AN INTEGRATED APPROACH TO USER-CENTERED PRIVACY FOR MOBILE INFORMATION SERVICES

Leonid Titkov, Stefan Poslad and Juan Jim Tan

Query Sheet

Q1 Au: ok?
Q2 Au: define?
Q3 Au: ok? sense
Q4 Au: city?
Q5 Au: spell out
Q6 Au: need full city.
Q7 Au: spell out name of this journal
Q8 Au: can this be table title
Q9 Au: need short table title.
Q10 Au: Spelling does not match with the reference.
Q11 Au: Supply high resolution graphic.
Q12 Au: Please cite figure 5.

3b2 Version Number : 7.51c/W (Jun 11 2001)
File path : {JournalProduction}Taylor&Francis/Uaai/v20n2-3/uaai148401/uaai148401.3d
Date and Time : 17/12/05 and 19:21

Applied Artificial Intelligence, 20:1–20

ISSN: 0883-9514 print/1087-6545 online
DOI: 10.1080/08839510500484181

Taylor & Francis
Taylor & Francis Group

AN INTEGRATED APPROACH TO USER-CENTERED PRIVACY FOR MOBILE INFORMATION SERVICES

Leonid Titkov, Stefan Poslad, and Juan Jim Tan □ *Department of Electronic Engineering, Queen Mary, University of London, London, United Kingdom*

□ *The evolution of powerful portable networked devices allows nomadic users seamless access to business and leisure information any time and any where. This drives service infrastructures to similarly evolve to meet new service management challenges in order to adapt information delivery to a richer spatial, temporal, and user context whilst balancing the concerns of the different stakeholders. This paper presents the use of a multi-agent system service infrastructure and associated privacy model to maintain the balance between the service provider's need to maximize its knowledge of the user context against the need for the user's privacy to be protected.*

Nomadic users can greatly benefit from location-awareness and personalization since these technologies offer a huge range of possibilities for accessing more relevant services that relate to their context. If the personal interests and current position of users are divulged then these can lead to an improvement of the user's enjoyment of mobile communication services, for example, a user could elect to be notified when friends are near and to meet at a proposed location. The use and reuse of the user's context has additional benefits for the user; it can reduce the information flow to the user. This is of particular benefit when delivering information to a low resource terminal, over a wireless link and to a mobile user whose attention may be constrained.

There are several user contexts that can be used. Personalization is based on the ability of the service to utilize the knowledge about a user's personal information. Location-aware services (LAS) can use a person's current location, destination, and route to deliver services. Terminal awareness orientates the information to be presentable on a range of terminals from mobile phones to high-end PDAs. Network link awareness takes into account the use of the QoS of the network link.

Address correspondence to Leonid Titkov, Department of Electronic Engineering, Queen Mary, University of London, Mile End Road, London E1 4NS, United Kingdom. E-mail: leonid.titkov@elec.qmul.ac.uk

Context awareness and in particular personalization and location-awareness, however, generates a range of privacy concerns (Kobsa 2002). If privacy issues within a context-aware service environment such as LAS are not properly addressed, users risk revealing their context, publicizing their personal details, and even compromising their safety at the current location, at the intended destination and on route. There is often a legal requirement in many countries to protect the privacy of the mobile users' information. If users perceive that the risks of using a technology outweigh their potential benefits, they may stop using that technology. The issue of keeping personalized, mobile, location-aware services (LAS) private has to be addressed in order to exploit their full market potential.

ANALYSIS OF PRIVACY REQUIREMENTS FOR CONTEXT-AWARE SERVICES

Example Context-Aware Service Interaction

In order to analyze the privacy requirements, we analyze a range of different services for mobile users based on those modeled in the EU CRUMPET, CReation of User-friendly Mobile services, PErsonalised for Tourism project (CRUMPET WP1). The list of user contexts that users may prefer to restrict includes: user ID, user location, routes, future destinations, service preferences, history of context modifications or use, personal details such as home address, telephone number and age, and payment credentials. Several main active participant roles are identified from these scenarios. A participant can play more than one role at any time. These include:

- *Owner*: Someone that owns the user context information.
- *Holder*: Some participant that collects and holds information about others, e.g., a *locator* such as a GPS device that determines a location, a *tracker* that maintains a history of user positions, or a *profiler* that monitors user interaction and builds a personal profile or user model.
- *Privacy Controller*: Someone whose duty it is to reveal as little of a context as possible to others. A privacy controller may also monitor the rights of others to share a context. The owner or a *broker* acting on his or her behalf often behaves as a privacy controller.
- *Requester*: A user of an owner's context. This may be the owner itself, another user, a provider, a *broker*, a *recommender*, a match-maker, an anonymizer or a geographical information server (*GIS*).

In addition, other middleware may be involved such as an *authorizer*. This is an authorization authority that is responsible for issuing the credentials used for the authentication and authorization of the other actors

TABLE 1 Summary of Features and Privacy Requirements for Different Context-Aware Service Use-Cases Q9 Q8

	Type of user context being transferred: the type of interaction in which it is transferred (conditions)	Context flow between the actors involved where the numbers and arrow indicate the sequence and direction of the control of flow
1.1	Location: Where am I? (Locator not under user control)	Owner 1> Broker <3,2> Locator Owner <6 Broker <5,4> GIS
1.2	Location: Where am I? (Owner controls locator)	Locator 1> Owner <5,2> Broker <4,3> GIS
2.1	Location: Use local service (Assumes 1.1/1.2)	Owner <6,1> Broker <5,2> GIS <4,3> Provider
2.2	Location: Use remote service (Assumes 1.1/1.2)	Owner <6,1> Broker <5,2> GIS <4,3> Provider
2.3	Location: Spatial navigation from location A to B (Assumes 1.1/1.2)	Owner <4,1> Broker <3,2> GIS
2.4	Location: Automatic awareness of Location (Assumes 1.1/1.2)	Tracker <6,1> Locator <5,2> GIS <4,3> Provider Tracker 7> Owner
3.1	Preferences: Owner directly describes preferences	Owner 1> Broker
3.2	Preferences: Profiler passively gathers preferences	Owner 2> Profiler
4.1	Preferences: Recommend services based on preferences (Assumes 1.1/1.2, 3.1/3.2)	Owner <4,1> Recommender Recommender <3,2> Broker
4.2	Location, time: Buddies are informed of each other's location (Assumes 1.1/1.2, 2.4)	Owner <4,1> Broker Broker <3,2> Tracker
4.3	General context: review history of context in an archive (Assumes 1.1/1.2)	Requestor, Owner <4,1> Tracker Tracker <3,2> Archive

involved within a particular interaction. Credentials can be issued specifying rights to use a context and the delegation of these rights to others.

At first sight, context sharing seems simple and controllable by the user: The user context resides under the control of the user or his or her authorized privacy controller. This is complicated when user contexts originate at participants other than the owner, for example, a location sensor or locator that resides in the network such as WLAN trilateration (use-case 1.1, Table 1). Similarly, a personal model of the user may be built through direct user interaction, by the user volunteering information (use-case 3.1, Table 1) or by some entity observing and profiling the user, possibly unknown to the user (use-case 3.2, Table 1). For push type services, user contexts need to be held by holders such as context trackers (use-case 2.4, Table 1). profilers may also model context usage via observation and compile a history of usage (use-case 4.3, Table 1). The authorization to use an owner's context is typically passed to an authorized principal and in some cases distributed further, e.g., a broker or recommender acting as a privacy controller may be required to pass a user context to a service provider (use-case 2.1–2.4, Table 1) or to other users (buddies use case 4.2, Table 1).

In all of the service interactions, some part of the user context may be revealed to others. In the context flow, the numbers and arrow indicate sequence and direction of control of flow, e.g., Owner <4,1> Broker indicates that the first message is from owner to the broker and the fourth message is from the broker to the owner.

Privacy Requirements

Privacy can be regarded both as a basic human right to be left alone or as a user freedom to not be observed by others. Typically a portion of privacy is traded in order to have greater convenience in order to interact with others such as other human users and service providers, e.g., in the buddy use-case, a user needs to reveal some of their user-context to others. Hence, in practice, the disclosure of the user context needs to be qualified. Golberg (2000) saw a different state of characteristics between true identity and true anonymity. He distinguished four levels of "nymity:" anonymity, non-reversible pseudonymity (not tied to a true identity), reversible pseudonymity (tied to a true identity), and identity. Maximum disclosure and invasion of user privacy occurs when a user's personal details are revealed and linked to additional contexts such as the current location or payment credentials. This suggests a common sense strategy and a set of requirements to protect user privacy as follows:

1. *Restricted disclosure*: when a user's personal details are given on a need to know basis.
2. *Owner-centered* context management: by default, the use and distribution of user contexts is governed by the owner or his or her designated privacy controller.
3. *Anonymity*: when user contexts are disclosed to others, they are linked to virtual IDs rather than real IDs if at all possible.
4. *Misuse of user context*: It should be defined, detected, and penalized.

ANALYSIS OF PRIVACY REQUIREMENTS FOR CONTEXT-AWARE SERVICES

Survey of Privacy Models

There are a variety of models to support user privacy (see Table 2). Conventional computer encryption and access control supports information privacy, which is confidentiality by restricting information access to those that can decrypt and access it. Conventional use of cryptography and access control typically involves public keys for confidentiality, digital signatures to

TABLE 2 Review of Privacy Solutions for Safeguarding User Context, Where √ Indicates that a Requirement is Present for a Particular Privacy Solution

Q1 Q8

	Restricted disclosure	Selective Nymity model	Owner-centered control of privacy	User context misuse: 1 defined, 2 detected & 3 penalized
Conventional	√			
Provide driven	√			1√
PET	√		√	
Privacy DRM	√	√		1–2√
Legislative	√	√	√	1–3√
Trust-based	√	√	√	
USHER	√	√	√	

guard against a misrepresentation of the information from the sender, and access control matrices. Such systems may be secure but have little privacy because of the use of signing and the lack of an anonymity model for the user.

Commercially, privacy is driven by provider-driven approaches such as P3P. This personalizes service access by allowing providers access to full personal information, acting as a trusted principal on behalf of owners. However the user has little control of the full context that the provider holds and it is thus open to abuse. The criticism of P3P includes: privacy is regarded as a negotiable commodity and the overall level of privacy depends on the user's ability to pay for it. Also, the process of sharing privacy information is provider-oriented. More specifically, P3P evaluation is done by matching user references against service provider policies. Users have no choice but to trust the service provider completely. Since P3P doesn't support authentication, there is no way to determine the legitimacy of the statements listed in service provider policy. Finally, although P3P documents are written in human readable format, they are difficult to understand and therefore machine assistance is required to evaluate them.

Privacy enhancing techniques (PET) seek to minimize or to eliminate the collection of personally identifiable information, for example, by acting as a broker to provide pseudo-anonymity by disassociating users' preferences from their real identity and by assigning users a pseudo identity. PET technologies are especially effective in minimizing unsanctioned user profiling and the gathering of private information by third parties. Although some sources consider PET to be a relatively effective solution to provide certain types of privacy (Kasanoff et al. 2001), PET does not address the overall problem. The main limitations of PET technology are that anonymizing a particular entity eliminates the possibility of information sharing and that the identity could be derived via monitoring the device. The latter operational mode requires a complex policy system to be capable of addressing user's guidelines regarding the sharing of private

information. These kinds of applications tend to be very complex and in order to be used effectively, some assistance is preferable. In some cases, in order to receive a particular service, some sharing of private information is required. Hence, what PET really needs is pseudo-anonymity.

Privacy can be regarded as a digital rights management or DRM problem (Gunter et al. 2004). This provides a strong model for owner context access to be accurately regulated but it blindly trusts that holders and authorized requestors act responsibly. This is open to potential abuse by context holders. There is no explicit user-centered "nymity" model.

Much of the world has progressively legislated broad fair information practices. These were codified in the OECD's (Organization for Economic Cooperation and Development) 1980 Guidelines on the Protection of Privacy and Transborder Flows of Personal Data (OECD). The main criticism of the legislative approach is its bounding to a particular geo-graphical region. The legislative approach requires an exchange of legal agreements between the user and service provider before any private information can be shared. The legal interpretation of these agreements can differ if the user is from the European Union (E.U.) and the service provider belongs to the United States.

Q10 The E.U.IST SECURE Project (Siegneur 2004) adopts a trust-based approach that supports an anonymity model but links it to evidence and to a persistent pseudonym to allow these to be evaluated and to be input into trust decisions.

The USHER (u-commerce services here for roamers) system presented in this paper is a hybrid approach that incorporates conventional security safeguards, such as confidentiality, integrity checks, authentication and authorization, and pseudo-anonymization, together with a policy-driven approach to personalization that is user-centered rather than provider-centered. The policy driven approach can be extended to utilize policy managers that can enforce an implementation of particular privacy legislation.

Survey of Privacy-Enabled Context-Aware Service Infrastructures

Our survey focuses primarily on agent-based context-aware service infrastructures. Whilst there are several prototype agent-based context-aware service infrastructures and some proposed privacy models (Kagal et al. 2003), few of the implemented context-aware infrastructures describe the integration of a comprehensive deployed privacy framework in sufficient detail. There have been a number of attempts to employ agent technology to deliver a location-aware system. Recent developments include projects such as CRUMPET (Poslad et al. 2001), HIPS (Benelli et al. 1999), and Gulliver's Genie (O'Hare et al. 2001). Although, these systems differ in terms of functionality and the agent architectures chosen for deployment,

certain common threats to the privacy of the user exist. All the agent systems use some kind of personal context sharing to deliver a range of services. This context is usually used by a second party to minimize and optimize the information flow. In terms of the services that each offers to tourists, Gulliver's Genie focuses more on the navigation and cultural requirements of tourists. This approach assumes the need to transfer a large amount of multimedia information and as a result supports intelligent pre-caching. In contrast, USHER takes a broader and more holistic view of the services that it provides to tourists, preferring to concentrate on the broad area of location-awareness.

Two general modes of operation can be identified: anonymous browsing and personalized service retrieval. Anonymous browsing assumes that the user ID is not revealed and all of the personal content sharing is done anonymously. Personalized service retrieval requires the real identity of the user to be revealed. Within agent-based frameworks, privacy issues can be addressed by having some kind of middle agent acting as a system portal that is user controlled and that is capable of anonymizing the user to the rest of the system. According to Decker et al. (1997), there are different patterns of interaction based upon who to reveal what to whom. In his model there are three main actors: users, requestors, and middle agents. He considers that there are three main types of middle-agents: matchmakers or yellow page agents that process advertisements, blackboard agents that collect requests, and brokers that process both. Taking privacy consideration into account, the most suitable solution is to employ broker-type middle agents. These agents can conceal a requester's identity from service providers and vice versa. They are also able to cope more quickly with a rapidly changing agent workforce. A problem here is that use of a hidden identity can reduce personalization.

CONTEXT-AWARE SERVICE INFRASTRUCTURE

The overall system design of the USHER system can be partitioned into the context-aware service infrastructure (this section) and the privacy model (following section). There exists a range of possible context-aware service infrastructures that includes combined location-sensors and PDAs, wireless transmission trilateration infrastructures, and personalized services where user preferences can be gathered indirectly by observation of user activity or directly via user participation.

Benefits of an Agent-Based Infrastructure

Using an intelligent agent-based infrastructure to deliver and access these services has several benefits. The use of agent autonomy coupled with

the use of agent semantic and social interaction protocols enables heterogeneous stake-holders such as users, mediators, and providers to control the service from their own perspective yet interoperate in a highly sophisticated manner. Agents can decide how much information they will reveal to providers, they can negotiate to get the best service level agreement, or refuse if their minimum constraints are not met. Agents can use a set of standard agent protocols, such as FIPA-ACL—a set of service neutral protocols—to support rich information and task sharing. Inherent in the FIPA-ACL model used is the notion of an ontological-based semantic model. Multiple domain specific ontologies, e.g., for tourism, personalization, and system management, can be integrated using ontology mediation and using a logic-based semantic language to reason about the ontology and the interactions. Also, inherent in the standard FIPA-ACL protocol model is a belief-based mentalistic model. For example, agents can share information when they believe that the other party does not already know the information and has a need to know it; agents can share tasks when they believe that another party has the capabilities and the intention to help them. Agents can be deliberative, they can build models of the environment in which they are situated and then can adapt their behavior to the state of the environment such as the network QoS, person, location and terminal. A key question is how much agent intelligence to situate in the mobile terminal—this depends upon the resources available in the mobile terminal and the desired level of software maintenance. Very small mobile terminals do not have enough resources to run sophisticated agents. Larger mobile terminals can run agents and therefore reap the benefits of agent-based deliberation to adapt content delivery to the network QoS and to personalize it.

Agents are especially suitable for the privacy protection system because of their ability to decompose and distribute tasks and to mask the identity of the originator of the task and information. Within a security system, the complexity is one of the major issues that can lead to intentional or unintentional misuse. Agent mediators can mask task complexity for users and thereby reduce the issue of security mis-configuration by users. Another important issue which justifies the use of intelligent agent technology is the agent's social abilities. By communicating with each other, agents can learn about new threats, revoke credentials more effectively, and help maintain user trust.

Agent-Based Context-Aware Service Infrastructure

An example of an agent-based infrastructure to support context-awareness is the CRUMPET (Creation of User-friendly Mobile Services PErsonalized for Tourism) project (Poslad et al. 2001) see Figure 1. This consists of

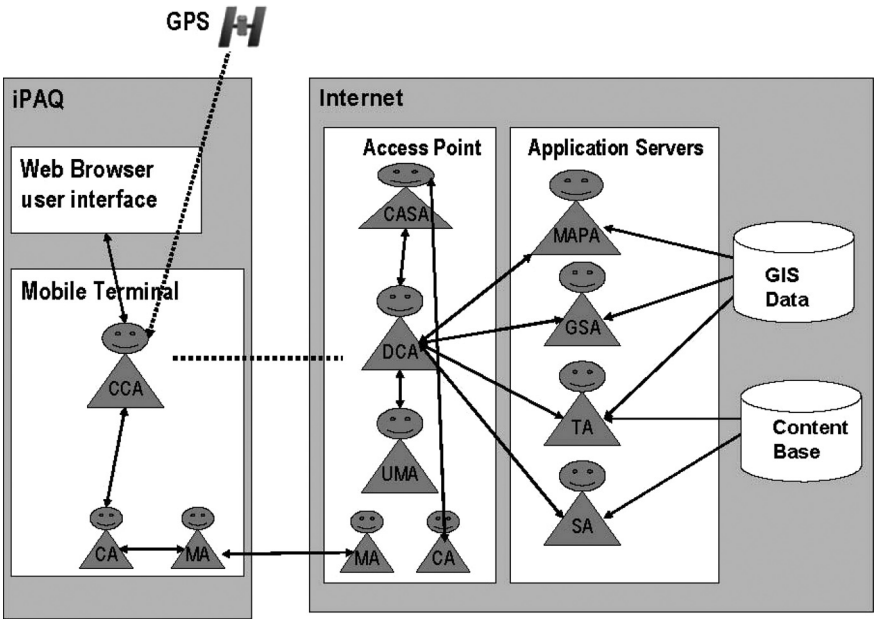

FIGURE 1 An agent-based architecture to support LAS services.

several FIPA agent platforms: one on the mobile terminal, one on the wireless access node, and one on the main wired Internet. These platforms support several agent services: CA and MA are network control and monitoring agents; the DCA is a broker agent; the CCA is a client agent, the UMA is a user modeling agent, and the MAPA, GSA, TA, and SA are the map agent, GIS agent, tour agent, and sights agent, respectively.

The original structure was based upon the microFIPA-OS agent middleware for small terminals and FIPA-OS agent middleware running on the Internet access node and servers (see Figure 1). Although, this system was successfully trialed in two cities (Schmidt-Belz et al. 2003), it was very complex and has proved to be very difficult to maintain. A second version of the system has been re-engineered and re-implemented using the JADE-LEAP agent platform (JADE 2004). This second version of the LAS system adheres to the basic architecture of the original system. It has been used during several M.Sc. projects and is being used to research privacy issues as part the USHER system.

PRIVACY MODEL

The requirements given previously suggest that a privacy system should integrate several kinds of safeguards (see Table 3) such as access control and a system to delegate and manage access rights. However, these two

TABLE 3 Relation of Context-Aware Privacy Requirements and Safeguards to Meet These

Requirement	Safeguard
Restricted access to user context	Access control matrix
Owner-centered management of owner's contexts	Use of a trusted principal and policy framework to manage access rights including creation, delegation
Pseudonymity	Broker/pseudonymizer
Management of misused user contexts	Social model of context sensors, context holders including monitors and penalisers

safeguards alone are not sufficient as we can't assume that all the parties involved act responsibly in their use of an owner's context or that we can adequately supervise access to an owner's context. In order to reduce the consequences of misuse, further safeguards are needed. This includes the use of anonymity (Kobsa et al. 2003), e.g., it is not necessary for every requester such as a GIS server, locator, or provider to hold a binding between a real owner ID and an owner's context. A further safeguard is to monitor misuse of contexts and enforce counter-measures to penalize misuse. Before discussing each of these safeguards in detail, the use of a semantic holistic model of the privacy to underpin these requirements is outlined.

A Semantic Model to Support Privacy

Because of the richness of the interaction between the different actors involved in exploiting and safeguarding the owner contexts, problems due to lack of a shared semantic model for privacy can arise (Tan et al. this volume). These include: the inability to interlink service processes that access the user-context with the security processes that safeguard them; the lack of a common terminology to allow owner contexts to interlink to management concepts that are understood by the different actors; and the lack of rich meta-data to manage the operation of the privacy safeguards. Hence we argue for not just a formal policy-model to support the management of owner preference access but also for a semantic model to support it. A semantic model of user preferences was modeled in the Protégé Ontology Development Environment.

A computational model of the preference model was exported as a Java-class hierarchy for use by JADE agents using the JADE plug-in for Protégé. A schematic model of the ontological model is given in Figure 2.

The V-SAT Model

The user context ontology links to a more general security model called V-SAT (viewpoints of safeguards) that protect assets (the items of value in

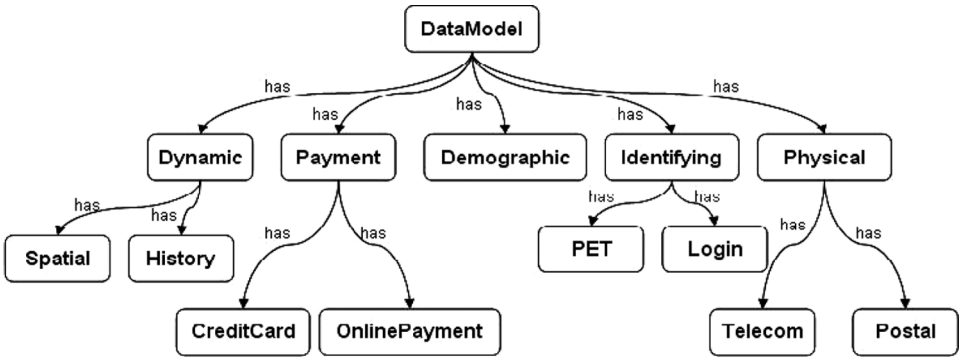

FIGURE 2 A schematic diagram of the ontology model of the user context.

the system) against threats (Tan et al. this volume) by considering the user context as a form asset (see Figure 3). Here, the assets are not just a client's personal record stored in a client device, they also include location, profiled personal information, client preferences, and personal details. As these are collected and used in different places, we need a sophisticated approach to safeguard them.

Restrict Access to User Contexts

The authority to access an owner's context is defined by the owner's preference policy. Similarly, a holder entity is involved in the information sharing process. However, in order to do so a specific permission has to be granted. The proposed solution is a unified approach based on P3P

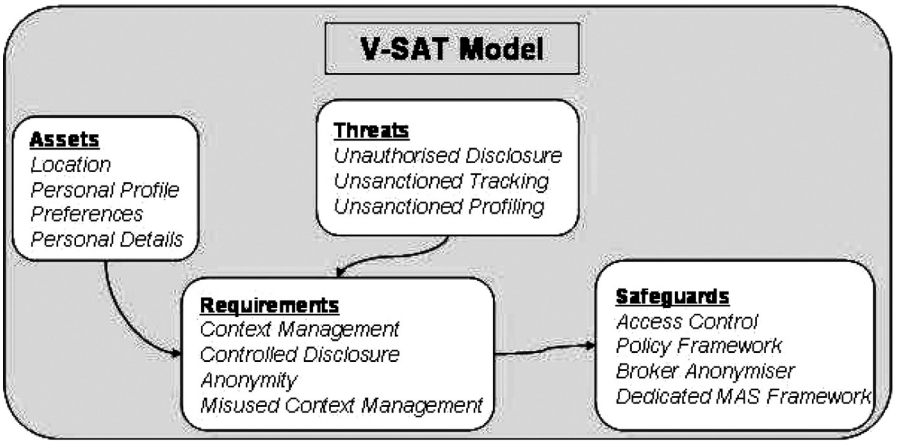

Q11 FIGURE 3 V-SAT model.

TABLE 4 Access Control Diagram

Group access level	Generic	GIS	Educational	Personal	Commerical
Level 1	R1			R1	
Level 2	R2			R2	
Level 3	R3			R3	
Level 4	R4			R4	
Level 5	R5			R5	

and enhanced by access control mechanisms to specify and interpret preference policies and information requests within a location–aware environment.

The policy consists of four main concepts. The first two refer to the roles within the system representing the authorizer and the owner of the policy. The other two concepts refer to the requester's clearance levels and the rule set. Requesters are separated into groups based on the service type such as generic, GIS, educational, personal, and commercial (see Table 4). The user assigns each requester an access level depending on its group membership. The combination of group and access level determines the set of rules {R1..Rn} relevant to a particular requester based upon a grid-based access control table of requester group versus access level.

When making a request, a requester specifies their ID, the pseudonym of the owner (or holder), the assets requested, and the purpose of the request. A credential-based mechanism is used to bind the service identifier to a specific type of credential. The token plays a role similar to the traditional $\times 509$ certificate but it is more general; it can bind any credential type to a service identity.

Policy-Based Management for Access to User Contexts

This subsection presents a formal description of the suggested policy-based privacy system, which is based on an access control matrix.

Roles

Every entity can belong to one of the following roles: *Owner* (e) $||Holder(e)||PrivacyController(e)||Authorizer(e)||Requester(e)$Holder. The requester and owner roles can overlap.

Information Owner. For every entity performing an owner role, the following actions can be performed: agree, deny, specify, and reveal.
Agree allows an owner to reveal information, deny allows an owner to deny information sharing, specify allows a holder to specify rules regarding

sharing, and reveal reveals an owner's policies to authorized parties for evaluation purposes.

Information Holder. The holder entity performs two major functions: it collects the user context information and it shares it with other entities. When it comes to information sharing the action set for an information holder is identical to the actions that an owner entity performs. The collection of information is done by a monitor action. An information holder can perform the following actions: agree, reject, specify, reveal, and monitor.

Privacy Controller. This is the system entity that is responsible for policy evaluation procedures. The following actions can be performed by the privacy controller: anonymize, requestcredentials, requestpolicy, evaluate, notify, reject, reveal, and requestInfo.

The anonymize action creates a pseudo identity that is used for later communication. Requestcredentials is performed in order to obtain the credentials of a requester entity. Similarly, requestpolicy and requestinfo actions are executed to retrieve the appropriate information from the information holder. The evaluate action is concerned with making a decision regarding sharing personal information. The notify, reject, and reveal actions deal with handling the results of the evaluation.

Requester. An information requester can perform two main actions: request information about a particular entity and revealcredentials.

Authorizer. The Authorizer entity performs a role similar to a traditional certification authority. It issues credentials and can revoke them. Credentials can be requested from the authorizer with the grant permission field set. Only authorized users such as owners can grant permission. This promotes authorized sharing of the user context.

Entity Classification Model

This is defined as follows:

$$g - \textit{is an EntityGroup};\ a - \textit{is an AccessLevel};\ p = \textit{PrivilegesSet}$$
$$p = \textit{Clearance}(g, a)$$

Privileges set p consists of subsets s1 and s2.

$$p = s1 + s2$$

s1 is a group level set of privileges, where s2 contains privileges corresponding to a specific member of the group. This approach makes the system more scalable. Entity specific rules have precedence within the system and have the ability to overwrite group specific rules.

$r1 -$ *rule belonging to the subset* $s1$;
$r2 -$ *rule belonging to the subset* $s2$;
$o -$ *request for information*;
$(\ Evaluate(o, r1) = reject\ )\&\&(\ Evaluate(o, r2) = reveal\ )$
$\Rightarrow Overwrite(r1, r2)$

The other important system rule states that even if different entity groups have the same access level their rule sets will be different but may overlap.

Policy Evaluation

Policy evaluation is done by evaluating requests and credential tokens against the information owner's preference policies (Titkov et al. 2004) using the algorithm given next. The policy rules are defined in the software. When a request is made for access to the user context via the broker, the broker requests the policy for the use of that user context from the owner. The broker then evaluates the request based upon the policy rules, the requestor's credentials, and the description of the user context. There are three possible outcomes of the evaluation: reject service requests that are not permitted by the client or if the service does not have a valid token; notify the client of service requests; and reveal the requested data when the requester meets all the owner's criteria for accessing the data in terms of a valid token, access level, entity group (entity id), constraints, time, location, and purpose. The algorithm for policy evaluation consists of three steps and is given in pseudo-code. Let (X) be the information requester, (Y) be the information owner (or holder), (T) the requester's token, (Rule) is a rule from the preference policy, and (Assets) are assets.

1. Validate the requester's token:
   1.1 Check the validity period if it has expired.
   1.2 Check if the request ID matches the Token ID.
2. Collect list of privileged information available for the requester:
   2.1 Get service entity type (commercial, GIS, education or personnel).
   2.2 Gets clearance level from clearance database.
   2.3 Compare rule clearance level with requester clearance level.
   2.4 Get generic privileged assets available for the requester.
   2.5 For every rule get the effect type (notify, reveal or reject).
   2.6 Add privileged assets based on the rule priority.
   2.7 Remove privileged assets based on the rule priority.
3. Validate the requested data with the compiled list of privileged data. Rules are executed based on priority of the effect type such that: reject is first priority, notify is the second priority, and reveal is the third priority.
   3.1 Get the information request

3.2 Check service request against privileged assets based on priority. If a request is not included in a privileged assets list then reject the request.
3.3 If a request includes data in the notify assets list, then notify the owner.
3.4 If a request includes data in the privileged assets list then reveal the assets.

Anonymization via Authorized Brokers

Generally, third parties such as brokers are used in service interaction to facilitate the interaction of users and services and to mediate between different service providers. In order to deliver such services some sharing of personal information is required (Kasanoff et al. 2001). Generally, a broker can be designed to act as a trusted third party, to be primarily user-driven or to be provider-driven. The USHER broker system treats the broker as user driven. It acts as the user's principal and serves to support persistent pseudo-anonymity for the user to service providers.

During online transactions there is a possibility that a third party can gather information about the user based on observing a user's actions, i.e., create a user model. This information can then be sold on and shared. The user has little control over this particular process. To partially resolve this issue the broker uses pseudonyms to anonymize the user to any user modelers (see Figure 4) and to reduce the possibility of information misuse by service providers. It can provide results based upon the user's personal preferences, without identifying the user. A persistent pseudonym value allows the broker agent to maintain personal information for users for use with successive service invocations.

Figure 4 demonstrates a scenario where a user agent wants a service to be personalized without revealing his identity to the service provider agent. The interaction is as follows:

1. User updates the personal preferences through the broker.
2. User sends a request through the broker.

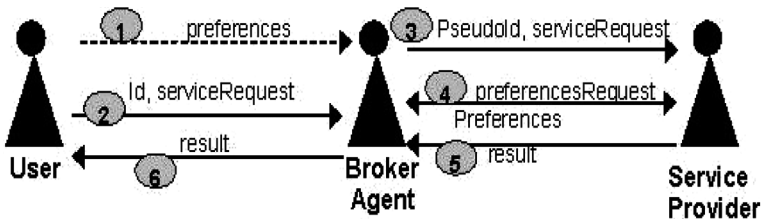

FIGURE 4 Pseudonymous identification scenario.

3. The broker uses the pseudonym of the user to request a personalized service from the service provider.
4. The service provider requests preferences from the broker using the pseudonym value.
5. The service provider sends the results to the broker.
6. The broker subsequently sends the results to the user.

In practice, a broker agent is more than a unilateral or bilateral separator of privacy between a user and provider, it controls how to match service requests and user preferences with the provider capabilities. The broker agent can selectively reveal information about an owner's context to others. It can act as a semantic mediator, converting data from a user to a provider semantic context. In most cases, users use a pull type of interaction with the broker agent, delegating to it to select a suitable service provider. The user can also delegate to it to use a push type of interaction when registered interests are triggered. Apart from this, the broker can also monitor network bandwidth, handle potential service and network provider failures, and handle content adaptation. Middle-agents may also have to be trusted by providers. A broker may find more than one suitable provider. A trustworthy agency has to guarantee providers that it can exploit competitive services in a disinterested and unbiased way (Pearson 2003).

System Implementation

The computational part of the privacy framework resides within an authorized principal designated by, trusted by, and under the control of the user. In our framework, this is a broker agent that provides privacy services to the other agents. A privacy control module is responsible for the actual privacy control. The internal architecture of this component can be depicted as follows:

- *Identity Manager* is used to support user anonymity. It is responsible for pseudonymous communication within an USHER service. Identity Manager generates pseudonyms and provides a mapping between user's real ids and their pseudonyms stored within the ID Repository.
- *Ontology Manager* is responsible for mapping different domain ontologies into a shared ontology. A shared ontology allows the broker agent to match and relate semantics within the system.
- *Access Control Manager* deals with requests for user's private information and reveals this data when necessary. The Access Control Manager updates the user's rule set with the appropriate rule whenever the user changes his or her privacy settings.

During a third-party request, the Access Control Manager receives the information request and token to be evaluated against the user's rule set. The evaluation engine evaluates the pattern part of each rule against requests and tokens to determine if there is a match. After evaluating the rule set and finding a match, the evaluation engine returns the behavior of the rule (reveal, reject, or notify). The Access Control Manager then provides the appropriate access to a resource based on the return value of the evaluator engine.

System Evaluation

A Location Aware Religious Services (LARS) application was implemented to demonstrate use-cases 2 and 3 (Table 1). The particular LARS developed is targeted mainly at the Muslim student population of Queen Mary College, but other religious content is also supported. It is a well known that Muslims have to pay their duties to God five times a day. Prayer times vary depending on the time of sunrise and the religious sect. It is challenging for mobile users to identify the times and location of religious services when traveling in an area not familiar to them. It was envisaged that both "pull" and "push" type of interaction, e.g., notification when religious service events are starting, would be useful. An example use-case is the following: Muhammad Smith from Birmingham, England is a first year student visiting Queen Mary College in London. While looking around campus, he realizes that the next prayer time is fast approaching, so he subscribes via his PDA to a LARS for aid in finding a mosque nearby. To wrap the religious service, a religious service agent (RSA) represents the religious institutions and acts as the interactive "service provider" for that

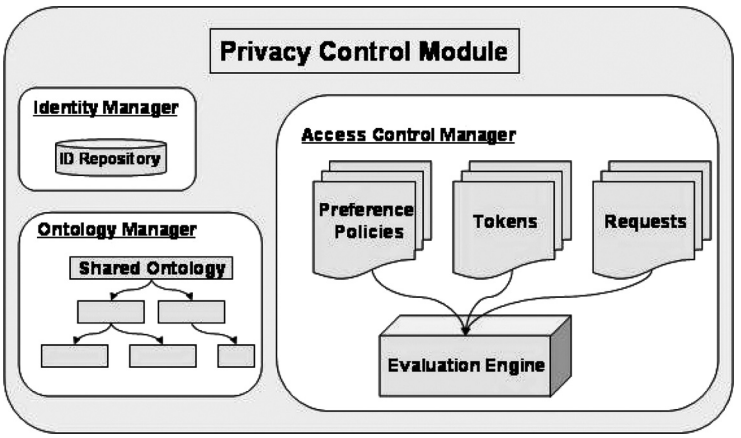

Q12 FIGURE 5 Privacy control module detailed architecture.

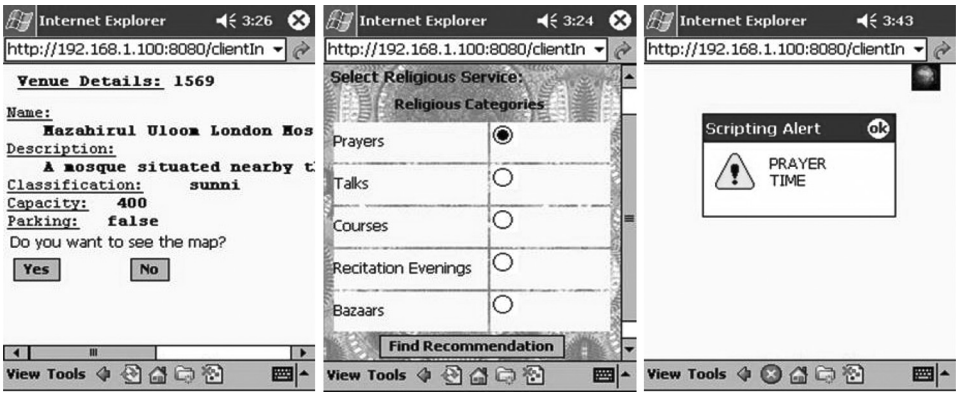

FIGURE 6 Screen shots of a location-aware religious service that supports use-case 2.4 (Table 1).

specific religious institute. The information generated by the RSA for this scenario is given in Figure 6.

In another application, a restaurant services application, the policy restaurant policy evaluation is done by evaluating request and credentials against preference policies. There three possible outcomes of the evaluation: reject, reveal, and notify. Finally, a tourism sights services for Queen Mary was evaluated by 20 users using an abbreviated SUMI (software usability measurement inventory) questionnaire. The usability scores were above average.

FURTHER WORK AND CONCLUSIONS

Further Work

The approach adopted can be extended in several ways. User modeling can reduce the information overload to mobile users whose locus of attention is constrained because they are often non-stationary, they are sensitive to disturbances in their variable physical and social environment, and their communicator's resources for display and processing is limited. However, the acquisition of the user model and maintenance of a fresh user model is challenging. In the USHER system, the user model is acquired through direct interaction with an owner. In the previous CRUMPET system (Poslad et al. 2001), the user model was derived through an analysis of observation of the user's interaction with the system. The issue is that much user interaction may be needed before an accurate model of the user can be captured. Some form of aggregation of the individual user models coupled with unsupervised learning may help a new user model to be derived using information from a related user cluster. To do this and protect privacy, the

owner would need to explicitly agree to the broker sharing information with some entity that supports user cluster modeling.

One of the privacy requirements given previously was that the misuse of user contexts should be defined, detected, and penalized, but this has not yet been addressed. There are several ways that a penalization mechanism could be implemented. One way is that reputation models could be generated for each requestor. Misusers of a user context could have their reputation lowered when infringements are detected. The owner or his or her principal could take into account the reputation of the provider when granting access to his or her user context. The difficult part of this is the detection of infringements. Self-interested requestors could sell user-contexts to other unauthorized requestors and they could use a variety of techniques to evade detection such as hiding the distribution of user contexts using encryption. One possible way to detect this is that requestors must be licensed and agree to their processes for using user-contexts to be auditable and tamper-proof.

CONCLUSIONS

Privacy is a major challenge when offering context-aware services for mobile users. The complexity and dynamicity of context-aware services requires a composite approach to providing privacy. In order for users to more conveniently access services, some sharing of private information, including a partial disclosure of the user context, but not necessarily the real user identity, may be required. In the USHER system, where location-aware tourism information is personalized for mobile users, a variety of approaches have been used to integrate and to safeguard personalized information. A granular model of anonymity coupled with a richer model of policy-based access control has been used. It can support some trading of partial privacy protection against more convenient service access. Although this overall domain is extremely complex, an agent-based approach seems successful in reducing its complexity. Agents' rich communication, and their task and function decomposition ability coupled with a strong privacy model can enable user privacy to remain user-driven in the face of commercially driven provider-centered service access.

REFERENCES

Berners-Lee, T., J. Hendler, and O. Lassila. 2001. The semantic Web. *Scientific American* 284(5):34–43.

Benelli, G., A. Bianchi, P. Marti, E. Not, and D. Sennati. 1999. HIPS: Hyper- Interaction within physical space. In *Proceedings of the IEEE International Conference on Multimedia Computing and Systems*, pages 1075–1078, Florence, Italy.

CRUMPET project (IST-1999–20147), Work Package 1, Deliverable D1.1, www.ist-crumpet.org.

Decker, K. 1999. Middle agents for the Internet. In *Proceedings of the 15th International Joint Conference on Artificial Intelligence*, pages 578–583, Nagoya, Japan, Morgan Kaufmann.

Q4 Ferber, J. 1999. *Multi-Agent Systems*. ■: Addison-Wesley.

Goldberg, I. 2000. *A Pseudonymous Communications Infrastructure for the Internet*, Ph.D. thesis, Univ. of California at Berkeley.

Gunter, C. A. and M. J. May. 2004. A formal privacy system and its application to location based services. In *Workshop on Privacy Enhancing Technologies (PET-2004)*, Toronto, Canada.

Hazas, M. and A. Ward. 2002. A novel broadband ultrasonic location system. In *Proceedings of the UbiComp*

Q5 *2002*, pages 264–280, Göteborg, Sweden.

The JADE, Java Agent Development Environment, http://jade.cselt.it.

Kagal, L., S. Cost, T. Finin, and Y. Peng. 2003. A policy language for pervasive systems. *4th IEEE Int. Work-*

Q6 *shop on Policies for Distributed Systems and Networks*.

Kasanoff, B. 2001. *Making it Personal: How to Profit from Personalization without Invading Privacy*. ■ Perseus

Q4 Publishing.

Kobsa, L. 2002. Personalised hypermedia and international privacy. *Communications of the ACM* 45(5): 64–67.

Kobsa, L. and J. Schreck. 2003. Privacy through pseudonymity in user-adaptive systems. *ACM Transactions on Internet Technology* 3(2):149–183.

Liu, J. and V. Issarny. 2004. Enhanced reputation mechanism for mobile ad hoc networks. In: *Trust Management*, eds. C. D. Jensen et al., LNCS vol. 2995, 93–108, Springer Verlag.

O'Hare, G. M. P. and M. J. O'Grady. 2003. Gulliver's genie: A multi-agent system for ubiquitous and intelligent content delivery. *Computer Communications* 26(11):1177–1187.

Understanding Privacy OECD Guidelines, Stanford, http://xenon.stanford.edu/ruchika/p3p/under standingPrivacyOECDGuidelines.html.

The Platform for Privacy Preferences Specification, http://www.w3.org/TR/P3P.

Pearson, S. 2003. A trusted method for self-profiling in e-commerce. *LNCS* 2631:177–193, Springer

Q7 Verlag.

Poslad, S., H. Laamanen, R. Malaka, A. Nick, P. Buckle, and A. Zipf. 2001. CRUMPET: Creation of user-friendly mobile services personalised for tourism. In *Proceedings of the 3G2001 Mobile Communication Technologies*, London, pages 28–32, London, United Kingdom.

The Protege Editor, http://protege.stanford.edu/index.html.

Schmidt-Belz, B., H. Laamanen, S. Poslad, and A. Zipf. 2003. Location-based mobile tourist services- first

Q5 user experiences. In *Proceedings of the ENTER 2003 Conference*, pages 115–123, Helsinki, Finland.

Signeur, J.-M. and C. D. Jensen. 2004. Trading privacy for trust. In: *Trust Management*, eds. C. D. Jensen

Q4 et al., LNCS vol. 2995, 48–62, ■: Springer Verlag.

Tan, J. J., S. Poslad and L. Titkov. 2004. A semantic approach to harmonising security models for open services. *J. Applied Artificial Intelligence* 20(2–3).

Titkov, L., S. Poslad, and J. J. Tan. 2004. Enforcing privacy via brokering within nomadic environment. In *Proceedings of the 4th International Symposium From Agent Theory to Agent Implementation*, (AT2AI-4), In *Cybernetics and Systems 2004*, Volumes 1 + 2, Vienna, Austria. Austrian Society for Cybernetic Studies.

Want, R. 1992. The active badge location system. *ACM Transactions on Information Systems (TOIS)* 10(1):91–102.