

Towards Context Adaptive Privacy Decisions in Ubiquitous Computing

Florian Schaub*, Bastian Könings*, Michael Weber*, Frank Kargl†

*Institute of Media Informatics, Ulm University, Germany

Email: { florian.schaub | bastian.koenings | michael.weber }@uni-ulm.de

†DIES Group, University of Twente, The Netherlands; Institute of Distributed Systems, Ulm University, Germany

Email: f.kargl@utwente.nl

Abstract—In ubiquitous systems control of privacy settings will be increasingly difficult due to the pervasive nature of sensing and communication capabilities. We identify challenges for privacy decisions in ubiquitous systems and propose a system for *in situ* privacy decision support. When context changes occur, the system adapts a user's privacy preferences to the new situation. As a consequence, recommendations can be offered to the user or sharing behavior can be automatically adjusted to help the user maintain a desired level of privacy. The system learns from user interaction and behavior to improve decision accuracy. In this paper, we outline the main components of our system and illustrate its operation with an ambient assisted living use case.

I. INTRODUCTION

Future ubiquitous or pervasive computing environments will be characterized by the integration of sensing, processing, and communication capabilities in the physical environment. Everyday objects will become *smart* [1], [2] in the sense that they can perceive their environment and communicate and interact with each other. The main goal is to facilitate more natural and less obtrusive interaction between users and computing systems in order to better support users in their activities. Context awareness and implicit interaction are seen as basic ingredients towards realizing this vision.

But with the promises of ubiquitous computing also arise novel privacy issues as pointed out early on by Weiser [3]. Almost invisible sensing capabilities ubiquitously embedded in the environment, paired with increasing storage and processing capabilities, also enable large scale surveillance of users [4]. Research on privacy in ubiquitous computing has been mainly focused on how sensitive information can be protected and shared in a privacy preserving manner [5], [6]. As location awareness is one of the key characteristics in ubiquitous computing, location privacy became a dominant topic [7], [8]. One common approach is the obfuscation of information to reduce its quality. Other approaches aim to hide the information owner via anonymization or the use of pseudonyms. Iachello and Hong [9] give a good overview on privacy controls for information sharing, e.g., based on privacy policies. Approaches based on spheres [6] or territorial privacy [10] employ physical metaphors to reduce complexity of privacy controls in ubiquitous systems.

While these approaches offer means for controlling privacy, their configuration is commonly left to the user. But users will have a hard time defining privacy settings that match their

actual privacy preferences due to the complexity of systems, the multitude of entities (technical and human), and changing context. Especially manually pre-defining what information should be available to which entity (human or technical) in any given situation will be infeasible in such scenarios. The abstract nature of privacy implications makes it difficult to presuppose the desired level of privacy for a specific situation [11]. Instead, privacy decision making should be supported *in situ*, i.e., in the situation where a privacy decision is required [12]. The current context needs to be considered in order to help users effectively control their desired level of privacy in any given situation.

Existing approaches for supporting *in situ* privacy decision making either focus on enhancing awareness of information flow or on controlling specific information items, such as location. Most systems [13], [14] only become active when triggered by external requests for the user's information. However, we argue that passive observations dominate in ubiquitous computing scenarios, i.e. entities sense, process, and communicate information about users without explicitly requesting access or interacting with them [10]. Therefore, not only interactors but also other physically or virtually present entities must be taken into account for privacy decisions as well as privacy controls. We propose to utilize context-awareness to dynamically adapt privacy preferences in a privacy decision model to help a user maintain a desired level of privacy, either by providing recommendations to the user or via automatic reconfiguration. By adapting to the user, the system can learn the user's preferred trade-off between user involvement and automated enforcement.

In this paper, we first discuss privacy decision issues in ubiquitous systems (Sec. II). Then, we outline our system and privacy decision process (Sec. III) and exemplify it with a use case (Sec. IV). We conclude the paper with an outlook on challenges and future work (Sec. V).

II. PRIVACY DECISIONS IN UBIQUITOUS COMPUTING

Making adequate privacy decisions in information systems is already difficult today. For example, when sharing information on social networking sites [15]. The characteristics of ubiquitous systems further increase the difficulty of specifying privacy settings that match the user's actual privacy preferences. The following challenges can be identified:

Information explosion: The amount of information about a user grows with the number of sensors. Information exists on multiple levels of abstraction with varying semantic richness. Privacy decisions for sharing this information must be made on relevant abstraction levels to be comprehensible to users. For example, privacy decisions for indoor location sharing should rather be based on rooms than on specific sensors.

Context explosion: With an increasing number of *smart* entities and environments, situations in which privacy decisions are required explode due to exponential growth in potential context configurations. Taking all context situations into account a priori is infeasible and would prevent efficient privacy decisions. Instead, privacy decisions must be adapted to previously unknown situations.

Physical boundaries dissolve: The integration of sensors and communication capabilities into the environment enables virtual entities to participate in physical environments. Users may not know that they are being observed or by whom. Thus, sharing decisions must consider physically present entities as well as virtual and remote entities, which may be unknown. Users must be made aware of the extent of their physical-virtual mixed environment.

Observations and disturbances: Physically and virtually present entities may observe users without actively interacting with them. Furthermore, not only exchanged information may be privacy sensitive. A user's activities, interaction partners, external state, posture, and behavior may also need to be protected. Therefore, privacy decisions must not only pertain to sensitive information but also observations of the user. Smart entities may also cause disturbances in a user's physical environment, e.g., via audiovisual output, automation, or robots. Such disturbances can impact a user's privacy and must be included in privacy decisions in ubiquitous systems.

Abstract privacy implications: Due to the complexity introduced by the characteristics above, implications of privacy decisions become too abstract to be estimated properly by the user in advance. Therefore, a continuous re-evaluation process is required to support the user [12].

III. IN SITU PRIVACY DECISION SUPPORT

We propose a system for supporting a user's privacy decisions in situ, i.e., in the context they are required in, following the notion of contextual integrity [11]. Instead of requiring static definition of privacy settings beforehand, our system approximates the user's privacy preferences and adapts them to the current context. The system can then either recommend sharing decisions and actions or autonomously reconfigure privacy settings. The goal is to support users in maintaining a level of privacy that adequately fits their privacy needs for the current activity in the current context. This means that privacy settings should neither be too tight nor too open [12].

In order to provide adequate decision support and properly adapt privacy preferences, we take into account the issues named in the previous section. Especially territorial privacy aspects, such as physical-virtual mixed environments, observations and disturbances, factor into privacy decisions, besides

information privacy aspects. Furthermore, the system adapts to the specific user by learning from explicit sharing decisions, implicit user behavior, and reactions to system actions. The personalization to a specific user has the advantage of better emulating that user's privacy decision process. It also helps to decide when to involve the user in the decision process by providing recommendations only and when privacy decisions can be realized autonomously.

We assume that the system is implemented as a personal trusted agent and supports privacy decisions of a single user. The system's main components are the *context model*, the *privacy decision engine*, and *realization and enforcement* of adapted privacy preferences. In the following, we will discuss them in more detail together with the privacy decision process. Figure 1 provides an overview.

A. Context Model

To facilitate privacy decisions on an appropriate abstraction level, we distinguish between *decision level* and *system level* (see Fig. 1). The system level handles context acquisition and provides semantically enriched information to the decision level. Thus, the system level enables context awareness but also filters context information and maps it to semantic concepts required for decisions. Semantic mappings can be derived from a pre-defined or learnt world model. On the decision level, the context model only contains components relevant for privacy decision making. The main components on the decision level are the *user* that performs an *activity* in an *environment*.

The user has *resources*, which can be information items but also devices or sensors, and an observable *state*, i.e., the user's posture or bodily expression. An *activity* is always user-centric and has some abstract *goal*. An activity involves the user, user resources, and *interactors*, i.e., entities the user engages with.

The environment spans the physical and virtual realm and is user-centric in the sense that it is defined by the user's physical location. It is assigned a *type*, i.e., a semantic label, such as home or work, based on system level input. The environment consists of physical and virtual *entities* and *staging*. Humans, software agents, services, and smart devices are all represented as *entities* with the ability to perceive, process and communicate information. The staging defines an environment's physical and virtual configuration. For example, the position of walls, windows, and screens, and how virtual entities are connected to the physical environment. While the staging shares system level characteristics, it can influence privacy decisions (e.g., an unfamiliar office) and is therefore represented on the decision level.

Entities can have different roles. In our previously defined territorial privacy model [16], an *observer* is an entity that can perceive the user, either directly or via other observers forwarding information. A *disturber* is an entity that has the ability to disturb a user's activity, e.g., a person walking into a room disrupting a conversation or a cleaning robot becoming active when the user watches television. When the user engages with an entity in an activity, the entity becomes an

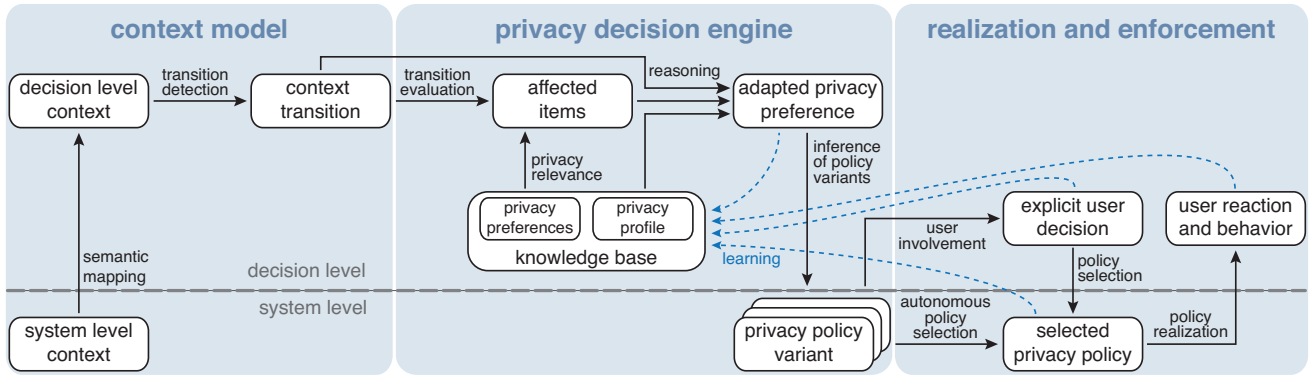


Fig. 1. Overview of the privacy decision process.

interactor. An entity can be observer, disturber, and interactor at the same time. Entities can be organized hierarchically to simplify privacy decisions. For example, sharing decisions can reference a person, while privacy settings will also apply to the person's devices, e.g., a mobile phone.

Privacy decisions are further facilitated by assigning trust to entities in the context model and ambient trust to the environment, i.e., determined by its type and staging. For this purpose, we are currently developing entity trust evaluation mechanisms based on social trust concepts [17].

B. Privacy Decision Engine

The context model allows to reason about which context items are affected by a context transition. A change in context C leads to a new context situation C' . The transition $T_{C \rightarrow C'}$ is defined as the difference between C and C' and captures all changes between them. When a transition occurs, the privacy decision engine (PDE) evaluates $T_{C \rightarrow C'}$ to determine which protection worthy context items are affected by $T_{C \rightarrow C'}$ (cf. Fig. 1). Protection worthy items can be the user's resources, but also the user's state, the activity, or its interactors. Protection worthiness (or privacy relevance) of context items for a given context are determined by the user's privacy preferences.

As a user's true privacy preferences exist only inside the user's mind, our system can only approximate them. Thus, the PDE matches privacy preferences stored in the knowledge base to the context transition $T_{C \rightarrow C'}$ and the new situation C' in order to infer an adapted privacy preference for C' in the reasoning step (see Fig. 1). We employ case-based reasoning [18] to avoid extensive a priori knowledge acquisition. Privacy preferences are stored as cases, which are retrieved by evaluating the similarity of previous situations with the current one. Similar cases are then adapted to the new context resulting in an adapted privacy preference that is retained as a new case. We plan to use personality-based privacy profiles to govern adaptation of privacy preferences. The main idea is to determine the user's personality type [19] before initial system use to select a privacy profile in order to speed up the bootstrapping process of learning the user's personal preferences. The privacy profile then serves as a basis for adapting privacy preferences and is subsequently

further adjusted to the user by learning from the user's explicit decisions, behavior, and reaction to system actions.

Based on the adapted privacy preference the PDE infers multiple privacy policy variants (see Fig. 1). While a privacy preference describes the user's privacy goal on the decision level, a privacy policy describes one way of achieving the privacy preference on the system level. Thus, C' may offer different alternatives, represented by privacy policies, for realizing the privacy preference. It may also be possible that the privacy preference cannot be realized in the current context. In that case, the privacy policy would suggest terminating the activity. For each privacy policy variant a confidence score is calculated based on how well it fits the adapted privacy preference. Based on the confidence scores, the PDE selects the most appropriate policy candidate or triggers user involvement if the confidence is below a certain threshold (see Fig. 1). The specific threshold is determined by the user's personality and previous privacy decisions.

The system learns from explicit user decisions, as well as reactions to the system's realization of privacy policies. These cues are used to adjust stored privacy preferences and tailor interaction thresholds to the user's expectations.

C. Realization and Enforcement

Next, the selected privacy policy must be realized on the system level (see Fig. 1). Our privacy policies combine territorial privacy and information privacy aspects. First, territorial privacy mechanisms [10], [16] are employed to prune the number of physical and virtual entities granted access to the user's private territory. The private territory is defined by a territorial privacy boundary that separates desired and undesired entities. The entities remaining inside the private territory have defined observation or disturbance channels to the user, or more specifically, to protectionworthy context items.

Next, our privacy policies define granularity adjustments for specific information items. For example, instead of the user's exact position only the street address or city can be provided. Similarly, the granularity of the user's identity can be adjusted, e.g., anonymous, pseudonymous, or full identity. Granularity adjustments can also be defined for other information types.

Depending on the environment, different strategies for policy realization and varying degrees of enforcement are possible [10]. In personal and shared personal environments, such as the home, system components are under the user's control allowing trust assumptions in terms of policy realization. In shared and public environments, the user has generally less control. Yet, trusted computing or collaborative mechanisms can support enforcement of privacy policies.

IV. USE CASE: AMBIENT ASSISTED LIVING

In the following we discuss a use case to illustrate the proposed system. Alice is an elderly person living alone in her home equipped with ambient assisted living technology. One day, Alice falls on her way to the bathroom and remains unconscious. The monitoring system (MS) detects Alice's fall and wants to inform Dan, her doctor, with a warning message including her vital signs.

Our system detects multiple context transitions. Before the fall (C_0), Alice's activity is *walking to bathroom* and the MS is active but can only process information locally. In transition $T_{C_0 \rightarrow C_1}$, Alice's fall is detected on the system level. Her state on the decision level changes to *unresponsive*. The PDE analyzes $T_{C_0 \rightarrow C_1}$ and determines that no privacy adaptation is required. Next, MS initiates a new activity on behalf of Alice in $T_{C_1 \rightarrow C_2}$. This activity involves entities Alice, MS, and Dan and requires access to Alice's vital sign sensors. The PDE matches this activity to the *emergency help* activity which allows sharing of vital signs with remote medical personnel. In the current context, the PDE derives a privacy policy that allows the MS to pass information on to Dan. The selected privacy policy also includes Dan's warning system, which the message is sent to to reach Dan. Note that this system is not part of the adapted privacy preference, but is required on the system level to realize the adapted privacy preference.

Dan's warning system receives the emergency message and informs Dan. Dan immediately drives to Alice and arrives at her front door shortly after. Dan's arrival triggers another context transition $T_{C_2 \rightarrow C_3}$. A new entity is added in C_3 as a disturber. The system authenticates the entity as Dan. The PDE matches Alice's state *unresponsive* and Dan's attribute *doctor* with Alice's privacy preference to receive medical help in emergencies. The derived privacy policy states that Dan can enter the house. The policy is realized by the house automation system that automatically opens the door for Dan. Dan enters the house and can help Alice. Note, that if Alice would have been conscious, the PDE could have involved her in the privacy decision by asking her if she wants to let Dan in. The adapted privacy preferences would be stored and validated later on when Alice is well again.

V. OUTLOOK

In this paper, we proposed a system for context adaptive privacy decision support in ubiquitous computing. Our system takes both physical and virtual entities into account and extends privacy decisions not only to information sharing but also to observations and disturbances. The system's privacy

decisions can be used to support the user in her privacy decision making or to autonomously reconfigure the system. The goal is to maintain a privacy level that aligns with the user's privacy preferences and activity in the current context, and is neither too open nor too restrictive. While focusing on a single user enables improved personalization, it also introduces challenges, such as handling shared resources, e.g., devices owned or operated by multiple users.

The development of the system is currently work in progress. Next, we plan to further refine the context model and validate its generality and expressiveness by applying it to a set of versatile use cases. At the same time, we are working on the integration of trust aspects in the context model. For the privacy decision engine, we are currently working on suitable context and knowledge representations to enable efficient reasoning and inference of preferences and policies on different abstraction levels. We plan to implement the privacy decision engine in a prototype system to evaluate the accuracy of privacy decisions in user studies.

REFERENCES

- [1] M. Weiser, "The computer for the 21st century," *Scientific American*, vol. 265, no. 3, pp. 94–104, 1991.
- [2] J. Bohn, V. Coroamă, M. Langheinrich, F. Mattern, and M. Rohs, "Living in a world of smart everyday objects – social, economic, and ethical implications," *Human and Ecological Risk Assessment*, vol. 10, no. 5, p. 763–785, 2004.
- [3] M. Weiser, "Some computer science issues in ubiquitous computing," *Communications of the ACM*, vol. 36, no. 7, pp. 75–84, 1993.
- [4] M. Langheinrich, "Privacy by Design - Principles of Privacy-Aware Ubiquitous Systems," in *UbiComp'01*. Springer, 2001, pp. 273–291.
- [5] —, "A privacy awareness system for ubiquitous computing environments," in *UbiComp'02*. Springer, 2002, pp. 315–320.
- [6] J. I. Hong and J. A. Landay, "An architecture for privacy-sensitive ubiquitous computing," in *MobiSys'04*. ACM, 2004, pp. 177–189.
- [7] A. Beresford and F. Stajano, "Location privacy in pervasive computing," *IEEE Pervasive Computing*, vol. 2, no. 1, pp. 46–55, 2003.
- [8] J. Krumm, "A survey of computational location privacy," *Personal and Ubiquitous Computing*, vol. 13, no. 6, pp. 391–399, 2008.
- [9] G. Iachello and J. Hong, "End-user privacy in human-computer interaction," *Foundations and Trends in Human-Computer Interaction*, vol. 1, no. 1, pp. 1–137, 2007.
- [10] B. Könings and F. Schaub, "Territorial Privacy in Ubiquitous Computing," in *WONS'11*. IEEE, 2011, pp. 104–108.
- [11] H. Nissenbaum, *Privacy in Context - Technology, Policy, and the Integrity of Social Life*. Stanford University Press, 2009.
- [12] L. Palen and P. Dourish, "Unpacking "privacy" for a networked world," in *CHI '03*. ACM, 2003, pp. 129–136.
- [13] M. Wu, "Adaptive Privacy Management for Distributed Applications," Ph.D. Thesis, Lancaster University, 2007.
- [14] N. Sadeh, J. Hong, L. Cranor, I. Fette, P. Kelley, M. Prabaker, and J. Rao, "Understanding and capturing people's privacy policies in a mobile social networking application," *Personal and Ubiquitous Computing*, vol. 13, no. 6, pp. 401–412, Aug. 2009.
- [15] A. Acquisti and R. Gross, "Imagined communities: Awareness, information sharing, and privacy on the facebook," in *PET'06*. Springer, 2006, pp. 36–58.
- [16] B. Könings, F. Schaub, M. Weber, and F. Kargl, "Towards Territorial Privacy in Smart Environments," in *Intelligent Information Privacy Management Symposium*. AAAI, 2010.
- [17] C. Castelfranchi and R. Falcone, *Trust Theory: A Socio-Cognitive and Computational Model*, 1st ed. John Wiley & Sons, 2010.
- [18] A. Kofod-Petersen and A. Aamodt, "Contextualised ambient intelligence through case-based reasoning," in *ECCBR'06*. Springer, 2006.
- [19] M. Korzaan and N. Brooks, "Demystifying Personality and Privacy: An Empirical Investigation into Antecedents of Concerns for Information Privacy," *Journal of Behavioral Studies in Business*, pp. 1–17, 2009.