols and coordination procedures) of the security management systems were still missing. This paper is therefore intended to encourage further research in this direction as especially pervasive computing becomes more mature. In addition, the usability and user-resource interaction elements of security are often neglected. The controller model was especially designed with this in mind, offering a system that allows specification of usability rules independently of the authorization rules, yet a protocol for their integration. Although today's access control systems are sufficient for single workstation environments, the model presented in this paper proved to be backward compliant with the needs of such resources. However, it is evident that as resource become more and more distributed and autonomous that the traditional approaches do not scale.

The next stage for this work is to implement the system in order to do application, performance and user evaluations. Secondly, the robustness and extensibility of the system need to be also evaluated from the perspective of distributed multiple resources as well as multiple access or interaction controllers. The effects of other aspects of authorization management, such as delegation and revocation, need to be also considered in the implementation.

6. Acknowledgements

The ideas with regards to coordination came about during the problem definition phase of the ongoing TrustCoM project www.eu-trustcom.com. I also thank my colleagues at TecO - Uni. Karlsruhe and SAP Research for their useful comments and time to discuss the ideas of this paper, as well as the feedback from reviewers and general discussion during the workshop. Finally, thanks to Nicole Robinson for her assessment of "real world" relevance of the paper.

Appendix: Remarks

Note that the contents and structure of the paper changed significantly after discussion in the workshop as well as subsequent discussions with colleagues about the direction of this work. The original workshop paper presented the architecture and protocol in still a very requirements gathering manner. The major problem cited from reviews of the first draft was the explicit focus on a particular problem and contribution. After the workshop the idea that addressing operational matters as opposed to purely design and implementation of adaptive security came to mind.

References

[1] J. Barwise and J. Perry. *Situations and Attitudes.* MIT Press, Cambridge, MA, 1983.

[2] M. J. Covington, P. Fogla, Z. Zhan, and M. Ahamad. A Context-aware Security Architecture for Emerging Applications. In *Annual Computer Security Applications Conference (ACSAC),* 2002.

[3] Anind K. Dey. Understanding and Using Context. *Personal and Ubiquitous Computing Journal,* 5(1):4–7, 2001.

[4] D. Hutter, W. Stephan, and M. Ullmann. Security and Privacy in Pervasive Computing: State of the Art and Future Directions. In *First International Conference on Security in Pervasive Computing.* Springer, 2003.

[5] Tim Kindberg, Kan Zhang, and Narendar Shankar. Context authentication using constrained channels. In *WMCSA,* 2002.

[6] Patrick McDaniel. On context in authorization policy. In *8th ACM symposium on Access control models and technologies,* 2003.

[7] Ghita Kouadri Mostfaoui and Patrick Brzillon. A Generic Framework for Context-Based Distributed Authorizations, 2003.

[8] Brian Noble and Mark Corner. The case for transient authentication. In *10th ACM SIGOPS European Workshop,* 2002.

[9] Johannes Petersen. Modelling Control Situations for the Design of Context Sensitive Human-Machine Systems. In *21st conference on Human Decision Making and Control,* 2002.

[10] William T. Powers. A Brief Introduction to Perceptual Control Theory. Copyright 2003 William T. Powers.

[11] Philip Robinson and Michael Beigl. Trust Context Spaces: An Infrastructure for Pervasive Security. In *First International Conference on Security in Pervasive Computing.* Springer Verlag Press, 2003.

[12] David Scott, Alastair Beresford, and Alan Mycroft. Spatial Policies for Sentient Mobile Applications. In *The 4th International Workshop on Policies for Distributed Systems and Networks,* pages 147–157. IEEE Policy, 2003.

[13] Gyula Simon, Tams Kovcshzy, and Gbor Pceli. Transient Management in Reconfigurable Control Systems. Technical report, Budapest University of Technology and Economics, 2002.

[14] M. S. Sloman. Policy driven management for distributed systems. *Journal of Network and Systems Management,* 2(4):333–360, 1994.

[15] F. Stajano. The Resurrecting Duckling – what next? In *The 8th International Workshop on Security Protocols,* April 2000.

[16] Erkan Tin and Varol Akman. Computational Situation Theory. *SIGART Bulletin,* 5(4):4–17, 1994.

[17] Albert van Breemen and Theo de Vries. An Agent-Based Framework for Designing Multi-Controller Systems. In *Fifth International Conference on The Practical Applications of Intelligent Agents and Multi-Agent Technology,* 2000.

[18] X. H. Wang, D. Q. Zhang, T. Gu, and H. K. Pung. Ontology based context modeling and reasoning using OWL. In *Workshop on Context Modeling and Reasoning at 2nd IEEE International PerCom,* 2004.

[19] Ka-Ping Yee. User Interaction Design for Secure Systems. Technical Report CSD-02-1184, University of California Berkeley, May 2002.

[20] Guangsen Zhang and Manish Parashar. Context-aware Dynamic Access Control for Pervasive Applications. In *Communication Networks and Distributed Systems Modeling and Simulation Conference (CNDS),* 2004.

This page intentionally left blank

IV

SOCIAL AND TECHNICAL APPROACHES TO PRIVACY PROTECTION

This page intentionally left blank

OVERVIEW

While it might be difficult to predict when, why and to whom people are willing to surrender their privacy, it might be easier to explain what people are willing to pay for the *gain* of privacy. Sometimes, it is worth a lot to them. For example, when it comes to the question of why people are buying houses instead of just renting them, thereby giving up flexibility, one prevailing answer is privacy. There seems to be a fundamental human need for a place where one can be for oneself, without being bothered by anybody else. Generally, people like to retain control over the scope of their actions. Most social transactions are inherently limited in scope, since they affect only those who are in immediate proximity or, at least, the involved parties are known. On this basis, it is possible to make reasonable decisions regarding the disclosure of personal information. With the advent of ubiquitous and pervasive computing, social borders become increasingly blurred. The gap between the physical and the virtual world is narrowed by new technologies such as sensor networks, RFID tagging and smart environments. Human activity is being constantly monitored and fed into processes that are designed to help people interact, communicate, work or play. Nevertheless, acting anonymously or in a transient way is becoming more challenging. Digital traces are left behind in a variety of devices and background services, and by putting them together, like the pieces of a puzzle, a personal profile can be established - something most people would disapprove of.

This inherent danger of pervasive computing is potentially limiting the acceptance of novel services and products. People might prefer to abstain from using these services if they feel that their privacy is subject to compromise. On the other hand, most pervasive computing services are based on close interaction with the human user and can only be substantially useful if the user entrusts some of his or her personal likes and dislikes to them. This is the motivation for research in technologies that are compatible with the desire for privacy. So far, research results indicate that there are no pure technological solutions to this dilemma, but tight interactions between technical mechanisms and social norms are necessary. In this session, the problem was looked at from both sides, admittedly with an emphasis on the technological aspects. But

whenever a technological solution is proposed, its social limitations are stated and it should be clear that we have arrived at no final answers yet.

In their paper, "A Social Approach to Privacy in Location-Enhanced Computing", the authors ask for the reasons for which people are willing to give up their privacy. What do they expect to gain in exchange for giving away personal information that could potentially lead to an undesired privacy breach? The authors argue that this issue cannot be resolved by relying on rational, i.e. economic, reasoning only. They say that there are far too many factors involved in such decisions and the social life of people is far too complex to be described in terms of simple rules. Instead, privacy should be described as a concept of life that emerges from all the tiny bits that are accumulated during social interactions.

In the next paper, "Safeguarding Personal Data with DRM in Pervasive Computing", an architecture is presented for the protection of personal data that is based on the same technology being intended for the protection of copyrighted material. Such a design can help build devices that respect the privacy concerns of their users and enable services that deal with personal data in public environments. For evaluation purposes, a prototype was built based on the NetBSD operating system and the Veriexec extension, which provides similar functionality as the digital rights management chip TPM. A public terminal is used to display personal medical information. To access this data, it acquires license information from the personal device of an authorized person.

The third paper in this session, "Maintaining Privacy in RFID Enabled Environments - Proposal for a Disable Model", deals with the problem of digital traces stemming from the use of RFID tags on consumer products. Permanently disabling ("killing") RFID tags after purchase seems unrealistic, since potential benefits would be lost, for example easy access to product information or inventory management. Instead, the authors argue for a password mechanism that allows reenabling an RFID tag after it has been disabled. This would allow consumers protect their privacy but still benefit from RFID functionality if desired.

The papers were presented by Ian Smith, Adolf Hohl and Sarah Spiekermann. The discussion after Spiekermann's presentation centered around the password issue. It was regarded unrealistic that a consumer could keep track of a password for each and every single item he or she keeps at home. Therefore, such a scheme would be feasible only if there was either a single password for all items, or if password management could be automated. This could be done, for example, as the user takes

items into and out of his home: the tags are automatically enabled when he brings them home and disabled when he takes them out. Thus, he can benefit from the RFID functionality in the secure environment his home offers without being subject to surveillance in the public.

Smith's presentation was followed by questions regarding profiles that could help define what data should be disclosed in which situations. Profiles are made up of rules that provide the basis for such decisions based on situational input. The presenter however rejected this idea, stating that rules may be suitable for average case scenarios and situations, but not for extreme cases, which are much more important, after all.

Hohl's presentation triggered concerns about the actual implementation of a DRM-based public terminal. It seems that there may be issues with standard soft- and hardware that could make hardening such a public terminal quite challenging. For example, swapping of main memory contents to the hard disk could leave traces of private information that could be accessed later by an unauthorized person. Another problem is controlling access to personal information. The hardware token that issues the license to the public terminal is a security risk itself.

This page intentionally left blank

MAINTAINING PRIVACY IN RFID ENABLED ENVIRONMENTS

Proposal for a disable-model

Sarah Spiekermann[1], Oliver Berthold[2]
Humboldt University Berlin, Unter den Linden 6, 10099 Berlin, Germany

[1] *Institute of Information Systems*
sspiek@wiwi.hu-berlin.de

[2] *Department of Computer Science*
berthold@informatik.hu-berlin.de

Abstract The presence of RFID technology in every-day life is expected to become a reality in the near future. Yet, as RFID tags enter consumer households and threaten to identify their owners' belongings, whereabouts and habits concerns arise about the maintenance of privacy. People are afraid of being 'scanned' or tracked with the help of a technology that is invisible to them and not under their control. To address this consumer concern standardization bodies such as the Auto-ID Center have proposed to integrate a kill functionality into RFID tags. The present article argues that killing tags at the store exit is, however, not a viable long-term strategy to ensure default privacy. Too many business models and services are already in the pipeline to use RFID functionality after a purchase has taken place. Economic interest and consumer benefits risk undermining widespread tag killing. As a response to this dilemma we propose a simple disable/enable mechanism. Our suggestion is to disable all tags by default as part of the shopping check-out process and provide consumers with a password that enables them to re-enable their objects' tags if needed.

Keywords: RFID, Privacy Enhancing Technologies, Privacy

1. Introduction

RFID technology is a major enabler of ubiquitous computing environments or the pervasive Internet as described and researched by technologists. Today, the technology is introduced to facilitate supply chain

management. Yet, as the technology's cost decreases it also allows for new business models and applications beyond logistics. In fact, manufacturers, retailers and consumers can all take advantage of the technology's ability to uniquely identify objects, view their characteristics and relate to their owners. Intelligent home environments, improved reclamation and recycling processes, brand protection, safety- and security applications, but also less queuing time in super markets and more personalized information services count among the myriad benefits expected from 'living' RFID tags at the item level. Due to these benefits we argue that it is unrealistic to expect RFID tags to be systematically killed at store exits.

As this is true, considerable privacy concerns are accompanying the introduction of RFID technology. Public debate is rising over the potential presence of smart chips in all of peoples' belongings. Privacy rights organizations call for a complete abolition of tags in all those areas where they can be in touch with people [12]. *Uncontrolled* technology surrounding us and even in our cloths opens up a whole new dimension for the privacy debate which has the potential to considerably damage well established brands (figure 1). As a result, we argue that industry investment in privacy enhancing technologies (PETs) along with proactive transparency should be part of any RFID introduction strategy.

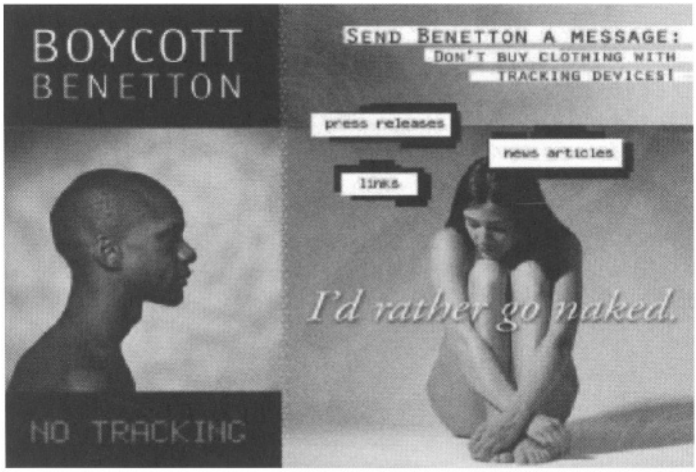

Figure 1. Example of how the privacy debate can impact the brand

Certainly, privacy is a multi-layered challenge when it comes to RFID. Section 2 will give a brief overview of the issues raised. Yet, what has

strongly dominated public debate so far is peoples' fear to be spied on by others, to be scanned and tracked. The immediate response by technology developers and early adopters has been to integrate a kill function into the specification of RFID tags (see section 3). Yet, even though 100% killing of all tags would be the perfect privacy solution economic interest risks undermining this approach to be used by *default.* In section 3, the current article therefore suggests to replace the kill function with a disable/enable function. The disable/enable model does not prohibit RFID tags' after-sales benefits. Instead it puts the use of RFID tags under peoples' control who can re-enable tags any time they need to (with a personal password). With this, our solution integrates industry, consumer and visionaries' interests.

One major cornerstone of our proposal is the *default* disabling process we recommend for supermarket check-out systems. Even though the discussion of such an automated check-out system is not subject of this article it still is an important requirement to make our solution work from a privacy perspective. While tag killing could only be applied to those goods where there are definitely no after-sales use scenarios, tag disabling can always be applied to all goods without after-sales sacrifices.

Section 4 closes with an acknowledgement of the challenges accompanying our proposal, especially password management, tag- and infrastructure cost.

2. Impact of RFID on Privacy

Consumer privacy is discussed today on the basis of three distinct temporal phases (see figure 2): in the retail outlet, at the retailers' check-out and outside of the retailers' premises.

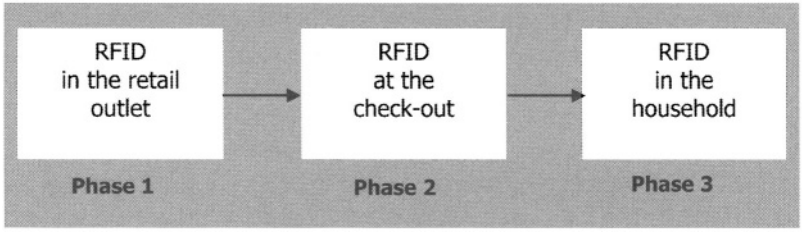

Figure 2. The 3 phases of the RFID Privacy Debate

RFID in supermarket premises (phase 1) allows for the creation of comprehensive profiles on how people move through the store [14]. These can be used to analyze how they buy in a similar way as is the case

today for web click-stream data collected in Internet stores. Privacy activists consequently call for not using RFID tags in retail outlets [12]. Especially when RFID tracking data is combined with video surveillance techniques concerns are high [4]. Early in-store trials, e.g. at the GAP were stopped [9].

Assuming that RFID will be used in retail outlets, privacy activists have stressed the point that at least when paying for goods (phase 2) it is unacceptable to have consumers queue again for deactivating tags [4]. Other sources call to "… prohibit merchants' pressure tactics to coerce keeping the tag alive…" ([11] citing a hearing before the California State Senate). As a result, deactivation, whether this means killing or disabling of the tag, needs to be integrated into the payment process.

Another thread of fear relevant at the supermarket check-out (phase 2) is concerned with the combination of highly granular EPC data with personal identity data[1]. Personal identity data is usually collected with the help of loyalty cards. Combining a person's identity at the moment of purchase with such detailed product information allows for a degree of psychographic segmentation of individuals that has not been available before.

Finally, direct abuse is feared of RFID tags' being read out uncontrolled and unnoticed of by unauthorized readers (phase 3). Thus, privacy could be intruded if people or institutions with readers were able to read out unrecognized on the belongings and whereabouts of others. This fear is fueled by the fact that information stored on an RFID tag can be read out unnoticed and from a distance.

3. Privacy Enhancing Functionality For RFID Tags

3.1 Background

The discussion has shown how and why privacy concerns arise around RFID technology. In the remainder of this paper we will focus on how privacy could be enhanced in phase 3, thus when people take RFID tags home and are tracked or read out unnoticed by others.

Version 1.0 of the EPC Network Specification [6] distinguishes several tag classes (currently form 0 to 5) depending on their sophistication as far as memory, power supply and communication range is concerned. A kill function is foreseen in conjunction with an 8- or 24 bit password scheme even for the simplest and lowest-cost type of tag class 0 and 1.

However, as we have outlined in section 1, economic interest is likely to impede a widespread killing of tags by default. Therefore, we would suggest replacing the kill function in the specification with a password

protected enable/disable mechanism. Depending on product nature and value we would propose two types of privacy enhancement with different levels of security and tag cost attached to them. In essence we argue that read-only chips (classes 0 and 1) should not be the long-term mass market solution for item level tagging. From a privacy perspective we strongly believe in the necessity to use tags with some write-capability in order to integrate long-term viable privacy functionality. Table 1 gives a requirements overview of privacy functionality foreseen for class 0 and 1 tags in comparison to our proposal described below.

RFID tag specification	**Class 0/1**	**Type 1 priv. enhmt.**	**Type 2 priv. enhmt.**
Memory			
ROM	**X**	x	x
EEPROM	x	**X**	**X**
Objects in Memory			
8 or 24 bit password related to kill function	**X**		
24 bit password to disable or enable the EPC		**X**	**X**
Status (enabled/disabled)		**X**	**X**
Operations			
Kill EPC function	**X**		
Verify (kill-)password	**X**		
Cryptographic one-way function			**X**
Disable function (to disable EPC based on password)		**X**	**X**
Enable function (to enable EPC based on password)		**X**	**X**
Verify password to disable/enable EPC		**X**	**X**
Change password		**X**	**X**
Generate random number (pseudoRNG)			**X**
XOR function			**X**

Table 1. RFID tag functionality relevant in the privacy context and potential enhancements

3.2 The Disable Model

The enhancement we propose is to integrate a disable/enable-function instead of a kill function into tags. We distinguish between two types of disablement. Type 1 implies a simple exchange of the kill function with a disable-function. The goal here is not to provide for perfect cryptographic protection of tag information, but to have good-enough protection in place to prohibit wide-spread tracking and spying. This is suitable for low-cost goods. Type 2 privacy enhancements include a more sophisticated crypto-based password scheme similar to proposals

of other researchers (e.g. [15]). This type of more cost intensive privacy enhancement only makes sense in the context of high value goods.

- **Type 1 privacy enhancements**

The way we envision the disabling process to flow is as follows: Instead of storing the kill password and function, the RFID tag stores a 24-bit *enable/disable* password and function. When a consumer pays for his products all tags are by default and automatically disabled. The disabling process is handled by the cash-registrar in order to avoid consumer time cost. With disablement a new password is randomly set on all tags. This one password is printed out on the customer's receipt[2]. It can be used by the new product owner to potentially re-enabling the EPC if needed for recycling, reclamations or intelligent home applications.

If unauthorized reader devices request the EPC from a disabled tag without the correct password the tag denies access to the EPC stored on it. From a layman perspective this means that *by default* objects bought do not communicate with any reading device except at one's personal request. The approach thus lends itself to calm all those privacy concerns related to unauthorized tracking and spying. At the same time, all economically driven intelligent home appliances and future consumer information needs are maintained. Trust in back-end reader architecture is not required. Control resides completely with the user.

From a technical perspective, of course, the tag still reacts to process re-enable requests. At this point several issues can arise from a security perspective: The most important one is that it is possible for an adversary to not decipher the password, but instead mime an anti-collision procedure. Anti-collision is a function used to uniquely recognize and communicate with one tag when several tags respond at the same time. If anti-collision is now based on the EPC - the structure of which is standardized - our disable-proposition could be circumvented. Our solution therefore relies on the fact that the EPC itself is not used for anti-collision. At first sight, this may be considered a major drawback of our solution. Yet, requirements in logistics suggest that full EPCs are not suited as a numbering scheme for anti-collision anyways. Forging through a full EPC is too time consuming. Therefore, other numbering schemes have been proposed for anti-collision including EPC dependent hash-values, a random number pre-integrated in the tag, RNG integrated into tags or a 12 to 14 bit serial number extract from the full 96 bit EPC [5]. For all these suggestions, our solution is feasible.

The second security weakness that may be argued is that a 24-bit password scheme is not a 'good-enough' protection. We argue that the

effort required by an adversary to decipher a 24 or 32 bit password is not worthwhile if the result is nothing, but the EPC of a low-cost/low involvement product. We therefore argue that the cost-benefit rationale of most adversaries in most situations will effectively protect consumers.

The third drawback is that there will be authorized readers (e.g. at the cash-register or in the consumer's smart home) which send the new owner's password around in plain text without encryption. A serious attacker, e.g. a thief, could therefore sniff on the cash-register or home environment and retrieve the password. Again, we would argue that for low-cost products the incentive for thieves or other adversaries is rather low.

Yet, for higher value objects (such as CD players, TVs, etc.) a systematic 'spying-attack' of this sort, e.g. on private homes could be realistic. Consequently we argue that for higher-value goods another (more sophisticated) password scheme may be necessary referred to here as type 2 privacy enhancement.

- **Type 2 privacy enhancements**

In order to defeat sniffing practices on high value goods, type 2 privacy enhancements foresee a challenge response method to verify the password. This method is based on a typical cryptographic one-way function [3]. First the tag sends a randomly generated value to a reader. Here, a pseudoRNG may be the most realistic solution for 'good-enough' security, where a standard RNG solution is too costly. The reader answers with a combined hash from the random value and the password. Using the same one-way function, the tag can then verify the reader's password.

The vulnerability of this procedure is that in the moment of resetting the password the new password is transmitted in plain text. An adversary could thus sniff on the new password (e.g. at the cash-register). In order to defeat such an attack the Vernam-Chiffre, a simple XOR function using the old password as the key to encrypt the new password can be applied for password re-set [13].

Compared to solutions proposing published hash functions or symmetric encryption for RFID environments [10, 8, 15] our solution does not require a database for personal tag management. Only one common password is used by a consumer or household. Switching product ownership implies just two password changing steps using a randomly selected temporary password. Key management is equally not required. This makes our solution more cost efficient and less complex.

4. Discussion

Obviously, both types of privacy enhancements imply additional cost for tag manufacturers. The most important cost driver is that the privacy enhancements we propose require tag manufacturers to use non-volatile and re-writable memory (e.g. EEPROM) instead of ROM for all item-level tags. Even though this is generally foreseen for tags of class 2 and upwards, the current specification does not include it for those low-cost tag classes 0 and 1. In addition to this memory cost the tags would need to be able to integrate two (or even five) additional functions[3].

Disabling a tag as we propose here only from time to time does not make sense. Our proposal integrates the requirement that the disable process itself takes place automatically when goods are checked out at the cash-register. While the disable model allows for default privacy and is therefore superior to the kill function industry players will argue that integrating disablement in cash-registers is costly. We argue that this may be true, but privacy needs justify the investment. With RFID cash-registers will undergo considerable technical changes in any way. Disabling will only be an additional requirement.

Password management can be a challenge in moving goods through the supply chain as well as in the user domain. Yet, as far as logistics is concerned our proposal is identical to the kill model. Probably, password information is transferred along with EPC information. When consumers take products home future scenarios foresee home agents and identity management systems [2] which manage peoples' assets, data and access rights[4]. In our thinking, such an agent could check new goods into the home system and set all devices to one common home password. Consequently, future consumers would not have to remember a myriad of passwords for each product. We believe that common password architecture for home readers or smart homes makes sense as consumers can access their devises more easily. A back-end database containing all tag data as proposed by Weis [15] as well as processing infrastructure to test all possible passwords [7] is not required. In the short- and mid-term, passwords printed out on receipts also don't increase consumer transaction cost since proof concepts in recycling and reclamation have been based on receipts for the last decades.

Finally, from a security perspective our proposal does, of course, not allow for highest level protection as is needed in many application areas. However, we also do not believe that military-level security is required for yogurt cans or even stereos. Even if the 'one-common-home-password'

we suggest would be decrypted, what would the thief learn more about my belongings than if he just unlocked the window and stepped in?

5. Conclusion

RFID technology will be a ubiquitous reality in every-day life in the future. This paper argues that economic interest seeks to maintain an RFID tag's functionality after a purchase has been made. On this basis it is argued that killing RFID tags is an unrealistic solution to preserve *default* privacy in the long run. The authors conclude that mass market RFID should be enhanced with privacy functionality which in our proposal implies write-enhanced memory. Two types of privacy protection are suggested implying different cost and sophistication.

The major benefit of the solution outlined is that the disable-model puts RFID communication into the sole control of the user. With this, the solution embraces current thinking in the development of PET technologies which takes a user-centric view. Secondly, a compromise is made between state-of-the-art security and what is economically feasible. Only 'good-enough' security is used to develop a proposition that will meet the privacy needs in a majority of situations. Finally, the model is the only proposition to our knowledge which allows for a realistic compromise between RFID-based market aspirations and business models on one side and peoples' desire for privacy on the other. Consequently, we believe that the disable-model is a good road to take.

Notes

1. Similar to the bar code, the Electronic Product Code, EPC, contains a serial number that can be related to a product category and a manufacturer. However, the EPC also contains a unique serial number associated with more detailed and comprehensive back-end data. This allows for retrieving an object's detailed characteristics, history and potentially other related data [1].

2. Long-term, the password will probably be transferred to an identity device such as a PDA owned by the consumer.

3. In fact, the low cost RFID tags "Philips I-CODE SL2 ICS10/11" already contains all components needed for type 1 privacy enhancements, needing only a few design changes.

4. For a reference on agent solutions currently developed to address the challenge of increasingly complex password management see e.g. HP's work on the 'e-person': http://www.-hpl.hp.com/research/iil/themes/eperson/eperson.htm

References

[1] Auto-ID Center. Technical memo - physical mark-up language update, p.5, 2002.

[2] S. Clauß and M. Köhntopp. Identity management and its support of multilateral security. *Computer Networks,* (37):205–219, 2001.

[3] Ivan Bjerre Damgård. Collision free hash functions and public key signature schemes. In *Eurocrypt '87,* volume 304 of *LNCS,* pages 203–216. Springer-Verlag, 1988.

[4] FoeBuD e.V. Positionspapier über den Gebrauch von RFID auf und in Konsumgütern, Presseerklärung. http://www.foebud.org/texte/aktion/rfid/positions-papier.pdf, 2003.

[5] EPC Global. Specifications for 900 MHz Class 0 RFID Tags, page 15. http://www. epcglobalinc.org/standards_technology/Secure/ v1.0/UHF-class0.pdf, 2003.

[6] EPC Global. Version 1.0 Specifications for RFID Tags. http://www.epcglobalinc.org/standards_technology/specifications.html, 2003.

[7] A. Juels. Privacy and Authentication in Low-Cost RFID Tags. Submission to RFID Privacy Workshop @ MIT, 2003.

[8] Shingo Kinosita, Fumitaka Hoshino, Tomoyuki Komuro, Akiko Fujimura, and Miyako Ohkubo. Nonidentifiable Anonymous-ID Scheme for RFID Privacy Protection. To appear in CSS 2003 in Japanese, 2003.

[9] Meg McGinity. RFID: Is This Game of Tag Fair Play? *Communications of the ACM,* 47(1):15, 2004.

[10] Miyako Ohkubo, Koutarou Suzuki, and Shingo Kinoshita. Cryptographic Approach to "Privacy-Friendly" Tags. Submission to RFID Privacy Workshop @ MIT, 2003.

[11] Gregory J. Pottie. Privacy in the Global E-Village. *Communications of the ACM,* 47(2):21, 2004.

[12] Peter Schüler. Dem Verbraucher eine Wahl schaffen - Risiken der RFID-Technik aus Bürgersicht. *c't,* (9), 2004.

[13] C. E. Shannon. Communication Theory of Secrecy Systems. *The Bell System Technical Journal,* 28(4):656–715, 1949.

[14] S. Spiekermann and U. Jannasch. RFID in the retail outlet: implications for marketing and privacy. IWI Working Paper, 2004.

[15] S. Weis. *Security and Privacy in Radio-Frequency Identification Devices.* PhD thesis, Massachusetts Institute of Technology (MIT), 2003.

SAFEGUARDING PERSONAL DATA USING TRUSTED COMPUTING IN PERVASIVE COMPUTING

Adolf Hohl[1], Alf Zugenmaier[2]

[1]*IIG dep. Telematics*
University of Freiburg
adolf.hohl@iig.uni-freiburg.de

[2]*Microsoft Research Cambridge*
alfz@microsoft.com

Abstract Pervasive computing can be divided into computing on personal wearable devices and computing in a smart infrastructure. When a wearable device communicates personal data for further processing to the infrastructure, privacy concerns arise. This paper presents an approach to dispel concerns relating to improper use of personal data based on digital rights management technology. A prototype implementation of this approach in a smart hospital environment is described.

Keywords: Privacy, policy enforcement, trusted computing, attestation

1. Introduction

The paradigm of ubiquitous and pervasive computing [16] leads to a much greater intrusion of information and communication technology into the personal life of everyone than what we experience today. The users of pervasive computing will use many smart personal objects, in addition, many services will be provided by a smart environment that will surround us. However, fears of users about the misuse of their personal data prevents the acceptance of these services and technologies. This is especially the case, when an agent running on a personal digital assistant is acting on behalf of the user and can autonomously release sensitive information to communicating partners such as service providing devices in the environment. Nearly everybody has had experience of misused personal information in the Internet such as unwanted adver-

tisements and spam. This is only the tip of the iceberg. More serious abuse of the information may involve selling it to rating agencies, resulting in unwanted "personalization" of prices, interest rates, denial of credit, etc. Therefore, it is essential that devices providing services handle their users' personal data with care. If it is not possible to ensure this, fear of misuse and privacy concerns remain with the user.

1.1 Problem Statement

In this paper, we address the problem of giving users of pervasive computing environments more control over their data after they are transmitted, e.g., during the use of a service or an application. Privacy issues can never be addressed completely without looking at the application domain [5]. Therefore, we make use of the scenario provided by the project EMIKA at the University Hospital of Freiburg [11][1]. In the hospital scenario, patients are equipped with a smartcard which can store the patients' health history or parts thereof[2]. In this scenario, patients can have access to the content of their smartcards and supplemental information, which is linked to other sources of information on this card or external to it. EMIKA envisions an infrastructure of public terminals or displays in the hospital in addition to the patients' personal devices[3]. It is necessary that the personal health information which can be processed by an application on a public terminal or display cannot be misused. Misuse can take two forms: alteration of the stored information by unauthorized parties, and privacy of the patient's health history. Potential solutions to the first problem were proposed by introducing different types of acces control models, see e.g. [7] and [3]. Therefore, this paper focuses on the second problem: how to make sure that the patient's information is not misused. The example which will be used throughout this paper is a public terminal with a browser that allows viewing of the information stored on the patient's smartcard and on file in the hospital database (cf. Figure 1). The public terminal or display has to forget the content and the actions performed after the patient ejects her smartcard, leaving no information about her health history in the browser cache. The same applies to a printer which may have been used during the session. In general, a service or an application is used, which may not be in the patient's or hospital's administration or trust[4] domain, therefore it is uncertain that sensitive personal data are treated in the expected way. The public terminal in the untrusted zone[5] enables access to files on the trusted smartcard and access to linked external information, for instance, X-ray images. The smartcard is viewed as trusted because it is owned by the user. The external database is

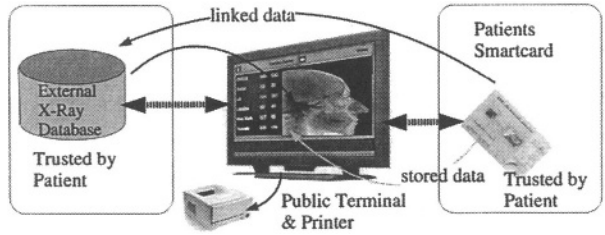

Figure 1. Architectural Overview

maintained by the hospital and is therefore also trusted. The terminal communicates directly with the smartcard and external sources.

2. Attacker Model

The aim of an attacker in this scenario would be to gain access to private health information. The attacker may gain control over some of the software on the public terminal, or gain complete control over the terminal after the user left it. He may read and insert communication between the smartcard and the terminal, or read and insert communication between the terminal and the backend database. In addition, the attacker may also introduce a fake terminal. An attack that involves an attacker looking at the display of the terminal is not considered.

3. The Approach

Our approach follows closely the idea presented by Korba and Kenny in [8] for solving the problem that a user can keep control over transmitted personal data is based on the following observation: the interests a service or application user has in dealing with sensitive data are similar to those of providers of copyrighted digital contents. Both, the copyrighted content provider and the patient, i.e., the personal data provider, are interested in making the supplied data available only for limited use and processing. Furthermore, unauthorized onward transmission and use should be prevented. Subsequently, control over transmitted data or contents has to be enforced.

This parallelism of interests between content providers and patients (service users) with regard to the processing of data makes rights management mechanisms a suitable toolset for the protection of sensitive personal data. Personal data is sent in a protected way to the service-providing device preventing unauthorized usage and information leakage.

Sensitive personal data has a license attached to it when communicated to the service providers. The license is issued for a single device only and limits the use of this personal data to this device. The service user now takes the role of a content provider and license issuer. Because it would be unmanageable if every patient had her own slightly different license attached to her data, patient interest groups should act as liaison and offer standardised licenses.

This is a contrary approach to classical anonymization techniques with the concepts of data minimality and data obfuscation, because a technical implemented temporal extension of the domain of trust is used, which prevents misuse.

4. Technical Solution

Successful deployment of a system which enforces the processing of personal data under given usage restrictions requires an independent processing or reporting component. This component can ensure or report that the applications which are executed are untampered with and provide a safe execution environment. The Trusted Computing Group [15] is developing extensions to computing platforms to ensure this. Because major industry players, including hardware and software manufactures and content providers, are involved in specifying this platform one can assume that devices with Trusted Computning or DRM capabilities will become pervasive. The TCG platform can produce signed attestations of the integrity of the software and report by this the execution of an untampered application.

Technically the TCG specifies hardware extensions by which different stages of starting and running a platform can be verified by measurement functions and reported to the TPM. By this, the trusted domain is extended with every successful verified component (BIOS, firmware of devices, bootloader, operating system). This extension of trust is illustrated in figure 2. If the platform has successfully started and all the hash values of the measured components matches the expected values of a known state platform, the TPM unlocks signing functions to be able to prove its known state. Microsoft proposes an operating system with the so called Next Generation Secure Computing Base NGSCB [10] which extends the existing context in which a process can be executed with a secure context environment. Only verified code can be executed in this protected context. Debugging or accessing other processes' memory is not possible and should be supported by a special processor mode in future. ARM the well known microprocessor designer as well proposes a

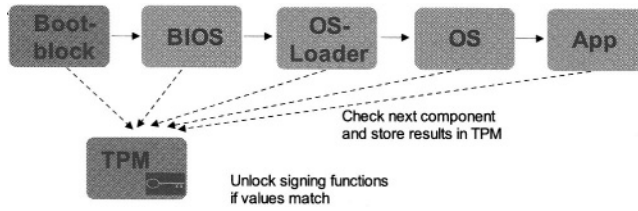

Figure 2. Boot procedure with code verification

model [2] with a couple of similarities, especially the division of context in a normal side and a secure side.

To decide if the platform which should process sensitive personal data behaves as it claims to, one has to know about the software and the platform. Trusted Computing mechanisms can guarantee a proper and verified execution. But it will be hard to know about all soft- and hardware components and about different versions of them. This makes a third party necessary to classify software and hardware components as trustworthy or possible to build trustworthy platform on it.

4.1 The implementation

In our proof of concept implementation a trustworthy platform e.g. with TCG compliant TPM wasn't available. This means the core root of trust cannot be the TPM chip. Instead we treated the used operating system with code integrity checking functionality as the core root of trust and the information about executed software on this system are reliable. We also excluded the use of a third party software component evaluator. The user, respective the users device knows how the terminal has to look like. A secure execution context comparable to NGSCB was also not available. To simulate the attribute of obliviousness (after the terminal was used, it should forget about everything) the application with the user data is executed from a ramdisk which is reformatted after the usage. To allow rapid protoyping, the smartcard functionality was implemented on a PDA. In figure 3 the interaction between the smartcard simulating PDA and the terminal is shown.

For the hospital environment, we extend the functionality of the smartcard with the capability of verifying these attestations and, thus checking the integrity of the public terminal or display. The current implementation of the public terminal is based on the NetBSD operating system [13] with a trusted path execution functionality, the so called

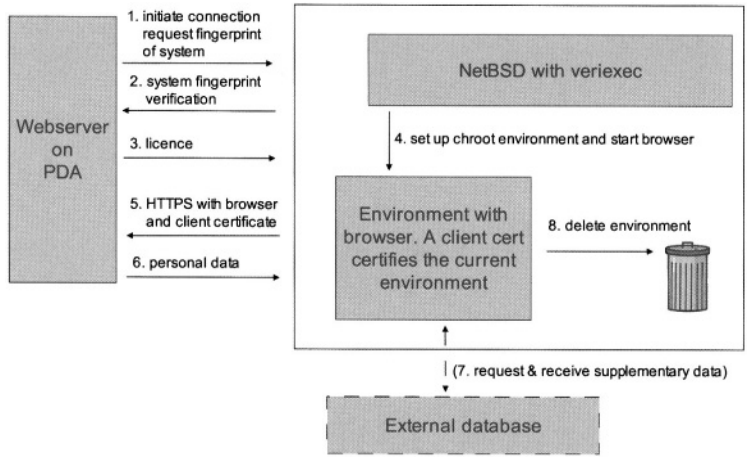

Figure 3. Phases during the use of the terminal

VERIEXEC option which supports the execution of applications with a valid hash signature and location only.

The public terminal runs a trusted daemon waiting for events to to interact (step 1 in figure 3) with the PDA used to emulate the smartcard. An IP connection between the terminal and the PDA is established. It is assumed that the operating system itself is a root of trust and attestations concerning executed software are valid. By this connection, the PDA requests a hash of the system of the public terminal. The daemon responds over a secure connection and reports the system state via sending a fingerprint of the software on the public terminal (step 2). After that, the PDA decides it the public terminal is known and respects given usage limitations on transmitted personal data. If it is a known one, the PDA issues a certificate with a licencse (step 3). The licence contains a list of access rights which in this implementation is either *view* or *view and print.* Based on these rights, the daemon sets up a chroot environment on a ramdisk (step 4) with or without a printer device. This certificate allows the applications on the public terminal to access personal data via a secure HTTP connection on the PDA (step 5 and 6). The personal data may contain links to external documents, like X-ray images (step 7). These hyperlinks are HTTPS hyperlinks with embedded login information to the external patient information

database. The daemon that set up the environment continually polls the PDA to find out if there still is a contact. If the contact is lost for more than five seconds, the ramdisk is deleted and, thus, no trace of the personal data left on the device (step 8).

5. Discussion

The implementation represents a first step towards using cooperative mechanisms to protect the privacy of users which are reported and enforced technically at the public terminals. The used operating system supports a verified execution but in itself can not represent the same core root of trust as trusted computing hardware. The PDA can issue the right to view and print. Printing is a digital transfer of sensitive data to another device, the printer. This means that the printer itself should have to respect the terms of the licence. Currently, a printer without permanent storage is used.

The implementation described in the previous section does not address the threat that the browser may be tricked into posting sensitive information to untrusted sites. To this end, further isolation of the network environment is required, similar to the isolation of the filesystem provided by the chrooted ramdisk.

The use of stunnel [14] and HTTPS is very computation intensive for the user's device. Using NGSCB-like enforcement mechanisms could reduce this load and lead to a solution closer to the capabilities of a real smartcard.

6. Related Work

There is some work that is related to the approach presented here. As stated before, the idea of using DRM like mechanisms and rights expression languages for the protection of personal data was discussed by Korba and Kenny [8]. However, they did not present a working system architecture or proof of concept implementation. Bussard et al. [4] demonstrate how to display sensitive information in federations of devices. However, their approach doesn't work if the information is too complex to be displayed on a limited screen (e.g. x-ray pictures). Kohl [7] pointed out that privacy is in fact a big issue in a hospital environment, but assumed a central organization for data storage and processing. Privacy through the use of identity management in a mobile computing environment is proposed in [6]. It is based on the retention of personal data and can not be controlled once they are given in foreign hands. Agrawal et. al [1] attach a licence to data in a database. This

approach is a good way of ensuring privacy as long as the data does not cross administrative domain boundaries.

Closer to the method presented here is the suggestion of Langheinrich in [9]. A policy is attached to personal data to create a sense of accountability. The approach of Mont et al. in [12] uses a third party to trace and audit the use of personal information.

7. Conclusions and Future Work

The results from the first trials are encouraging and lead us to believe that mechanisms similar to rights management can be used to enforce privacy. The setting in the hospital is almost ideal for such an approach of cooperative privacy. It can be expected that only few different companies will provide equipment for the hospitals. Hospitals are highly regulated and, therefore, there is interest by the hospital to ensure privacy. Additionally, this approach can be used to shift the work of ensuring the correct handling of data from the person installing and maintaining the pervasive computing environment to the software vendor for the viewer of the data.

Future work includes a port of the current implementation to NGSCB and a closer look at certificate management and revocation. In addition, different DRM systems approaches have to be evaluated to find out which one supports the need of handling of personal data. It will also be interesting to implement the certificate validation on a smartcard to verify the performance.

Notes

1. We would like to point out that these issues are not limited to the hospital environment and also appear in other areas, for instance in e-commerce and web-services in general.
2. Smartcards like this are currently being specified and will be used in the near future in the German health system under the name "Krankenkarte".
3. While every patient will be supplied with a smartcard, not every person will own a PDA.
4. Here, trust is defined as the patient being confident that her data is not misused.
5. Current terminals has to be considered untrusted as long as a user of the terminal has no chance to convince hisself. It is easy to tamper a terminal with given its public location, while it is very hard to administrate it such that it remains tamper resistant.

References

[1] Rakesh Agrawal, Jerry Kiernan, Ramakrishnan Srikant, and Yirong Xu. Hippocratic Databases. In *28th Int'l Conf. on Very Large Databases (VLDB), Hong Kong,* 2002.

[2] ARM. TrustZone Technology - Secure extension to the ARM architecture. 2004.

[3] G. Brose, M. Koch, and K.-P. Löhr. Entwicklung und Verwaltung von Zugriffsschutz in verteilten Objektsystemen - eine Krankenhausfallstudie. 2003. Praxis der Informationsverarbeitung und Kommunikation (PIK).

[4] Laurent Bussard, Yves Roudier, Roger Kilian-Kehr, and Stefano Crosta. Trust and Authorization in Pervasive B2E Scenarios. In *Proceedings of the 6th Information Security Conference (ISC'03) Bristol, United Kingdom, October 1st-3rd,* 2003.

[5] G. Iachello and G. D. Abowd. Security requirements for environmental sensing technology. 2003. 2nd Workshop on Ubicomp Security, Oct. 2003, Seattle, WA, USA.

[6] Uwe Jendricke, Michael Kreutzer, and Alf Zugenmaier. Mobile Identity Management. Technical Report 178, Institut für Informatik, Universität Freiburg, October 2002. Workshop on Security in Ubiquitous Computing, UBICOMP.

[7] Ulrich Kohl. From Social Requirements to Technical Solutions - Bridging the Gap with User-Oriented Data Security. In *Proceedings IFIP/Sec '95, Cape Town, South Africa, 9-12 May,* 1995.

[8] Larry Korba and Steve Kenny. Towards Meeting the Privacy Challenge: Adapting DRM. 2002. ACM Workshop on Digital Rights Management.

[9] Marc Langheinrich. A Privacy Awareness System for Ubiquitous Computing Environments. 2001.

[10] Microsoft Corporation. *NGSCB: Trusted Computing Base and Software Authentication,* 2003.

[11] Guenter Müller, Michael Kreutzer, Moritz Strasser, and et al. Geduldige Technologie für ungeduldige Patienten, führt Ubiquitous Computing zu mehr Selbstbestimmung? In *Total vernetzt. Springer: Berlin, Heidelberg, New York,* pages 159–186, 2003.

[12] M. Mont, S. Pearson, and P. Bramhall. Towards Accountable Management of Identity and Privacy: Sticky Policies and Enforceable Tracing Services. 2003. HPL-2003-49.

[13] NetBSD. http://www.netbsd.org. 2004.

[14] stunnel. www.stunnel.org. 2004.

[15] Trusted Computing Group. TCG Backgrounder. May 2003.

[16] Marc Weiser. The Computer of the 21st Century, 1991. Scientific American, vo.265. no.3, Sept.1991, pp 66-75.

Part of this work was funded by the DFG / Gottlieb Daimler and Carl Benz Foundation.

This page intentionally left blank

A SOCIAL APPROACH TO PRIVACY IN LOCATION-ENHANCED COMPUTING

Ian Smith[1], Anthony LaMarca[1], Sunny Consolvo[1], Paul Dourish[2]

[1] *Intel Research Seattle 1100 NE 45th St, 6th Floor, Seattle WA, 98105, USA*
{ian.e.smith, anthony.lamarca, sunny.consolvo}@intel.com

[2] *School of Information and Computer Science U.C. Irvine, Irvine CA, 92697, USA*
jpd@ics.uci.edu

Abstract Place Lab is a system for positioning a user based on passive monitoring of 802.11 access points. Place Lab seeks preserve the user's privacy by preventing disclosures, even to "trusted" systems in the infrastructure. We are pursing two avenues to explore these and other privacy issues in the domain of socially-oriented applications. We are doing fieldwork to understand user needs and preferences as well as developing applications with significant, fundamental privacy concerns in order to expose the strengths and weaknesses in our approach.

Keywords: Privacy, location-based services, location-enhanced computing, ubiquitous computing, context-aware services

1. Introduction

Privacy has long been recognized as a central concern for the effective development and deployment of ubiquitous systems [2, 3, 13, 15, 17]. As both a technical problem and a social problem, it is difficult to deal with, to design for, and to model coherently.

The traditional frame within which privacy arguments are cast is a trade-off between risk and reward. This is a popular approach in a range of fields from public policy to cryptography. The risk/reward framework, in the pervasive computing context, suggests that individuals make decisions about technology use by balancing perceived risks against anticipated benefits-that is, in a fundamentally economic approach, they trade off costs against benefits and adopt technologies in which the benefits outweigh the costs, while rejecting those in which the

costs outweigh the benefits. Therefore, many have argued, creating successful location-enhanced computing requires finding the most effective balance between risks and rewards [10, 25].

This approach has a number of problems, though, both as a conceptual framework and, consequently, as a model for design. Studies of actual practice fail to display the sort of rational trade-off that this model would suggest. There are a number of possible reasons.

First, it is likely that the model is over-simplified and neglects a number of related factors that are important for decision-making about technology adoption and use. For example, we have found that naturally-occurring accounts of privacy behaviors depend on *recourse* as much as risk and reward. By recourse, we are referring to the actions that can be taken by users in the event that others misbehave.

Second, recent research in the area of behavioral economics suggests that traditional rational actor approaches fail to adequately account for everyday behavior even within their own fairly limited terms of reference [22]. The notion of stable exchange-values for goods, services, and labor upon which conventional economic modeling is based seems to fare poorly when applied to human actors who are meant to embody these principles. Instead, psychological and social factors seem to interfere with the mathematical principles of neoclassical economics. In a simple example, while you might pay a neighborhood kid $20 to mow your lawn, you would be less likely to mow your neighbor's lawn for $20. Recent approaches that attempt to incorporate psychological elements into economics models, such as prospect theory, revise traditional notions of commodity and value.

Third, and perhaps more fundamentally, studies of technological practice suggest that technology adoption and use should be seen not simply in terms of individual decisions about costs and benefits, but rather in terms of broader patterns of participation in cultural and social life. For example, in Harper's (1992) study [11] of the use of active badges in research laboratories, it is telling that a number of people report participating in the use of the system in order to be seen as team players, in order to provide support to others, etc. In other words, social actions have symbolic value here, and these are frequently the more salient elements of adoption decisions. Ito's studies of mobile messaging amongst Japanese teens [14], or the studies by Grinter and colleagues of the use of SMS and Instant Messaging amongst teens in the US and the UK [8, 9] describe the use of messaging technologies as cultural practices, essentially casting the adoption of these technologies as forms of participation in social life. To use the technologies is simply part and parcel of appropriate social practice. As technologies become increasingly integrated

into everyday practice, rational decision-making about privacy trade-offs is increasingly irrelevant.

Fourth, studies of privacy management in the everyday world, drawing on studies in social psychology, suggest that privacy management is a much more nuanced and contingent phenomenon. Drawing on the work of Irwin Altman, Palen and Dourish [20] present a model of privacy as a continual and dialectical process of boundary regulation. These boundaries are not simply barriers to information flow, but are also the boundaries between self and other through which differentiation and affiliation are achieved, and boundaries between past and future that reflect the emergence of genres or conventions for information practice. Some of this can be seen in studies of personal web pages [6] and increasingly, lately, studies of blogs (e.g. Nardi et al [19]) where genres arise that provide both expectations and interpretive norms for understanding information disclosure. For instance, where most personal web pages are unlikely to state the details of where people can be found at particular hours of the day, that is an appropriate and indeed expected form of disclosure for college professors. The dialectic model that Palen and Dourish propose has a number of important implications for design that are quite at odds with traditional rational actor approaches. Principally, they situate information disclosure settings within the immediate circumstances of activity, suggesting that the "costs" and "benefits" of information disclosure are continually subject to negotiation and change.

Finally, one implication of these broader perspectives for traditional approaches to the specification and description of location-based or context-aware computing is that the very notion of "context" may be problematic - it may not be something that can be uniquely defined, but depends on the person to whom the context is being disclosed, or the specific features of the setting in which the formulation is made [7]. We will return to this later, in describing a field study of the ways in which context is formulated [24].

Accordingly, in our work, we have been developing an alternative to traditional formulations of privacy, both as a conceptual framework and a technical approach. Our approach in Place Lab [26] attempts to avoid the abstract formulation of privacy needs and the "disclose and hope" model that requires them (see below).

Our essential argument, then, is that there are no abstract rules by which privacy is formulated; rather, the information practices that we refer to under the rubric of "privacy" are emergent phenomena of everyday social action.

One common objection to this argument is that, while rules and resources may not be part of our conscious experience of information prac-

tice, they must nonetheless be underlying factors, which we have learned and internalized so that they are no longer consciously available to use. We all had to be taught these rules, once upon a time; every one of us, after all, has a story of the moment when, as a young child, we loudly make some remark that was wildly socially inappropriate and embarrassing to our parents. So, the basis for our current practice must be rule-based, even though those rules are no longer part of our conscious experience.

However, this objection is fallacious. It is broadly equivalent to this argument-that when learning to ride a bicycle, we managed to stay upright through the use of training wheels. Once we became competent bicyclists, we no longer used training wheels but, even though the training wheels are no longer visible, they must, nonetheless, be the basis of our balance.

2. Classes of Location Applications

Place Lab is a research effort to build a low-cost, widely-available, indoor-outdoor positioning system [1, 4, 16, 12]. Devices running Place Lab use radio beacons in the environment (such as 802.11 access points) as known "way points" that can be used to triangulate location. Since an increasing percentage of computation devices are shipping with some radio sensing capability (such as 802.11 or Bluetooth), a map of known beacons allows them to position themselves with no additional hardware. One advantage Place Lab has over many other location technologies is that it is based on passive monitoring of radio signals and local lookups and computation. As such, devices running Place Lab can position themselves completely locally and only need disclose their location when it is desired by the user[1].

Our initial explorations with Place Lab have shown that location-enhanced applications fall broadly into three classes: institutional, social, and personal. These classes of applications differ based on the person or organization to whom location information is disclosed. A *personal* application is one that does not need to disclose location information to anyone to be effective. An example is a pedometer or other personal fitness applications. Another set of personal applications are way finding or route planning applications. These types of applications may need the user's location to function properly, but it is not necessary to communicate that location to anyone given local storage and possibly a cache of content.

Institutional applications are a more common arrangement, requiring that people disclose information to a central authority (normally, an or-

ganization) in return for some service. Active Badge systems [27] and related context-based services operate according to this model; information about location is relayed to a central server, while then makes contextualized services available to clients and users. This architectural approach made sense when both client-side computation and network bandwidth were limited, and so has been a common structure in prototype ubicomp systems. However, given the relentless march of time and Moore's Law, alternative technical approaches are now more feasible, and avoid the sorts of privacy commitments being made in this architecture.

It should be noted that it is possible to build the same institutional application with varying degrees of disclosure on the part of users. For example, if Google made their index of web pages publicly available, one could turn Google into a personal application since a user could do their searches while disclosing little to no personal information. In this scenario, one could download the entire medical index and then search locally for a specific condition, revealing the possible interest in a medical condition, but no more beyond that. However, in most cases, institutional applications have substantial commercial, public interest, or intellectual property barriers that prevent them from being organized in this open fashion.

The final class of applications in our taxonomy is *social.* These applications require disclosure to people, rather than institutions to work effectively. Many ubiquitous-computing services, such as friend finder [26] or context-aware chat [23], are examples of social applications. A friend finder is an application that alerts you when one of your "friends" is nearby, facilitating serendipitous social interaction. Clearly, this requires at least that the user and her friend's locations be exchanged in some way.

There are risks in social applications, although they are not as clear as some other scenarios. In the friend-finder example, by what mechanism should "friends" be designated? Certainly, it should require some type of mutual acceptance, otherwise the system can and will be abused by anyone with the technology. Avenues for recourse are also unclear. Are the forces of recourse-such as social isolation or embarrassment-strong enough to affect user behavior? With due respect to considerations of risks and recourse, we are more interested in how this technology will be adopted be social actors. It easy to imagine that being on someone's "friends list" in a friend-finder application might be as important as being in someone's cell-phone address book. Studies of the gift-giving practices of teens [28] have revealed the social impact of being "in" the social space of someone's cell-phone address book to be significant.

3. Social Applications, Privacy, and Place Lab

Previously, we argued that simple models that imply that people are rational actors making a narrow choice such as "will I give away this information for that commodity" are insufficient to explain the privacy-related behaviors we observe. If there are areas in which people can be seen as close to making rational choices it is the area of personal applications. Because disclosures to others are not required for personal applications, fewer social forces come to bear and an individual can make decisions "flying solo." This is not say that a simple risk versus reward calculation can be employed to predict user behavior-that would ignore issues such as user-interface concerns that still exist in personal applications. In the case of the pedometer personal application, issues such as size, weight, visibility to others, and battery life are quite significant to ultimate user adoption. Even the designation "personal" is troublesome here; if a pedometer is implemented in a "disclose and hope" fashion, the personal application takes on social dimensions as it can be used to track those that are walking with you.

Institutional applications are also problematic unless situated in their social context. Consider workplace-safety applications of location-tracking technology. Organizations and institutions might view this as a positive development, decreasing accidents or preventing workplace violence. Individuals who work for these organizations are likely to have many complex relationships to the deployment of such a technology [11] and the institution that deploys it. Yet again, the individual user's relationship to the organization and the deployed technology is not a simple matter of a trade-off in risk versus reward.

One of the goals of the Place Lab project is to build location infrastructure that will foster the development of successful applications. Unfortunately, as we have argued in Section 1, the inherent value of an application is a complex and unpredictable metric to predict. Of the three classes of application, the value of those in the social domain is the most unpredictable and often counter-intuitive. For this reason, we have chosen to initially focus our study of privacy and its relationship to location-enhanced computing on this class of application.

In the Place Lab project, we have begun two efforts to better understand the future space of social, location-based applications and how people will formulate the social norms governing their use. The first is a field-study to expose situated user concerns and the second is an application to help us directly experiment with these issues.

3.1 A Field Study Of Privacy Concerns

We are conducting a user study to understand people's perceptions about privacy and how time, place, and other people affect the types of disclosures they might make. In other words, we are trying to understand the social factors that would affect our future application development. Our study design uses the experience sampling method [5] or ESM (often called a "beeper study"). In an ESM study, a participant is given a mobile device such as a PDA that periodically alerts the user and asks a question(s). While incurring more overhead and interruption than techniques such as diary studies, ESM data is typically highly accurate as it is collected in situ and does not require recall.

In our case, this allows the participant to answer questions about location in the actual location, not in a lab or conference room days later. An additional advantage arises from the fact that participants will be carrying a computational device with them during the course of the study. Since we already have the Place Lab positioning infrastructure running on small devices, we can create questions that are customized based on the user's location. For example, through Place Lab, our ESM application might know the user's location and look up that location in a database of business records. We can then discover if the city or county business records, perhaps "Smith's hardware store," matches well with how users self-report their location. We believe that this comparison will shed light on how users' perception of risk varies with time and physical location.

Some examples of the types of questions we are designing into our study are:

- If your boss asked you for your location right now, how would you answer? Your spouse?
- If your mother asked where you were right now, would you answer 'a bank,' 'the corner of 45th and Vine,' or something else?
- Would you tell Alice your location right now in exchange for hers? If so, what would you be comfortable telling Alice? What would you want to know about Alice's location?

3.2 Ambush: A Dangerous, Yet Privacy-Aware Application

Rather than trying to develop locations-enhanced applications that skirt privacy issues, we have chosen the opposite approach. We have devised an application that we believe offers substantial new functionality

while at the same time presents significant privacy risks. In this way, we hope to attack the privacy issue "head on" by experimenting with privacy strategies and mechanisms.

Our application is called "ambush" and is based on the work of Mynatt and Tullio. In [18], Mynatt and Tullio describe an ambush as the use of a shared calendaring system to infer a person's probable location in the future with the intent of "ambushing" them for a quick face-to-face meeting. This process is used frequently in larger organizations, particularly by subordinates, to have brief conversations with senior managers who are between meetings.

We have generalized the notion of ambush to be any location, not just conference rooms visible in a shared calendar system at work. Our ambush application allows a user Alice to define a geographic region-say a public park-and ask to be notified anytime Bob enters that region. If Alice lives near the park and wants to visit with Bob, clearly both can benefit from the possible serendipitous, social encounter in the park. Another use of ambush is micro-coordination. Such tasks are common in urban environments, such as "Let me know when Charles or DeeDee get to the subway station so I can go meet them." Another use of ambush is the creation of social capital [21] through discovery of shared interests that are demarcated by places, such as bookstores, music venues, or civic organizations. It should be noted that current "friend finder" systems offered by cell-phone providers are actually corner-cases of our ambush application in which the only location that can be specified is "near me."

The potential for nefarious activities with ambush are rife, making risk a significant issue. As previously stated, we chose ambush as a test application because it forces to come to grips with the privacy concerns.

As an aside, we are not concerning ourselves right now with the security and authenticity issues of ambush. We are not addressing questions like, "How do I know that no malicious entity modified or hacked the users' devices to steal their location information?" or "How can I be sure that this geographic region is Green Lake Park as Alice purports and is not my home as I suspect?" Although these are interesting questions, we are focusing our initial investigations on the privacy issues.

We have devised several concrete strategies to help us address the privacy concerns in ambush. First, our privacy concerns field study with ESM mentioned above will include questions that are specifically tailored to an ambush-style application. This can help us craft our technical strategies to be sensitive to the social norms and perceptions of our user community.

Without going into tremendous detail, we are considering three significant techniques to blunt the privacy concerns in ambush. All of these are currently be explored through our early efforts.

- Reciprocity: If you get someone else's location you give up your own. Although this strategy is vulnerable to certain types of abuse-notably that people who do more things and go more places have more to lose than those that stay at home constantly-it offers some advantages. It allows those who disclose their location to know who requested the information; if the location offered in reciprocity is of little value ("always at home"), perhaps social norms of recourse can be used to deter abuse.

- Explicit acceptance: This seems central to our strategy of preserving privacy. You have to take explicit action to disclose your location, so it is at least possible for you to be aware of others' attempts to observe you, for good or ill. This has the obvious problem that it does not scale well to large numbers of disclosures of your location. Either you will become irritated with the frequent disturbances or become "numb" to the action and cease to really make a decision about the disclosure. Both this technique and the previous one are situated primarily the social domain for both the user's understanding of the disclosures being made as well as the possibilities for recourse.

- Indirection: Perhaps Alice should "make an argument" to Bob for the release of his information to her. In this model, Place Lab does not disclose Bob's location to Alice, but rather shows Bob Alice's argument (perhaps in text form) when he enters the park. "Bob: We should get our kids together in the park. Call me. -Alice." This technique can easily be combined with either of the first two for additional benefits. This is a similar to many systems that leave information at specific places in the world, but it is focused on the two users rather than leaving information "for anyone."

4. Conclusion

Despite being in the early stages of the Place Lab project, we know that accurately recognizing and addressing privacy concerns is critical to the success of our system as a platform for location-enhanced computing. Unfortunately, understanding disclosure of user's information and its relationship to an application's success is difficult to predict. This is especially true in the domain of social applications in which users disclose personal data to other individuals. To increase our understanding of

applications in this domain, we are running an ESM study to learn how location, context and place interact with a user's inclination to disclose information to others. To gain experience with a particular application, we are building and plan to deploy "ambush" a request-driven location service. By building and deploying a useful yet dangerous application like ambush, we hope to develop an understanding of how applications interact with social norms.

Notes

1. A number of other technologies including GPS have this same advantage that location is computed locally.

References

[1] P. Bahl and V. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. INFOCOM (2), p. 775–784.

[2] L. Barkhuus and A. Dey. Location-based services for mobile-telephony: a study of users' privacy concerns. In *Proceedings of INTERACT 2003, 9th IFIP TC13 International Conference on Human-Computer Interaction,* 2003.

[3] V. Bellotti and A. Sellen. Design for Privacy in Ubiquitous Computing Environments. In *Proceedings of The Third European Conference On Computer Supported Cooperative Work (ECSCW '93),* Milan, Italy, 1993. Kluwer Academic Publishers.

[4] P. Castro, P. Chiu, T. Kremeneck, and R. Muntz. A Probabilistic Room Location Service For Wireless Neworked Environments. In *Proceedings of Ubicomp,* Atlanta GA, 2001.

[5] S. Consolvo and M. Walker. Using the Experience Sample Method to Evaluate Ubicomp Applications. *IEEE Pervasive Computing Mobile and Ubiquitous Systems: The Human Experience,* 2(2):24–31, Apr-June 2003.

[6] K. Crowston and M. Williams. Reproduced and Emergent Genres of Communication on the World Wide Web. *The Information Society,* 16:201–205.

[7] P. Dourish. What We Talk About When We Talk About Context. *Personal and Ubiquitous Computing,* 8(1).

[8] R. Grinter and M. Eldridge. Y do tngrs luv 2 txt msg? In *Proceedings of the European Conference On Computer Supported Collaborative Work (ECSCW 2001),* Bonn, Germany, 2001.

[9] R. Grinter and L. Palen. Instant Messaging In Teen Life. In *Proceedings of the ACM Conference On Computer Supported Cooperative Work (CSCW 2002),* New Orleans, LA, 2002.

[10] M. Gruteser and D. Grunwald. A methodological assessment of location privacy risks in wireless hotspot networks. In *Proceedings of the First International Conference on Security in Pervasive Computing,* 2003.

[11] R. Harper. Looking at Ourselves: An Examination of the Social Organization of Two Research Laboratories. In *Proceedings of the ACM Conference Computer-Supported Cooperative Work,* Toronto, Canada, 1992.

[12] J. Hightower and G. Boriello. A Survey and Taxonomy of Location Sensing Systems for Ubiquitous Computing. Technical Report CSE 01-08-03, University of Washington, 2001.

[13] J. Hong, G. Boriello, J. Landay, D. MacDonald, B. Schilit, and J. Tygar. Privacy and Security in the Location-enhanced World Wide Web. In *Workshop on Ubicomp Communities: Privacy as Boundary Negotiation (Ubicomp 2003),* Seattle, WA, 2003.

[14] M. Ito and O. Daisuke. Mobile Phones, Japanese Youth, and the Re-Placement of Social Contact. Front Stage-Back Stage: Mobile Communication and the Renegotiation of the Social Sphere. Grimstad, Norway, 2003.

[15] E Kaasinen. User Needs for Location-aware Mobile Services. *Personal and Ubiquitous Computing,* 7(1):70–79, 2003.

[16] A. Ladd, K. Bekris, A. Rudys, L. Kavraki, D. Wallach, and G. Marceau. Robotics-based Location Sensing Using Wireless Ethernet. In *Proceedings of MOBICOM,* pages 227–238, 2002.

[17] S. Lederer, J. Mankoff, and A. Dey. Who Wants to Know What When? Privacy Preference Determinants in Ubiquitous Computing. In *Proceedings Of Extended Abstracts of CHI 2003, ACM Conference on Human Factors In Computing Systems,* pages 724–725, Fort Lauderdale, FL, 2003.

[18] E. Mynatt and J. Tullio. Inferring calendar event attendance. In *Proceedings of the ACM Conference on Intelligent User Interfaces (IUI 2001),* pages 121–128, Santa Fe, New Mexico, 2001.

[19] B. Nardi, D. Schiano, M. Gumbrecht, and L. Swartz. I'm Blogging This: A Closer Look at Why People Blog. *Communications of the ACM,* to appear.

[20] L. Palen and P. Dourish. Unpacking "privacy" for a networked world. In *Proceedings of ACM Coneference on Human Factors In Computing Systems.*

[21] R. Putnam. *Bowling Alone.* Simon and Schuster, New York, 2000.

[22] M. Rabin. Psychology and Economics. *Journal Of Economic Literature,* 36:11–46.

[23] A. Ranganathan, R. Campbell, A. Ravi, and A. Mahajan. ConChat: A Context-Aware Chat Program. *IEEE Pervasive Computing,* pages 52–58, July–Sept 2002.

[24] E. Schegloff. Notes on Conversational Practice: Formulating Place. In D. Sudnow, editor, *Studies in Social Interaction,* pages 75–119. New York, 1972.

[25] B. Schilit, J. Hong, and M. Gruteser. Wireless location privacy protection. *IEEE Computer,* 36(12):135–137, 2003.

[26] B. Schilit, A. LaMarca, G. Boriello, W. Griswold, D. MacDonald, E. Lazowska, A. Balachandran, J. Hong, and V. Iverson. Ubiquitous Location-Aware Computing and the Place Lab Initiative. In *First ACM International Workshop on Wireless Mobile Applications and Services on WLAN Hotspots (WMASH),* San Diego, CA, 2003.

[27] M. Sprietzer and M. Theimer. Providing Location Information In A Ubiquitous Computing Environment. Panel session in Proceedings Of Fourteenth ACM Symposium on Operating Systems Principles. Asheville, NC, 1993.

[28] A. Taylor and R. Harper. The Gift of the Gab: A Design Oriented Sociology Of Young People's Use Of Mobiles. *Journal of Computer Supported Cooperative Work (CSCW),* 12(3):267–296.

This page intentionally left blank

Author Index

This page intentionally left blank

Topic Index