Privacy Context Model for Dynamic Privacy Adaptation in Ubiquitous Computing

Florian Schaub*, Bastian Könings*, Stefan Dietzel*†, Michael Weber*, and Frank Kargl*†

*Ulm University
Ulm, Germany
{ firstname.lastname }@uni-ulm.de

†University of Twente
Enschede, The Netherlands
{ s.dietzel, f.kargl }@utwente.nl

ABSTRACT

Ubiquitous computing is characterized by the merger of physical and virtual worlds as physical artifacts gain digital sensing, processing, and communication capabilities. Maintaining an appropriate level of privacy in the face of such complex and often highly dynamic systems is challenging. We argue that context awareness not only enables novel UbiComp applications but can also support dynamic regulation and configuration of privacy mechanisms. We propose a higher level context model that abstracts from low level details and contains only privacy relevant context features. Context changes in our model can trigger reconfiguration of privacy mechanisms or facilitate context-specific privacy recommendations to the user. Based on our model, we analyze potential privacy implications of context changes and discuss how these results could inform actual reconfiguration of privacy mechanisms.

Categories and Subject Descriptors

H.1.2 [**Models and Principles**]: User/Machine Systems

Keywords

Context awareness, personalization, pervasive computing, privacy, privacy impact assessment, ubiquitous computing

1. INTRODUCTION

Ubiquitous computing (UbiComp) envisions intelligent environments and smart objects that support users in their daily activities [29]. Central aspects are context awareness and implicit interaction, which can be realized by integrating sensing, processing, and communication capabilities into the physical environment, and modeling context on varying levels of abstraction [3, 9]. Besides the promises of UbiComp, novel privacy issues arise from embedded sensors paired with increasing storage, communication, and processing capabilities [30, 19]. Physical boundaries start to dissolve due to the fusion of virtual and physical environments through ubiquitous communication technologies. The amount of data that can be gathered and communicated to remote entities almost in realtime explodes [24]. The augmentation of the physical world with smart devices and corresponding virtual representations reintroduces a physical dimension to privacy not present in systems solely based on information exchange, e.g., online social networks. Users may not even be aware of sensors in their physical proximity that capture and relay information about them. While such *observations* may generate information, smart devices that act in the user's physical environment (e.g., a vacuuming robot [5]) can disturb the user's privacy expectations. Such actuators and autonomous devices are part of the UbiComp system and their potential *disturbances* need to be considered when studying privacy in UbiComp. We refer to such physical privacy aspects of UbiComp as *territorial privacy* [14]. Territorial privacy corresponds to a user-centric rather than information-centric understanding of privacy, in the traditional sense of having *"the right to be let alone"* [28]. To reflect these different aspects of UbiComp privacy, we distinguish between *observers*, which are entities that receive information gathered in the user's proximity, and *disturbers*, which are entities that can actively intervene in a user's environment [14].

Existing privacy research in UbiComp has been mainly focused on achieving awareness and control of information collection and processing. However, many approaches rely on pre-specified privacy policies or assume static or limited scenarios, and do not consider dynamic adaptation of privacy control strategies with respect to changing situations. Yet, privacy regulation is a dynamic and selective process [2]. Thus, privacy mechanisms should support dynamic privacy regulation [22] to ensure that privacy settings correspond to the user's privacy expectations. Context awareness can play a key role in this adaptation, as context changes can affect a user's privacy in a given situation. Knowing which context changes potentially affect privacy would facilitate dynamic adaptation of privacy mechanisms or enable tailored privacy decision support for users [24]. The question is, however, how relevant context changes and their privacy implications can be determined in UbiComp environments in the face of a large number of potentially unknown context situations.

In this work, we address this issue by deriving a generic privacy context model as a context abstraction for arbitrary scenarios. Our model facilitates identification of privacy relevant context changes and analysis of their potential privacy implications. The results allow to decide when to dynamically adapt privacy mechanisms and how.

UbiComp '12, September 5-8, 2012, Pittsburgh, USA.
ACM 978-1-4503-1224-0/12/09.

6th International Workshop on Context-Awareness for Self-Managing Systems (Casemans'12), UbiComp'12 workshops, Pittsburgh, September 2012
Author version.

2. RELATED WORK

Most existing UbiComp privacy approaches, such as pawS or Confab [20, 12], provide mechanisms to enhance privacy awareness and control of information collection and processing, but do not directly support users in the privacy configuration process. Privacy protection of collected context information has also been extensively studied, resulting in context model extensions to include ownership [6] and usage preferences [10]. Privacy of location information has especially attracted considerable research in this area [17]. Other work focuses on privacy sensitivity of context information [25] and privacy-preserving exchange thereof [11].

In contrast to context awareness and privacy of context information, how context awareness can support dynamic privacy management in general has received considerably less attention. Some approaches use context awareness to selectively control disclosure based on pre-specified context-dependent privacy policies. Jiang & Landay [13] propose information spaces with physical, social or activity-based borders to enforce privacy policies when accordingly tagged information crosses such borders. Moncrieff et al. [21] use a trained decision tree to match a user's location, social interactions, and critical activities in a smart home with privacy disclosure rules for care givers. In ATRACO, context changes trigger dynamic privacy policy evaluation [16]. Gong et al. [7] use active badges to enforce pre-configured spatial privacy policies while also offering an additional privacy now button. Formal access control models have also been extended for context awareness [1, 18] and Sigg [26] proposes to use context to improve security mechanisms. However, none of them support the dynamic configuration of privacy policies. In the Super-Ego framework [27], location requests from mobile applications are autonomously decided based on how often others have shared this location. Saleh et al. [23] enable users to set privacy preferences in and for specific context situations. Wu [31] suggests adaptive privacy management based on confronting users with privacy access requests and learning from decisions. Bünnig [4] describes an abstract disclosure decision model and argues that an appropriate context abstraction is required that matches a user's privacy preferences. Heiber & Marrón [8] model privacy influence of context but explicitly leave out aspects of the physical world. In contrast, we propose an abstract privacy context model to support dynamic privacy adaptation and reconfiguration as well as context-specific privacy recommendations for users.

3. PRIVACY CONTEXT MODEL

In the development of our privacy context model (PCM), we focus specifically on privacy-relevant context features in order to obtain a lean, yet expressive model. Our model can complement comprehensive context models for application adaptation by providing a high-level privacy abstraction. The question is what context features have privacy relevance and should therefore be represented in such a model?

In order to investigate this, we first look at an example for context-adaptive privacy: *location monitoring of field representatives*. Alice is a field representative for Bob's company. She uses a company car for client visits but may also use it personally. To keep an overview of his field representatives, Bob integrates a location sharing feature into company cars and phones. Bob's access to Alice's location should depend on her context. While driving to a client, Bob may access Alice's destination and her estimated arrival time but not Alice's exact position. When Alice is at a client (Charlie) Bob can know Alice's exact location. At the same time, all incoming calls for Alice should be blocked in order to not disturb the meeting. After work, Bob should not get any information about Alice's whereabouts.

Our example describes a selective location disclosure system, which also considers disturbances (blocking phone call during client meeting) to better motivate privacy relevance of context features in UbiComp systems.

The example shows that Alice has privacy sensitive context features, e.g., her location, and that granularity changes (preciseness of location) affect her privacy. However, privacy sensitive items only have actual privacy implications if they are accessed by another entity. We generalize persons, devices, and services in a UbiComp system as *entities*. Both Bob and Charlie observe Alice's location. Bob does so remotely via the location sensor and server, while Charlie physically sees Alice at his place without requiring access to a location sensor. However, Charlie could not observe Alice's location at other times. Thus, the *presence* (virtual or physical) of an observing entity makes a difference in the scope of the observation. We use *channels* to describe how entities and privacy sensitive items are connected as part of an observation. Alice takes a special role as we define privacy sensitivity of context items always for the current *user*. The user's *privacy sensitive items* (e.g., her location) are modeled explicitly. Only channels involving Alice's privacy sensitive items are represented in the PCM.

Disturbances are similar to observations but point in the other direction. Disturbances, such as calls to Alice's phone, originate from some entity and end at a physical entity in Alice's proximity, e.g., her phone's ringer.

To give current configurations of observations and disturbances a meaning, we have to take the user's *activity* into account as well. Activities allow us to distinguish situations, such as driving to Charlie and driving home, and corresponding privacy settings. The activity describes which privacy sensitive items and entities are essential for the user to perform the activity.

So far, we have identified three major privacy-relevant context aspects: the *user* and her privacy sensitive items, the user's *environment* containing virtual and physical entities and channels, and the user's *activities* in a specific situation. Next, we look at each aspect in more detail. Figure 1 shows two modeled situations from the example that we will use in the following for clarification.

3.1 User

The example shows that a user has two kinds of privacy sensitive items. Some items reveal information and others are targets for disturbances. Both kinds can be physical or virtual, and both dimensions need to be considered.

3.1.1 Information sources

An information source is the origin of some information about the user. In the example, Bob and Charlie observe Alice's location via different means (physical observation vs. relayed information from a location sensor). In terms of privacy, Alice's location is observable in both cases. Therefore, we model location as a single information source (see Fig. 1). How this information source is actually observed is part of

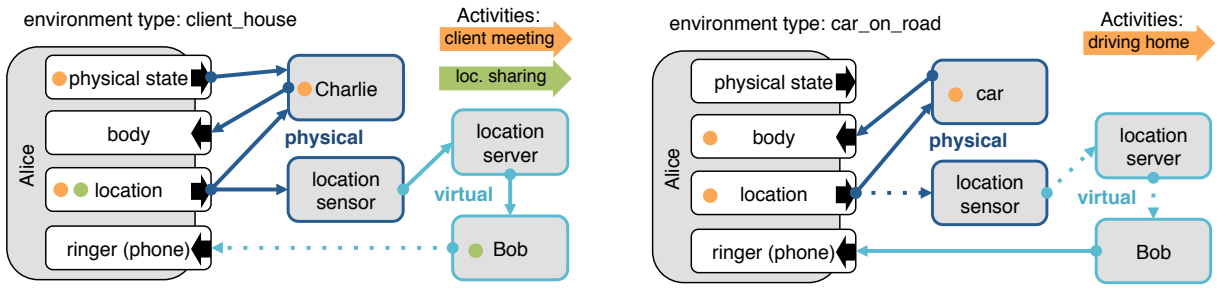

(a) User (Alice) is at client (Charlie).

(b) User (Alice) is driving home after work.

Figure 1: Example: context-aware monitoring of field representative.

the environment (see Sec. 3.2). Another example of an information source is the user's physical state, e.g., posture or facial expression. In our example, Alice's physical state is observable by Charlie but not Bob.

The previous information sources constitute time-variant data. For example, Charlie knows Alice's current location but not her future locations. Of course, dynamic information sources may facilitate prediction of future values to some extent. Static information sources, on the other hand, provide information that is statically associated with the user, such as Alice's name or phone number. Once disclosed, they are irrevocably known to the observer. We define a *type* to distinguish between *static* and *dynamic* information sources. Moreover, both types of information sources may provide information at a specific *granularity* in different situations.

3.1.2 Disturbance endpoints

We define a disturbance as any action of an entity occurring in the user's physical proximity. Hence, a disturbance must not have a negative connotation. The endpoint of a disturbance can be seen as a privacy sensitive item, as it potentially disrupts the user's tranquility. By default, the user's *body* and its senses are an endpoint for physical disturbances. For example, Charlie could touch Alice. In addition, we also consider device entities that are assumed to be controllable by the user as disturbance endpoints of the user. For example, the ringer of Alice's phone is one of her disturbance endpoints. By including such additional endpoints, we gain more expressiveness, because we can model changes of such endpoints and resulting privacy implications.

3.2 Environment

The environment is centered on the user's current location and represents the UbiComp system in which the user is acting. The environment models which entities surround the user and how they can observe or disturb privacy sensitive items. While typical environment models for context awareness can be very detailed, our example has shown that privacy relevant context features can be reduced to physical and virtual entities and their observation and disturbance channels [15]. We argue that this is sufficient to capture privacy relevant context changes, because any privacy relevant change of the physical environment would be reflected by changes of present entities and available channels. Similarly, the user's physical location does not need to be modeled explicitly as location changes cause changes in the environment model, i.e., the environment moves with the user. Uncertainty about unknown or undetected entities or channels can be modeled with respective placeholders. By assigning confidence values to each environment feature, we can model certainty about detected features and the risk of undetected ones.

3.2.1 Environment type

We assign a *type* to the environment. In our example, the environment types are *client house* (Fig. 1a) and *car on road* (Fig. 1b). The type gives the environment semantic meaning [21] to facilitate recognition by the user. It also supports differentiation of environments and allows to model how the user moves through them.

3.2.2 Entities

An entity has a *type* (person, device, or software service) that does not change [14]. The environment includes entities of different *presence*. For example, Charlie is a *physical* entity, while Bob is only a *virtual* entity (see Fig. 1).

An entity's *capabilities* (observe, disturb, both) determine how it can connect to privacy sensitive items of the user. As humans, Bob and Charlie can observe and disturb. The location sensor can only observe Alice's location.

3.2.3 Channels

Channels model how entities use their capabilities in the current context. We distinguish between two *types*: observation channels, which originate from an information source and end at an entity, and disturbance channels, which originate from an entity and end at a disturbance endpoint. Channels can consist of multiple *hops*, whereby each hop is a virtual or physical link between a *source* and *sink*. For example, Bob's observation channel of Alice's location consists of a physical hop from Alice's location information source to the location sensor (e.g., in her phone), and virtual hops from location sensor to location server and finally to Bob (see Fig. 1a). Thus, multi-hop channels define directed graphs between the user and multiple entities [15], which reflect a hierarchy of entities. Undesired channels can be removed according to Alice's privacy preferences (dotted lines in Fig. 1).

3.3 Activities

An activity is an abstract description of what the user is doing and marks privacy sensitive items and entities that are essential for performing the activity. However, activities and privacy preferences should not be confused. An activity describes what the user is doing in a situation while privacy preferences would describe how other entities would be allowed access to privacy sensitive items in that situation. A

user can engage in multiple activities at once. Each activity can consist of multiple actions.

3.3.1 Activity type

The *activity type* allows categorization of activities. In the example, Alice's activities *driving to client* and *driving home* allow to distinguish between two situations in an otherwise similar environment (her car). While activity detection is an active research topic, we assume its feasibility for our model. We assume further that an actvitiy's type is constant and does not change without the whole activity changing.

3.3.2 Actions

Actions mark the starting and endpoints of essential channels of an activity. For example in Fig. 1a, Bob's observation channel for Alice's location contains other entities and Bob could also relay the information further. Alice's activity *location sharing* has an action that marks Bob as the essential sink of location observation. Actions can be fulfilled with different privacy friendliness levels, e.g., encrypted or unencrypted connections, which are represented by different channels in the model.

An action consists of an *interactor*, a *privacy sensitive item*, and a *channel type* (observation or disturbance). When multiple interactors should access the same sensitive item, the activity would contain an action for each interactor.

4. CONTEXT CHANGES

A context change signifies an event with potential privacy implications. Thus, context changes mark occasions where privacy decisions and a re-adjustment of privacy settings may be required. We can simplify our analysis of privacy implications of context changes by representing the PCM as a directed graph $G = (V, E)$. V consists of disjoint sets of information sources (IS), disturbance endpoints (EP), and entities (EN): $V = IS \cup EP \cup EN$.

Observation (O) and disturbance (D) channels are directed edges, with $E = O \cup D$. IS vertices can only have outgoing edges, EP vertices can only have incoming edges. Entities EN can have incoming or outgoing edges according to their capabilities. Observations can be represented as paths $p_o : u_o \rightsquigarrow e$ ($u_o \in IS$, $e \in EN$), and disturbances can be represented as paths $p_d : e \rightsquigarrow u_d$ ($u_d \in EP$). Edge construction follows a set of rules to match the outlined model:

1. If $(v_i, v_j) \in O$, then v_j must be in EN and have capability *observe*.
2. If $(v_i, v_j) \in O$, then v_i must be in IS. Or v_i is in EN and has capability *observe* and a path $p_o : v_k \rightsquigarrow v_i$ exists in G from $v_k \in IS$ to v_i.
3. If $(v_i, v_j) \in D$, then v_i must be in EN with capability *disturb*.
4. If $(v_i, v_j) \in D$, then v_j must be in EP. Or v_j is in EN and has capability *disturb* and a path $p_d : v_j \rightsquigarrow v_k$ exists in G from v_j to $v_k \in EP$.

An activity is represented as a set of its actions: $A = (a_1, \ldots, a_n)$. Each action is a triple (u, e, c) consisting of a privacy sensitive user item ($u \in IS$ or $u \in EP$), an entity $e \in EN$ and a channel c exists between them (i.e., $p_o : u \rightsquigarrow e$ or $p_d : e \rightsquigarrow u$).

Based on this graph representation, all potential context changes can be characterized as ADD or REMOVE operations on edges and vertices. Changes of an information source, disturbance endpoint, or entity can be represented by combined ADD/REMOVE operations. For example, changing the granularity for location sharing entails removal of the current IS and simultaneous addition of a new IS with the same edges.

We assume that context changes can be serialized in order to analyze their atomic privacy implications. We further assume that new nodes are added before edges are established to them, and that a node is removed once all edges to it have been removed.

4.1 Privacy Implications

In order to assess privacy implications of context changes, we define the user's *exposure* in a context situation as an abstract measure for privacy. We distinguish between *potential exposure* defined by the cardinality of $IS \cup EP$ and *actual exposure* defined by the number of observations and disturbances $|p_o| + |p_d|$.

For the user, context changes are constituted by adding or removing information sources or disturbance endpoints.

ADD IS. Adding an information source increases potential exposure, as entities could establish observation channels to it. Static information sources pose a slightly higher privacy risk at this stage, because entities only need short access to the information source to learn its static content. Dynamic information sources are also critical but the privacy exposure posed by them increases over time when they are actually observed. Further, the granularity of an information source directly affects privacy.

ADD EP. Adding a disturbance endpoint increases exposure by increasing the risk of potential disturbances.

REMOVE IS/EP. The removal of a privacy sensitive item decreases potential exposure by the risk of observations or disturbances for the removed item. If observation or disturbance channels are connected with the removed item, they are removed as well. Thus, both potential and actual exposure are reduced.

Changes to IS and EP are intrinsic to the user, as they are explicitly or implicitly triggered by the user or her systems. Changing the set of privacy sensitive items directly affects the user's potential exposure. An adaptive privacy system could remove unnecessary privacy sensitive items in order to enhance privacy w.r.t. unknown entities in the environment.

Changes of *environment type* are also implicitly triggered by the user, because the environment type changes only when the user moves to a new environment. Thus, a change of environment type signifies a new situation that requires privacy re-evaluation, because the risk of unknown entities may have changed or different privacy preferences may apply to the new environment. Changes of the environment type are likely accompanied by entity and channel changes.

ADD EN. An added entity may establish observation or disturbance channels. However, addition of an entity facilitates evaluation of potential privacy implications before that entity actually establishes channels. An adaptive privacy system could evaluate that entity's possible channels and support or prevent channel establishment according to privacy preferences. The entity's capabilities determine if IS, EP or $IS \cup EP$ are at risk.

REMOVE EN. Removal of an entity also removes all channels associated with it. Further entities are removed if they are not connected to an IS or EP afterwards. Thus, removal

of an entity reduces actual exposure.

ADD O. In contrast to adding or removing privacy sensitive items, adding or removing channels effects actual rather than potential exposure. Addition of a channel signifies either a new one-hop (physical) channel or an extension of an existing path by adding a (virtual) hop. When a new observation channel is established, the connected *EN* (sink) learns the content of the user's *IS* (source). For static IS, the sink learns the provided information and may keep it indefinitely. For dynamic IS, the sink can gather information from IS as long as the channel persists. Especially physical observation channels have privacy implications as they require a physically present observer, which then may act as an interface between physical and virtual world. If undesired observation channels are created, the user's privacy system can try to exert control over channels or entities depending on the level of control in the current environment [14].

REMOVE O. Removal of an observation channel reduces actual exposure, similar to removal of entities. Note that potential exposure remains the same as long as the IS remains available. Monitoring the establishment and removal of observation channels allows analysis of how much information has been disclosed to connected entities through a channel while it existed.

ADD D. Adding a disturbance channel increases actual exposure by increasing the number of entities that are able to disturb the user.

REMOVE D. Removing a disturbance channel reduces actual exposure. Removing a disturbance channel also removes all disturbers attached to this channel if no other channel to the disturbance endpoint exists for them.

Entity and channel changes are extrinsic to the user. They represent the user's embeddedness in the context situation, i.e., how the user is connected with other entities. A change of the environment type should trigger general privacy reevaluation as the risk of unknown entities and channels can deviate between environments. Entity changes allow analysis of potential channels before they are established, while channel changes affect actual exposure.

ADD A. When an activity is added, multiple entities may be marked as interactors and the set of channels (of a specific type) between interactors and privacy sensitive items are marked as essential. The resulting privacy implication is that at least one channel must remain for each potential action to allow accomplishment of the activity. Thus, an adaptive privacy system should respect activities when managing the user's actual exposure. In addition, a new activity may be associated with privacy preferences, which a privacy system may try to enforce.

REMOVE A. Once an activity is removed, entities, channels, and privacy sensitive items may not be essential anymore and could be removed if corresponding privacy mechanisms exist.

In general, we identified intrinsic and extrinsic changes [12] with different privacy implications. Intrinsic changes affect the user's potential exposure. Extrinsic changes in the environment affect the user's actual exposure. Our generic analysis of implications shows that ADD operations have potentially negative impact on privacy, as more privacy sensitive items are exposed, or more channels and entities are introduced. REMOVE operations on the other hand have a positive impact on privacy as less privacy sensitive items are exposed and channels or entities are removed. Activities strike a balance between privacy and quality of service, because context components essential for the performance of the user's activities are specifically marked. While not considered in our current analysis, the PCM could also be used for analysis of long-term privacy implications by monitoring changes over time and applying pattern analysis.

4.2 Privacy Implications in the Example

Looking back at our example, we can analyze the privacy implications of the context changes between the two situations in Figure 1. When Alice leaves Charlie to drive home, no intrinsic changes occur. Information sources and disturbance endpoints remain the same. The environment type changes, indicating that the risk of unknown entities and channels may have changed. The entity *Charlie* is removed with corresponding channels. At the same time, a new physical entity *car* is added which can access the user's location (e.g., for navigation) and may disturb the user (e.g., with distance warnings or active safety systems). This entity and its channels are marked essential by the new *driving home* activity. As the *client meeting* activity has been removed, Alice can now be disturbed by phone calls. But Bob can no longer observe Alice's location, because the *location sharing* activity has been removed.

In the new situation (Fig. 1b), less observation channels to Alice's information sources exist, thus reducing Alice's actual exposure. However, the number of potential disturbances increases due to more disturbance channels. This small example shows that our model supports analysis of intricate privacy implications caused by context changes. Identified context changes identify *when* privacy adaptation is required and privacy implications can inform *how* privacy mechanisms could be adapted w.r.t. to privacy preferences or provide tailored user recommendations.

5. CONCLUSIONS & FUTURE WORK

We have shown how privacy-relevant context information can be abstracted to the user and her privacy sensitive items and entities in the environment that can observe or disturb such items. In addition, our privacy context model reflects the user's activities by marking items and entities that are part of an activity as essential. Due to the model's focus on privacy, only privacy relevant context changes are reflected while other more detailed changes without privacy relevance are filtered out. Our model has the potential to reduce the complexity of adaptive privacy systems by pre-filtering relevant privacy components and thus facilitating more concise context representations. Our analysis of generic privacy implications stemming from context changes shows that it is also feasible to determine the potential impact of even subtle context changes on the user's privacy. As a key contribution, our context model takes information privacy aspects as well as physical and territorial privacy aspects into account. It is therefore especially suited for ubiquitous systems in which virtual and physical realms start to merge. The generic nature of our model makes it applicable to diverse use cases and situations. Besides the example given here, we already successfully modeled privacy-relevant context in other applications, e.g., ambient assisted living, and plan to further assess the suitability of our model in future work.

A major challenge in the instantiation of the proposed model is the detection of physically and virtually present entities and channels. We are currently evaluating different

strategies and mechanisms for environments with varying levels of control or trust [14].

We are also working on the integration of trust in our model. Individual trust plays an important role for privacy decisions and is subject to similar dynamics. We plan to annotate environment components with trust in order to model trust changes as context changes thus facilitating analysis of privacy implications from trust changes in the same model.

Our model and the analysis of privacy implications can support either the autonomous reconfiguration of privacy mechanisms or properly support users in privacy decisions by providing context-aware privacy warnings and configuration recommendations. Thereby, a balance between user and system autonomy is important, in order to neither overwhelm users with decision requests nor alienate them with unexpected automated reactions. We are currently developing a privacy decision engine [24] that utilizes an implementation of the proposed privacy context model to support users in dynamic privacy management processes and decision making. Our system will adapt to individual users by learning their privacy preferences over time from explicit privacy decisions and implicit reactions to autonomous reconfiguration of privacy mechanisms. We plan to evaluate the effectiveness of our approach and the proposed privacy context model based on user studies in different scenarios.

6. REFERENCES

[1] A. Almutairi and F. Siewe. CA-UCON: a context-aware usage control model. In *CASEMANS'11*. ACM, 2011.

[2] I. Altman. *The Environment and Social Behavior: Privacy, Personal Space, Territory, and Crowding.* Brooks/Cole Publishing Company, CA, 1975.

[3] C. Bettini, O. Brdiczka, K. Henricksen, J. Indulska, D. Nicklas, A. Ranganathan, and D. Riboni. A survey of context modelling and reasoning techniques. *Pervasive and Mobile Computing*, 6(2):161–180, 2010.

[4] C. Bünnig. Smart Privacy Management in Ubiquitous Computing Environments. In *HCII'09*. Springer, 2009.

[5] T. Denning, C. Matuszek, K. Koscher, J. R. Smith, and T. Kohno. A spotlight on security and privacy risks with future household robots. In *Ubicomp'09*. ACM, 2009.

[6] A. Dey, G. Abowd, and D. Salber. A Conceptual Framework and a Toolkit for Supporting the Rapid Prototyping of Context-Aware Applications. *Human-Computer Interaction*, 16(2):97–166, 2001.

[7] N. Gong, M. Laibowitz, and J. Paradiso. Dynamic privacy management in pervasive sensor networks. In *Conf. Ambient Intelligence (AmI'10)*. Springer, 2010.

[8] T. Heiber and P. Marrón. Exploring the Relationship between Context and Privacy. In *Privacy, Security and Trust within the Context of Pervasive Computing*. Springer, 2005.

[9] K. Henricksen, J. Indulska, and A. Rakotonirainy. Modeling context information in pervasive computing systems. In *Pervasive'02*. Springer, 2002.

[10] K. Henricksen, R. Wishart, T. McFadden, and J. Indulska. Extending Context Models for Privacy in Pervasive Computing Environments. In *CoMoRea'05*. IEEE, 2005.

[11] C. Hesselman, H. Eertink, M. Wibbels, K. Sheikh, and A. Tokmakoff. Controlled Disclosure of Context Information across Ubiquitous Computing Domains. In *Intl. Conf. Sensor Networks, Ubiq., and Trustworthy Comp. (SUTC'08)*. IEEE, 2008.

[12] J. I. Hong and J. A. Landay. An architecture for privacy-sensitive ubiquitous computing. In *MobiSys'04*. ACM, 2004.

[13] X. Jiang and J. A. Landay. Modeling privacy control in context-aware systems. *IEEE Pervasive Computing*, 1(3), 2002.

[14] B. Könings and F. Schaub. Territorial Privacy in Ubiquitous Computing. In *WONS'11*. IEEE, 2011.

[15] B. Könings, F. Schaub, M. Weber, and F. Kargl. Towards Territorial Privacy in Smart Environments. In *Intelligent Info. Privacy Mgmt. Symp.* AAAI, 2010.

[16] B. Könings, B. Wiedersheim, and M. Weber. Privacy & Trust in Ambient Intelligence Environments. In *Next Generation Intelligent Environments: Ambient Assistive Systems*. Springer, 2011.

[17] J. Krumm. A survey of computational location privacy. *Pers. and Ubiquitous Comp.*, 13(6), 2009.

[18] D. Kulkarni and A. Tripathi. Context-aware role-based access control in pervasive computing systems. In *SACMAT'08*. ACM, 2008.

[19] M. Langheinrich. Privacy by Design - Principles of Privacy-Aware Ubiquitous Systems. In *UbiComp'01*. Springer, 2001.

[20] M. Langheinrich. A privacy awareness system for ubiquitous computing environments. In *UbiComp'02*. Springer, 2002.

[21] S. Moncrieff, S. Venkatesh, and G. West. Dynamic privacy assessment in a smart house environment using multimodal sensing. *ACM TOMCCAP*, 5(2), 2008.

[22] L. Palen and P. Dourish. Unpacking "privacy" for a networked world. In *CHI'03*. ACM, 2003.

[23] R. Saleh, D. Jutla, and P. Bodorik. Management of Users' Privacy Preferences in Context. In *Intl. Conf. Information Reuse and Integration*. IEEE, 2007.

[24] F. Schaub, B. Könings, M. Weber, and F. Kargl. Towards Context Adaptive Privacy Decisions in Ubiquitous Computing. In *PerCom'12 WiP*. IEEE, 2012.

[25] K. Sheikh, M. Wegdam, and M. V. Sinderen. Quality-of-Context and its use for Protecting Privacy in Context Aware Systems. *Journal of Software*, 3(3):83–93, 2008.

[26] S. Sigg. Context-based security: state of the art, open research topics and a case study. In *CASEMANS'11*. ACM, 2011.

[27] E. Toch. Super-Ego: a framework for privacy-sensitive bounded context-awareness. In *CASEMANS'11*. ACM, 2011.

[28] S. D. Warren and L. D. Brandeis. Right to privacy. *Harvard Law Review*, 4:193–220, 1890.

[29] M. Weiser. The computer for the 21st century. *Scientific American*, 265(3):94–104, 1991.

[30] M. Weiser. Some computer science issues in ubiquitous computing. *Comm. of the ACM*, 36(7):75–84, 1993.

[31] M. Wu. *Adaptive Privacy Management for Distributed Applications*. Ph.D. thesis, Lancaster University, 2007.