Classification of Privacy Management Techniques in Pervasive Computing

Stan Kurkovsky, Oscar Rivera, Jay Bhalodi

Central Connecticut State University
1615 Stanley Street, New Britain, CT 060505, USA
kurkovskysta@ccsu.edu

Abstract. Context awareness, the ability to adapt to the needs of each user, is a fundamental property of pervasive computing systems. Context information is created by tracking the actions and collecting real-time user data, such as location or body temperature. As the system collects more information about its users, the quality of its context-aware services increases, and so does a potential threat to the users' privacy if the context information is compromised. This paper surveys a wide range of existing pervasive computing systems that address this complex problem, and provides a generalized classification of privacy management techniques used by these systems. A comparative analysis of the surveyed systems and their privacy features is also presented.

Keywords: Privacy, security, management, pervasive computing.

1 Introduction

Since the earliest days of research in pervasive and ubiquitous computing, preservation of user privacy has been the subject of an active discussion [2,9,51]. On the one hand, a pervasive computing environment (PCE) needs to collect a substantial amount of information about its users to be able to adapt and respond to their demands without requiring much explicit user interaction. On the other hand, the more information a system has about its users, the greater is the potential threat to the users' privacy.

It is a core design feature of most PCEs to preserve the anonymity of the clients (users, services, or smart devices) interacting with service providers. However, an inherent contradiction lies in the fact that the service provider cannot fully trust the client that does not reveal its true identity. Without such a trust, the service provider cannot ascertain that the client has a rightful access to the requested resource or service. As a result, there is a perceived conflict between privacy and the technologies that enable the features of a PCE dealing with sensing and storing of user-related information. Personal privacy, in the aspects related to the information about the user's identity, past and current activities, and location, may be challenged because these kinds of information are typically among the data collected and stored by the sensors and services employed by most PCEs [30].

Large amounts of data potentially containing personal or personally identifiable information may become a serious temptation for misuse by a PCE service provider

or a third party intruder. Protection of such information as personal identity, location and past activities becomes more difficult due to the inherently distributed nature of PCEs and a continuous influx of newly collected data.

The main objective of this paper is to review and classify a broad range of privacy management techniques, methods and tools implemented by various existing and proposed pervasive computing systems and architectures. This paper is organized as follows. The background and motivation of this study is discussed in Section 2. Section 3 describes the identified privacy management techniques and illustrates them with the implementation details of one or more existing system. Section 4 concludes the paper with a summary.

2 Background and Motivation

Research literature provides a wide range of definitions of privacy in PCE: "control over information disclosure" [17], "privilege of users to determine for themselves when, how, and to what extent information about them is communicated to others" [29], and "ability of an individual to control the terms under which their personal information is acquired and used" [21]. A detailed discussion of many broad aspects of privacy in various contexts can be found in [5,9,18,38,39].

In PCEs, users may find an increasing amount of obstacles to maintain their privacy because of the growing volume of potentially identifiable data collected by the systems [1]. Data collected by perceptual interfaces that are capable of recognizing users, their gestures and facial expressions, may become a potential threat to the user privacy if it is somehow made accessible to a third party. Finally, some users may be fearful of compromised privacy and not be willing to accept and use a system that tracks and collects information about them, even though this data by itself may pose no privacy threats [1].

While there is a significant amount of research in pervasive computing aimed at design and implementation of privacy management tools and techniques, their practical usability and acceptance remains an important challenge [14]. In 2003, Computing Research Association identified the ability to "give computer end-users security they can understand and privacy they can control" in the "dynamic, pervasive computing environments of the future" as one of the four major research challenges of trustworthy computing [20].

Recently, a number of studies explored social implications of pervasive technologies and their impact on the changing perceptions of privacy [3,24,35]. As the technologies enabling PCE become ever more miniaturized and further blend with the surroundings, the users may not recognize or feel their presence and lose the awareness of the fact that some or all of their actions may be monitored and recorded by a PCE. Continuous collection of information about the users will inevitably expose the details of their personal traits: patterns of behavior, driving and walking routes, current location, shopping preferences, likes and dislikes, social associations, etc. Most studies suggest that some users may be willing to expose certain information about their behavior and actions in order to enable the adaptive nature of a PCE; at the same time, they may prefer to establish some boundaries and let other aspects of their

activities remain private. Users also may prefer to have the tools to limit the amount or granularity of the collected information, or to stop the monitoring altogether if they decide to do so [3]. Although the very notion of pervasive computing suggests that the technology should become invisible and blend into the environment, some users prefer to know that the technology actually is there [24].

Among many real world application domains where pervasive computing systems have been implemented, healthcare is one area in which privacy plays the most prominent role [23,44,53]. Hospitals, clinics and emergency rooms are a perfect testbed to implement a PCE: information services provide real-time data about patients, their medical history and treatment records; this information must be treated with confidentiality, yet available anytime to authorized users. Healthcare environment requires a high availability of information services, constant coordination among colleagues, and rapid response to emergencies [53]. A healthcare-oriented PCE must support a high degree of mobility for its collaborating users, who must have real-time access to the patient data protected by strong privacy safeguards. Depending on the specialization of a medical professional and their role in the treatment of a given patient, they may have access to different areas of records of that patient. In this application domain, privacy support may be further enforced by various legislative acts; in the US, protection of private medical records is mandated by the Health Insurance Portability and Accountability Act (HIPAA). Implantable medical devices used in a healthcare-oriented PCE present an even wider range of privacy-related challenges that are unique to this domain [23].

3 Privacy Management Techniques in PCE

In this section, we describe a number of privacy management techniques used in the design and implementation of existing pervasive computing systems and architectures. Due to space limitations, we omit a detailed discussion of each individual system; instead in each section describing a particular privacy management technique, we discuss the implementation of the corresponding technique by one or more system or PCE architecture.

3.1 Access Rights and Policy Management

It is often a challenging problem to control access to confidential information in a PCE [16,25,32]: users requesting access could lack knowledge about which access rights may be required, access control should enable easily adjustable and context-sensitive access rights, there may be a multitude of diverse information services making the management of access rights essential. Users may interact with many different smart devices and providers in order to obtain their services. Since these devices and services could be malicious or genuine, the user privacy may be compromised.

RAVE is a media space designed to support people who are distributed geographically, but work together on common tasks. Participants use video cameras, monitors, microphones and speakers, and can choose from a number of ways to

distribute their audio-video content or receiving the same from others in the media space [9,40]. The users of RAVE can control who can connect to them and what kind of connection is allowed. Feedback provided by this system alerts the users of the type of data being sent and who is allowed to access such data. In addition, one of the main benefits of RAVE is that it provides a privacy control at all levels of specificity, allowing users to make early decisions about granting permissions for specific kind of service to specific users.

The Trusted Platform Module (TPM) is a hardware solution for mobile computing devices that can provide common security architecture for pervasive environments [7]. TPM implements access rights without revealing user identities to external parties. A TPM-enabled system can authenticates a mobile user by issuing a challenge to the device, which is then signed by the access requestor. This technique provides robust security in the possible event of a third party stealing non-shareable information. The outside user will not able to authorize the operation signature and thus not able to enter a secure PCE.

The Privacy Awareness System (PawS) offers a set of tools that allows users protect their privacy while helping others respect it [37]. A user in PawS can specify the capabilities disclosed in the privacy policy such as who can update or delete the data versus who can only view the data. Thus, when a user enters any environment, a privacy proxy checks these policies against user's predefined privacy preferences. Services can collect information and users can utilize these services if the policy agrees, otherwise the user may not choose to use the service.

Privacy Violation Avoider (PriVA) is privacy-aware PCE model [4], aimed at avoiding information leaks while sharing resources and information among the users. PriVA has built-in generalized policies for sharing resources. It has default policy for certain resources and these policies cannot be changed to reduce the workload of the model. If a user does not want to share the resource they can tag for resource as non-sharable, rather than applying complex customized policies. Although PriVA provides a default policy for each resource, users might want to modify this policy. User can choose additional policies from the list of the customized policies provided in the model, which provides some policy flexibility. Different documents can have different policies, for example in teacher's handheld; the document containing grades have different level of importance versus the document with a list of reading materials.

Context-aware access policy management has been successfully implemented in the MOSQUITO project [36], which supports security requirements policies. Policy rules are expressed in XML and typically capture such context information as location and proximity. MOSQUITO's policies are aimed at evaluating trust between a consumer and a provider of a service in a PCE. The framework provides the flexibility to choose the implementation details of location-based trust evaluation policies, which can be customized to meet the application-specific requirements.

PerGym is a pervasive system for a gym that provides personalized services based on privacy-sensitive context information [46,47]. PerGym aggregates context information from distributed heterogeneous sources. Using a set of distributed policies, this information is aggregated by a context-aware privacy module, whose main responsibility is to enforce the system-wide compliance with a set of privacy policies set by each individual user.

Daidalos (Designing Advanced network Interfaces for the Delivery and Administration of Location independent, Optimized personal Services) is an infrastructure for enabling privacy in a pervasive environment [13,45,49]. Daidalos provides a Privacy Negotiation Manager, which is involved in privacy policy definition, structures and relations to personal information of the users. It also creates a negotiation protocol within the agent software where statements of privacy policy are exchanged between peers.

3.2 Classification of Resources

Mobile devices carried by the users of a PCE may contain a wide range of information, not all of which may be shareable. It could include personal names, phone numbers, addresses, appointment schedules, to do lists, bank accounts, as well as official or business information regarding customers, account holders, etc. Some elements of this information may be confidential and sensitive; others may be open for access. In an environment that supports resource sharing, unwanted parties could retrieve the confidential information causing information leakage and leading to the violation of user privacy. One method to assure the users that resource sharing will prevent such confidential information leakage is to provide a mechanism to classify all information resources by their degree of "share-ability."

PawS [37] provides the encoding of privacy policies into machine-readable XML to classify data. This structure indicates the users with whom an item of information may be shared.

PriVA [4] uses a resource manager that maintains a list of all resources available on each device. PriVA's TagR module tags these resources as shareable or non-sharable. This model provides the users the flexibility to share resources as they want, whenever they want. PriVA' ClsR module can be used to categorize all resources into different groups, such as images, videos, or text, adding to the simplicity of this model.

3.3 Data Persistence Control

In pervasive environments, information can be automatically collected and stored over a long period of time. Personal data is intrinsically valuable; it must be protected indefinitely against misuse and unauthorized disclosure. This task can be complicated by the fact that storage may be distributed among many systems and their components. Similarly to the ability to control the dissemination of information, the ability to control the persistence of data is also a fundamental property of privacy. Pervasive environments must provide tools to control the disclosure of collected and stored data. This can be done by placing time constraints on the data as defined by the owner or a service provider. A time to live attribute used in the Capability-Based Privacy-Preserving Scheme [34] or a restricted lifetime of access rights implemented in Hierarchical Identity-Based Encryption [27] could ensure that data does not exist beyond an established period of time. By limiting the data storage period, the risk of data leakage can be minimized.

PawS [37] uses a lifetime attribute for data that can be set by the owner, but it is the responsibility of a service provider or a user to observe that attribute. However, this framework stops short of providing the capability to virtually shred the data to satisfy the owner's privacy preferences. Such a feature can be the best tool to ensure confidentiality and privacy in pervasive environments since the capability to control information also implies the capability to extinguish private information. Therefore, PCE users should be in control of the lifespan of the information related to all private and personal activities.

3.4 Granularity Awareness

Information accessible in a PCE, such as time and location, can be grained down to details with different levels of privacy requirements. At the same time, there is a clear need to maintain the balance between the reasonable resolution of information the users are willing to provide and the amount of information needed by a particular service in a PCE. For example, some location-based systems may not function properly if they receive updates on a scale of kilometers and hours, as opposed to meters and seconds [10]. Similarly, privacy granularity restrictions used in SPARCLE [14] provide the ability to control database access down to a single field in the database, which may significantly hamper the functionality of the application on a large scale.

In Hierarchical Identity-Based Encryption [25,27], context information may be transformed along the way from the sender to receiver by decreasing its granularity level and therefore removing the details of information, for which the receiver may not have sufficient access rights. Additionally, information access requests may be evaluated and denied based on the granularity of the requested information. It is up to the owners of information to define the rules specifying the levels of granularity of the information they own and the associated levels of privacy.

The Publish/Subscribe Substrate model [43] is event-based; it uses three kinds of granularity levels to provide enhanced security and privacy guarantees to the users. System granularity assigns system events into categories that may or may not be received by third parties. Event type granularity determines whether subscribers of the system can receive events of a given type. Matching set granularity determines which set of interested and authorized subscribers can receive access to a particular set of events.

In PawS [37,38], users are allowed to adjust the granularity of their location data that is made available to the system; the resolution of the user location can be indicated in meters, or fractions thereof. PawS service providers can also adjust the level of location granularity they require by requesting only the necessary elements from the full set of location and/or information, e.g., only asking for the building name instead of the building and a room number.

3.5 Constraints and Permissions

Role Based Access Control (RBAC) [50] is one of the most widely used methods to control access to resources and services. In RBAC, users are associated with roles, which, in turn, are associated with a set of permissions. In pervasive systems, a straightforward implementation of RBAC may be difficult due to a large number of diverse service providers and transient users, which may lead to privacy violations and information leaks [26]. In a PCE, constraints on access rights often provide means to balance the tradeoff between the amount of privacy a user is willing to concede and the value of service that will be provided to the user. For example, users of an application may give control over choosing their location that will be available to others based on the identity of the person receiving the information. In order for PCE to be successful, owners of information should have a convenient way to stipulate the conditions under which their information can be accessed.

RAVE [9, 40] provides users with control over connections: which users can connect and what kind of connections are allowed. The users are made aware of what data is being sent via the feedback provided by the RAVE system. Therefore, RAVE makes its users aware of the technologies used in the PCE and their potential implications to their privacy.

Information Space Model [31] uses information boundaries and allows its users to define a set of permissions for access to information, resources, and services. The information space boundary acts as a trigger for privacy control to enforce permissions defined by the owners of each information space. A boundary, whether physical, social, or activity-based, delimits an information space. For example, a school principal can create an information space model that contains all the information in his office.

Context Models for Privacy [28] uses fact-based and situation-based approaches to manage the sets of constraints and permissions. Fact-based model provides a tool that application developers can use to explore and specify context requirements; it defines entities about which context information is required and types of information that interests the users. Developers may also choose to identify an appropriate source for each fact type using this fact-based model. Situation-based model provides a way to describe contexts at a higher level than the fact-based model. Situations are defined using predicate logic and can be combined easily with logical connectives to form rich descriptions. The situations help yield a truth value by expressing conditions on the context that can be evaluated against a set of variable bindings and a context to allow or deny access to shared information.

3.6 Ownership of Context Information

Features that allow or enable ownership of contextual information mark information in a manner that allows an association with an individual, behavior or action. This association gives the system the ability to protect data by implementing rules based on context, which can be used as an additional detail to be considered when evaluating and interpreting privacy rules. Caveats based on contextual information and its

ownership association with an individual user provides additional restrictions to data access rights.

Privacy Protecting Middleware [17] provides several mechanisms to control the conscious disclosure of contextual data. The rules for disclosure are negotiated and established at the time data is requested. Only the data that meets the privacy criteria set by the rules during the negotiation phase can be released. It is also capable of modifying the contextual data for a specific use or situation. These modifications remove privacy specific details that are not relevant to the nature of the data exchange. These mechanisms require a conscious awareness of contextual ownership in order to negotiate and publish data that is consistence with the privacy concerns of the owner. The state or fact of being owned must be a prerequisite to any viable negotiation between an entity that is requesting access to information and one that is safeguarding information.

RAVE [9, 40] provides a mechanism to establish the user context ownership and control. Users have complete control over the audio-video content in their personal space by controlling the state (on/off) and placement (location) of all audio-video equipment. In the RAVE virtual office space, workers choose to contribute information or opt out of the environment at any time. They can also choose who will receive their multimedia streams and what data to share. In multimedia spaces, context information is related to images in the video and sounds from the audio stream. Users have full authority over the contextual information in their personal office space but not in the community areas.

Participants in the PawS environment [37] can also assign ownership to context. PawS uses the W3C Platform for Privacy Preferences (P3P) to express ownership and control over user content [1, 48, 55]. P3P is both machine and human readable and can be used to mark content and communicate privacy policies. PawS users can also assign a privacy statement to context information using the P3P syntax. Although P3P can communicate user's privacy preference, it cannot enforce the preference: it relies on trust for enforcement.

Trusted Platform Module [7] can establish ownership by creating a non-migratable key to lock the information. This technique assumes that the user who holds the keys to unlock the information is also the owner of the information. TPM can ensure privacy if the distribution of keys to decrypt the data is controlled. Otherwise, an unauthorized party can use the keys to access the information.

Context Models for Privacy [28] addresses the challenges of assigning ownership to context information and enabling users to express privacy preferences for their own information. This approach offers a direct link between an information source and the entities that should be entitled to control the corresponding context information for privacy purposes. This ownership relationship describes the connection between an entity and the context attributes in which they have an interest in terms of privacy. With context models, ownership can be assigned to facts or objects. Ownership can also be assigned to one user or a group of users. Finally, ownership can be based on situation by applying rules associated with the fact or object types referenced by the situation.

3.7 Information Flow Obfuscation

Concealing the actual nature and extent of information flowing through a system is one of the techniques for protecting privacy in a PCE. All information flowing through the system should be intended for or restricted to the use of a particular entity that can be either the client or service provider. The disclosure of any information to a third party would be considered a violation of privacy. In a PCE, the role of the third party is not easily identified and guarding against information disclosure to an obscured entity presents many challenges. However, for the purposes of this work, the third party could be defined as any person, client, services provider or concealed observer who is not directly related to a service request. Strategies for obfuscating information flow can employ one or more of the following methods: concealing the existence of the services available from the service provider; concealing any information necessary to request services for the service provider; concealing any information concerning a service accessed by a client; concealing the exchange of contextual information between the client and the service provider; minimizing the amount of private information needed to fulfill the service request; encoding the contextual information such that it could only be understood by the parties involved in the service request; anonymizing the communication.

Mix networks and mix nodes [22, 33] provide a method that protects communication from third party observers, which eliminates the ability to trace the communication from the service provider to the client. This is achieved by routing all traffic through a set of specialized mix nodes that perform recoding of original messages and eliminate the possibility of establishing any linkage between incoming and outgoing messages. As a result, an anonymous communication network is established by obscuring the source and destination of the information flow.

Mist [2,16] facilitates anonymous communication for location-based applications. Mist-enabled systems separate the location information from the identity of the entity, thereby allowing the users to access the services while maintaining their privacy. The system uses a specialized encrypted communication protocol built on top of an overlay network of dedicated routers structured at a number of hierarchical levels. The level of information flow protection depends on the user's choice of the connecting router: routers at the higher levels of hierarchy provide a higher degree of privacy protection.

Obfuscation of the information flow in the Privacy Sensitive Information dilUting Mechanism (PSIUM) [19] is implemented using two methods. PSIUM reduces the usefulness of data that can be collected by third parties or the service provider by masking the actual nature of a service request. In PSIUM, the client queries the service provider for more services than is actually required. The service provider responds to all queries, but only one of the queries contains the desired service. The rest of the queries were designed to hide the true nature of the service request by including false information which reduces the usefulness of the data. PSIUM also protects the communication path by employing an anonymity enhancer which uses frequently changing pseudonyms to mask the client's identity and mixed zones to break the link between information source and destination.

3.8 Service Access Protection

A wide range of sensing technologies may provide context data used by a PCE, which may include location, time, user activities, etc. Despite the automated interaction with the context providers, users may be concerned about the leakage of private information, and thus they may want to know who can access the private data, as well as what kinds of data are being collected and how that data id being used. A number of PCEs provide different mechanisms to avoid the information disclosure about the services accessed by their users.

Privacy Protecting Middleware [17] provides a mechanism to avoid revealing any occurrences of access to context services by using cryptographic techniques such as a public- key encryption or secure communication protocols. Additionally, Privacy Protecting Middleware negotiates and adjusts the parameters of granularity and encryption for communication each session with the service consumer.

PSIUM [19] uses false data that appears to make sense, which makes it difficult for the service provider to distinguish the true data from the false data used in a request. To fulfill this requirement, PSIUM requires each client to have sufficient storage resources to maintain a database of genuine location coordinates and results that were used previously by the user. Thus, PSIUM prevents the misuse of the users' history of requests by a service provider while maintaining an acceptable quality of the service, but increasing the overall number of service requests.

Interaction History model [15] addresses the problem of service access protection by using authentication based on the history of user interactions with the system. After each interaction, credentials are created and are later used to prove the interactions. The credentials are encrypted and can be used only by their creator and other trusted members of the environment. In order to keep the privacy the creator of the credential cannot trace users and blind signatures can be used to ensure that signatures are not recognized. The history of user interactions is used in communication with the service provider or trusted partner without revealing the user identity.

The concept of mix zones together with frequently changing pseudonyms [10,11] can limit the possibility of network traffic analysis in order to disclose the access of PCE services. User privacy is enhanced by using a different pseudonym to access a different information service provider. Mix zones are used to further protect the network traffic from being analyzed. Mix zones are "connected spatial regions of maximum size in which none of these users has registered any application callback" [10]. All information service requests are forced to pass through one or more mix zone, which results in a disconnection between the user and any of their pseudonyms.

One of the main tasks of the Context-Aware Privacy Module employed by PerGym [46,47] is to transform service requests coming from each client and make them indistinguishable from those of other clients. PerGym anonymizes all requests by modifying the actual client identity information and by generalizing the context data included in the request. Generalization used by PerGym involves increasing the level of granularity of a particular value, e.g. a precise user location measured to a meter may be generalized into an interval or room in a building.

3.9 Information Disclosure Protection

In order to maintain nondisclosure of user information, service providers must only require the necessary data to perform a particular task, and the users should only need to provide that limited amount of information. The decision about what data elements are to be disclosed, in what format, and for what purposes, should be made in advance using protocols such as those employed in P3P [1,48,55] and PawS [37]. However, neither of these systems can support a perfect scenario for protecting against information disclosure. For example, P3P supports web interactions; therefore, the mechanisms are tightly coupled with web usage models and protocols. Also, the policy language contains constructs for expressing information for protecting information disclosed when involved in web browsing.

There are various techniques that can be followed to manipulate the personal and contextual data aimed at preventing information disclosure [17]: the requested information can be substituted with somewhat different information; requested information could be replaced with a set of data where the requested information in contained, or with another less specific data. Most of the systems surveyed in this work have information disclosure protection services built into their core foundation: both clients and service providers are designed only to disclose the information that is necessary for services request. Some of these systems deserve a special attention.

LocServ is a middleware service that facilitates the integration of location sensing technologies by location-based applications [42]. LocServ-enabled applications use multiple systems where users can specify a location query using the location models and help resolve queries using various location-tracking technologies. This sort of service requires mechanisms for controlled access to the user's location information allowing users to fine-tuned control over the amount of released information.

Privacy Protecting Middleware [17] provides mechanisms to control the disclosure of private information and enables users to reveal only the minimum amount of information required for a given service request. The SOA provides modules to publish, discover, compose, and access services. The middleware enforces access control, minimizes the amount of context information for each service request, and provides features for anonymization of data.

Geographic Location/Privacy service (Geopriv) collects the location information about its users by creating digitally signed location objects that encapsulate the user location data and their privacy requirements [54]. This model provides coupling of data when location information has been passed between multiple applications. In order to make sure the location server can make an informed decision it needs to possess rules indicating who has access to location items.

3.10. Protection of Information Usage

The ability to control the dissemination of information is a fundamental property of privacy. In a PCE, information can be collected and stored automatically for a variety of reasons. A service provider can collect personal information to customize the services presented to the user. Context data can be collected to enhance the user's experience with the system or to provide additional services. This collection of data

can be used to discover new facts, connections or relations previously unknown about the users and must be protected against misuse. If the data is saved for future processing, it must be protected from access by third parties and any unintended use by the service provider itself. When data is released, it must be released to the intended audience only. The collection and storage of data presents a challenge to privacy. If the user does not have control over the use or dissemination of their data, privacy cannot be ensured. For these reasons, PCE must provide enforceable restrictions on information usage by other users and service providers.

Privacy Protecting Middleware [17] monitors and regulates information exchanges between the user and the service provider in accordance with negotiated privacy rules and constraints. The middleware verifies data access according to identity-based access rights and available contextual information. It also filters the data according to privacy rules and contextual granularity constraints that are negotiable during the service request. It removes content from the data that is not required by the service request.

Godard, a component of RAVE [9, 40], provides tools to control the dissemination of audio and video information within the system and serves as a mediator for all service requests. Each service request indicates the intentions of the service initiator, which simplifies making a decision about allowing or denying the connection. Such an approach helps maintain a reasonable tradeoff between the user privacy and protection of access.

PawS [37] can be used to specify information usage restrictions. PawS includes collection and processing tools that enable service providers to communicate their information collection policies to the users. The announcement of the collection practices by the service provider enables the users to control the dissemination of private information by opting not to proceed with a data exchange. PawS also provides a mechanism to negotiate access to information in accordance to the user defined privacy policies. Although PawS can communicate information constraints between the service provider and the user, it cannot enforce those constraints. Instead, PawS must rely on the element of trust to protect against information misuse. Data is exchanged with the privacy policy attached and compliance is optional; therefore, PawS cannot fully protect information usage in pervasive environments where trust is not guaranteed.

The Quality of Privacy (QoP) [52] concept can be integrated into privacy aware architectures and can be used to protect information usage. QoP-enabled architectures provide a mechanism that allows both the user and the pervasive environment to use quantitative parameters to govern the level of privacy expected by the user. These parameters are based on five contextual variables: location, identity, access, activity and persistence. The parameters can be used as currency to determine the cost that the users might need to pay in regard to services offered by a pervasive environment and the privacy. The level of QoP demanded from the user will depend on contextual variables and the degree of privacy desired while using the pervasive application. On the other hand, the information that the user agrees to make available within the system determines the services the environment is willing to provide. In QoP, the context-aware privacy filter controls the information shared by the user with the pervasive environment according to a specific level agreed negotiated between the user and the service provider. A QoP-enabled system monitors the contextual

elements to determine the level of QoP to ensure that all parties adhere to the privacy agreement effectively enforcing the protection of information.

4 Summary

In this work, we surveyed a number of existing pervasive computing systems and architectures and examined a range of privacy management techniques used in their design and implementation. Table 1 summarizes this survey by indicating whether a particular technique is implemented be each system. It can be observed that while some of the techniques (e.g. access rights management or information disclosure protection) are prevalent in most of the systems, other critical features (e.g. data persistence control or granularity awareness) are implemented only by a handful of the systems.

Table 1. Summary of privacy management techniques in pervasive computing environments.

Technique / Architecture	Access rights management	Access policy management	Classification of resources	Data Persistence Control	Grananuality awareness	Constraints and permissions	Ownership of Context	Obscured information flow	Service access protection	Information disclosure protection	Protection of information usage
Capability-Based Privacy-Preserving Scheme	✓		✓	✓	✓	✓		✓	✓		
Context Models for Privacy	✓		✓		✓	✓	✓			✓	
Daidalos	✓	✓	✓						✓	✓	
Geopriv								✓	✓	✓	
Hierarchical Identity-Based Encryption	✓		✓	✓	✓	✓	✓	✓		✓	
Information Space Model					✓	✓		✓		✓	
Interaction History								✓	✓	✓	
LocServ								✓	✓	✓	
Mist								✓	✓	✓	✓
Mixed Networks								✓	✓		
MOSQUITO	✓	✓				✓				✓	
PawS	✓	✓	✓	✓	✓		✓	✓		✓	✓
PerGym/CARE	✓	✓			✓			✓			
PriVA	✓	✓	✓		✓	✓				✓	
Privacy Protecting Middleware	✓	✓					✓	✓	✓	✓	✓
Pseudonyms and Mix Zones								✓	✓	✓	
PSIUM	✓							✓	✓	✓	✓
Publish/Subscribe Substrate	✓	✓			✓	✓			✓		
Quality of Privacy (QoP)	✓					✓		✓			✓
RAVE	✓	✓				✓	✓				✓
SPARCLE	✓	✓			✓						
Trusted Platform Module	✓		✓				✓	✓	✓	✓	✓

The very concept of the PCE is based on anytime, anyplace availability of services, which implies that the system is always on, automatically collecting context data about its users wherever they may be. Completely disconnecting from all information services may not be desirable for many users. However, there is a clear need to for the users to be able to control their privacy in a flexible way: increase the granularity levels of spatial and temporal information that is collected about their activities [10,25,37,43], set the policies making them "invisible" to the system in certain areas or at certain times [4,9,36,37,45,47], or simply "disappear for a while" by feeding white lies into the system [6].

As pervasive technologies become more mature and become more widely adopted, the importance of preserving user privacy will continue increasing, which is clearly illustrated by pervasive computing systems in the healthcare application domain. This trend indicates that if pervasive computing systems are to succeed beyond the prototype stages and research labs, privacy management features should be given a very serious consideration.

5 References

1. M.S. Ackerman, "Privacy in Pervasive Environments: Next Generation Labeling Protocols," Personal and Ubiquitous Computing, Nov 2004, 8(6), pp. 430-439.
2. J. Al-Muhtadi, R. Campbell, A. Kapadia, M.D. Mickunas, Y. Seung, "Routing through the Mist: Privacy Preserving Communication in Ubiquitous Computing Environments," 22nd International Conference on Distributed Computing Systems, 2002, pp. 74-83.
3. D. Anthony, T. Henderson, D. Kotz, "Privacy in Location-Aware Computing Environments," IEEE Pervasive Computing, Oct-Dec, 2007, 6(4), pp. 64-72.
4. K.M. Asif, S.I. Ahamed, N. Talukder, "Avoiding Privacy Violation for Resource Sharing in Ad hoc Networks of Pervasive Computing Environment," 31st Annual International Computer Software and Applications Conference, 2007, pp. 269-274.
5. R. Babbitt, J. Wong, C. Chang, "Towards the Modeling of Personal Privacy in Ubiquitous Computing Environments," 31st Annual International Computer Software and Applications Conference, Jul 2007, pp. 695-699.
6. S.A. Bagüés, A. Zeidler, C.F. Valdivielso, I.R. Matias, "Disappearing for a while - using white lies in pervasive computing," ACM Workshop on Privacy in Electronic Society, 2007, pp. 80-83.
7. S. Balfe, S. Li, J. Zhou, "Pervasive Trusted Computing," 2nd International Workshop on Security, Privacy and Trust in Pervasive and Ubiquitous Computing, 2006, pp. 88-94.
8. R. Beckwith, "Designing for Ubiquity: The Perception of Privacy," IEEE Pervasive Computing, Apr-Jun 2003, 2(2), pp. 40-46.
9. V. Bellotti, A. Sellen, "Design for Privacy in Ubiquitous Computing Environments," 3rd European Conference on Computer Supported Cooperative Work, 1993, pp. 77-92.
10. A.R. Beresford, F. Stajano, "Location Privacy in Pervasive Computing," IEEE Pervasive Computing, Jan-Mar 2003, 2(1), pp. 46-55.
11. A.R. Beresford, F. Stajano, "Mix Zones: User Privacy in Location-aware Services," 2nd IEEE Annual Conference on Pervasive Computing and Communications, 2004, p. 127-131.
12. P. Bhaskar, S.I. Ahamed, "Privacy in Pervasive Computing and Open Issues," 2nd International Conference on Availability, Reliability and Security, 2007, pp. 147-154.
13. A.J. Blazic, K. Dolinar; J. Porekar, "Enabling Privacy in Pervasive Computing Using Fusion of Privacy Negotiation, Identity Management and Trust Management Techniques." 1st International Conference on the Digital Society, Jan 2007, pp. 30-35.

14. C. Brodie, C.-M. Karat, J. Karat, J. Feng, "Usable Security and Privacy: a Case Study of Developing Privacy Management Tools," Symposium on Usable Privacy and Security, 2005, pp. 35-43, 2005.
15. L. Bussard, Y. Roudier, R. Molva, "Untraceable Secret Credentials: Trust Establishment with Privacy," 2nd IEEE Annual Conference on Pervasive Computing and Communications, Mar 2004, pp. 122-126.
16. R. Campbell, J. Al-Muhtadi, P. Naldurg, G. Sampemane, M. D. Mickunas, "Towards Security and Privacy for Pervasive Computing", International Symposium on Software Security, 2002, pp. 77-82.
17. R.S. Cardoso, V. Issarny, "Architecting Pervasive Computing Systems for Privacy: A Survey," IEEE/IFIP Conference on Software Architecture, Jan 2007, pp. 26-29.
18. J. Cas, "Privacy in Pervasive Computing Environments – a Contradiction in Terms?" Technology and Society Magazine, Spring 2005, 24(1), pp. 24-33.
19. H.S. Cheng, D. Zhang, J.G. Tan, "Protection of Privacy in Pervasive Computing Environments," International Conference on Information Technology: Coding and Computing, Apr 2005, pp. 242-247.
20. Computing Research Association, "Four Grand Challenges in Trustworthy Computing," Nov 2003, http://www.cra.org/reports/trustworthy.computing.pdf
21. M.J. Culnan, "Protecting privacy online: is self-regulation working?" Journal of Public Policy Market, 2006, 19(1), pp. 20-26.
22. C. Diaz and B. Preneel, "Taxonomy of Mixes and Dummy Traffic," 18th IFIP World Computer Congress, Aug 2004, pp. 215-230.
23. D. Halperin, T.S. Heydt-Benjamin, K. Fu, T. Kohno, W.H. Maisel, "Security and Privacy for Implantable Medical Devices," IEEE Pervasive Computing, Jan/Mar 2008, 7(1), p. 30-39.
24. G.R. Hayes, E.S. Poole, G. Iachello, S.N. Patel, A. Grimes, G.D. Abowd, K.N. Truong, "Physical, Social, and Experiential Knowledge in Pervasive Computing Environments," IEEE Pervasive Computing, Oct-Dec 2007, 6(4), pp. 56-63.
25. U. Hengartner, P. Steenkiste, "Access Control to Information in Pervasive Computing Environments," 9th Conference on Hot Topics in Operating Systems, 2003, pp. 157-162.
26. U. Hengartner, P. Steenkiste, "Avoiding Privacy Violations Caused by Context-Sensitive Services," 4th IEEE International Conference on Pervasive Computing and Communications, 2006, pp. 222-233.
27. U. Hengartner, P. Steenkiste, "Exploiting Hierarchical Identity-Based Encryption for Access Control to Pervasive Computing Information," 1st International Conference on Security and Privacy for Emerging Areas in Communications Networks, Sep 2005, pp. 384-396.
28. K. Henricksen, R. Wishart, T. McFadden, J. Indulska, "Extending Context Models for Privacy in Pervasive Computing Environments", 3rd IEEE International Conference on Pervasive Computing and Communications, Mar 2005, pp.20-24.
29. J.I. Hong, J.A. Landay, "An Architecture for Privacy-Sensitive Ubiquitous Computing," 2nd international Conference on Mobile Systems, Applications, and Services, Jun 2004, pp. 177-189.
30. A.R. Jacobs, G.D. Abowd, "A Framework for Comparing Perspectives on Privacy and Pervasive Technologies," IEEE Pervasive Computing, Oct-Dec 2003, 2(4), pp. 78-84.
31. X. Jiang, J.A. Landay, "Modeling Privacy Control in Context-Aware Systems", IEEE Pervasive Computing, Jul-Sep 2002, 1(3), pp. 59-63.
32. L. Kagal, T. Finin, A. Joshi, "Trust-Based Security in Pervasive Computing Environments," Computer, Dec 2001, 34(12), pp. 154-157.
33. M. Klonowski, M. Kutyłowski, "Provable Anonymity for Networks of Mixes," 7th International Workshop on Information Hiding, Jun 2005, Lecture Notes in Computer Science, Vol. 3727, pp. 26-38.

34. D.M. Konidala, D.N. Duc, L. Dongman, K. Kwangjo, “A Capability-Based Privacy-Preserving Scheme for Pervasive Computing Environments”, 3rd IEEE International Conference on Pervasive Computing and Communications, Mar 2005, pp. 136-140.
35. V. Kostakos, E. O'Neill, A. Penn, “Designing Urban Pervasive Systems,” Computer, Sep 2006, 39(9), pp. 52-59.
36. S. Lachmund, T. Walter, L. Gomez, L Bussard, E. Olk, “Context-Aware Access Control Making Access Control Decisions Based on Context Information,” 3rd Annual International Conference on Mobile and Ubiquitous Systems, Jul 2006, pp. 1-8.
37. M. Langheinrich, “A Privacy Awareness System for Ubiquitous Computing Environments”, 4th International Conference on Ubiquitous Computing, Lecture Notes in Computer Science, Vol. 2498, pp. 237-245.
38. M. Langheinrich. “Personal Privacy in Ubiquitous Computing – Tools and System Support,” PhD thesis, ETH Zurich, Switzerland, May 2005.
39. S. Lederer, J.I. Hong, A.K. Dey, J.A. Landay, “Personal Privacy through Understanding and Action: Five Pitfalls for Designers,” Personal and Ubiquitous Computing, Nov 2004, 8(6), pp. 440-454.
40. W.E. Mackay, “Spontaneous Interaction in Virtual Multimedia Space: EuroPARC's RAVE system,” Imagina, 1992.
41. G.T Marx, “What's in a Name? Some Reflections on the Sociology of Anonymity,” The Information Society, 1 April 1999, Routledge, 15(2), pp. 99-112.
42. G. Myles, A. Friday, N. Davies, “Preserving Privacy in Environments with Location-Based Applications,” IEEE Pervasive Computing, Jan 2003, 2(1), pp. 56-64.
43. L. Opyrchal, A. Prakash, A. Agrawal, “Designing a Publish-Subscribe Substrate for Privacy/Security in Pervasive Environments,” 2006 ACS/IEEE International Conference on Pervasive Services, Jun 2006, pp. 313-316.
44. G. Pallapa, N. Roy, S Das, “Precision: Privacy Enhanced Context-Aware Information Fusion in Ubiquitous Healthcare,” 1st International Workshop on Software Engineering for Pervasive Computing Applications, Systems, and Environments, 2007, sepcase, p. 10.
45. E. Papadopoulou, S. McBurney, N. Taylor, M. H. Williams, K. Dolinar, M. Neubauer, “Using User Preferences to Enhance Privacy in Pervasive Systems,” 3rd International Conference on Systems, 2008, pp. 271-276.
46. L. Pareschi, D. Riboni, A. Agostini, C. Bettini. “Composition and Generalization of Context Data for Privacy Preservation,” 6th Annual IEEE International Conference on Pervasive Computing and Communications, 2008, pp. 429-433.
47. L. Pareschi, D. Riboni, C. Bettini, “Protecting Users’ Anonymity in Pervasive Computing Environments,” 6th Annual IEEE International Conference on Pervasive Computing and Communications, 2008, pp. 11-19.
48. J. Reagle, L.F. Cranor, “The Platform for Privacy Preferences,” Communications of the ACM, Feb 1999, 42(2), pp. 48-55.
49. I. Roussaki, M. Strimpakou, C. Pils, N. Kalatzis, M. Neubauer, C. Hauser, M. Anagnostou, “Privacy-Aware Modelling and Distribution of Context Information in Pervasive Service Provision,” 2006 ACS/IEEE International Conference on Pervasive Services, Jun 2006, pp. 150-160.
50. R.S. Sandhu, E.J. Coyne, H.L. Feinstein, C.E. Youman, “Role based access control models,” IEEE Computer, Feb 1996, 29(2): pp. 38-47.
51. M. Satyanarayanan, “Privacy: The Achilles Heel of Pervasive Computing?” IEEE Pervasive Computing, Jan-Mar 2003, 2(1), pp. 2-3.
52. M. Tentori, J. Favela, M.D. Rodriguez, V.M. Gonzalez, “Supporting Quality of Privacy (QoP) in Pervasive Computing,” 6th Mexican International Conference on Computer Science, Sep 2005, pp. 58-67.
53. M. Tentori, J. Favela, M.D. Rodríguez, “Privacy-Aware Autonomous Agents for Pervasive Healthcare,” IEEE Intelligent Systems, Nov/Dec 2006, 21(6), pp. 55-62.

54. H. Tschofenig, H. Schulzrinne, A. Newton, J. Peterson, A. Mankin, "The IETF Geopriv and Presence Architecture Focusing on Location Privacy," W3C Workshop on Languages for Privacy Policy Negotiation and Semantics-Driven Enforcement, Oct 2006.
55. T. Yu, N. Li, A.I. Antón, "A Formal Semantics for P3P", Conference on Computer and Communications Security, 2004, pp. 1-8, 2004.