

Semantic Web Technologies to Reconcile Privacy and Context Awareness

Fabien L. Gandon and Norman M. Sadeh

School of Computer Science - Carnegie Mellon University
5000 Forbes Avenue, Pittsburgh, PA 15213-3891, USA
{Fabien.Gandon, Norman.Sadeh}@cs.cmu.edu
<http://www-2.cs.cmu.edu/~sadeh/>

Abstract. Increasingly, application developers are looking for ways to provide users with higher levels of personalization that capture different elements of a user's operating context, such as her location, the task that she is currently engaged in, who her colleagues are, *etc.* While there are many sources of contextual information, they tend to vary from one user to another and also over time. Different users may rely on different location tracking functionality provided by different cell phone operators; they may use different calendar systems, *etc.* In this article, we describe work on a Semantic e-Wallet aimed at supporting automated discovery and access of personal resources, each represented as a Semantic Web Service. A key objective is to provide a Semantic Web environment for open access to a user's contextual resources, thereby reducing the costs associated with the development and maintenance of context-aware applications. A second objective is, through Semantic Web technologies, to empower users to selectively control who has access to their contextual information and under which conditions. This work has been carried out in the context of *myCampus*, a context-aware environment aimed at enhancing everyday campus life. Empirical results obtained on Carnegie Mellon's campus are encouraging.

1 Introduction

Increasingly, application developers are looking for ways to provide users with added levels of convenience and ease of use through functionality that is capable of capturing the context within which they operate. This may involve knowing where the user is located, the task she is currently engaged in, her eating preferences, who her colleagues are as well as a variety of other contextual attributes. While there are many sources of contextual information, they tend to vary from one user to another and also over time. Different users may rely on different location tracking functionality provided by different cell phone operators; they may use different calendar systems, *etc.* Traditionally, context-aware applications and services have been hardwired to predefined sources of contextual information (e.g. relying on a particular set of sensors and protocols to track a user's locations). As a result, they remain prohibitively expensive to build and maintain and are few and between. We argue that what is needed is a more open environment, where context-aware applications can automatically discover and access a user's personal resources such as her calendar or location tracking functionality. This can be done by viewing each source of contextual information (or personal resource) as a Web service. Unfortunately, current Web Services standards

such as UDDI [16] or WSDL [25] are not sufficient when it comes to describing a user's personal resources and to enabling automated access to them by context-aware applications. Another challenge, as we move towards more open platforms for access to a user's personal information, revolves around privacy issues. Users should be able to retain control over who has access to their personal information under different conditions. For instance, I may be willing to let my colleagues see where I am or access my calendar activities between 8am and 5pm on weekdays but not over the weekend. In addition, I may want to fine tune the granularity of the answer provided to a given query, depending on the context of that query. For instance, I may be willing to disclose the room that I am in to some people but only the city where I am to others. In fact, I may even want to give different answers to different people, telling my secretary I am off to see my dentist, while telling my customers I am busy in a meeting.

In this paper, we introduce a Semantic Web architecture aimed at supporting the automated discovery and access of personal resources in support of a variety of context-aware applications. Within this architecture, each source of contextual information (e.g. a calendar, location tracking functionality, collections of relevant user preferences, organizational databases) is represented as a Semantic Web service. A central element of our architecture is its *semantic e-Wallet*, which acts as a directory of contextual resources for a given user, while enforcing her *privacy preferences*. Privacy preferences enable users to specify what information can be provided to whom in different contexts. This also includes *obfuscation rules* that allow users to fine-tune the accuracy or inaccuracy of the information provided in response to different queries under different conditions.

We have validated our architecture in the context of *myCampus*, a context-aware environment aimed at enhancing everyday campus life at Carnegie Mellon University (CMU). The environment revolves around a growing collection of task-specific agents capable of automatically accessing a variety of contextual information about their users such as their locations, calendar activities as well as a variety of other attributes and preferences (e.g. context-aware restaurant concierge, context-aware message filtering agent, *etc.*). Students access the environment from PDAs over the campus's 802.11 wireless LAN. Empirical results obtained with a group of students over a period of several days are briefly summarized at the end of this article. While, in this paper, we focus on scenarios involving individual users, it should be noted that our architecture extends to scenarios where users are entire organizations. In this context, both organizations and individual users could each have one or more Semantic e-Wallets capable of leveraging a variety of individual or organizational knowledge subject to a rich set of privacy (or confidentiality) constraints.

The remainder of this article is organized as follows. Section 2 provides a brief overview of the state of the art in context-awareness, privacy, Web Services and the Semantic Web, emphasizing limitations of the work reported so far in the literature. In Section 3, we provide an overview of our Semantic Web environment for context-awareness and privacy. Section 4 focuses more specifically on the Semantic e-Wallet and includes a high-level scenario outlining its operation in response to a query about the current location of a user. Section 5 introduces the three layers of knowledge required to support the e-Wallet functionality. Section 6 discusses the e-Wallet's

current implementation, which is based on OWL Lite [24], XSLT transformations [27], and JESS [10]. Sections 7 and 8 provide further details on the e-Wallet's three layers. Section 9 summarizes what we view as the main contributions of our work. It includes a brief discussion of experiments carried out in our *myCampus* environment along with some concluding remarks.

2 Prior Work

Prior efforts to develop context-aware applications are many. Early work in context-awareness includes the Active Badge System developed at Olivetti Research Lab to redirect phone calls based on people's locations [23]. The ParcTab system developed at the Xerox Palo Alto Research Center in the early nineties relied on PDAs to support a variety of context-aware office applications (e.g. locating nearby resources such as printers, posting electronic notes in a room, *etc.*) [21, 22]. Other relevant applications that have emerged over the years range from location-aware tour guides to context-aware memory aids. More recent research efforts in context-awareness include MIT's Oxygen [6], CMU's Aura [11] and several projects at Berkeley's GUIR (e.g. [14]) to name just a few.

While early context-aware applications relied on *ad hoc* architectures and representations, it was quickly recognized that separating the process of acquiring contextual information from actual context-aware applications was key to facilitating application development and maintenance. Georgia Tech's Context Toolkit represents the most significant effort in this direction [7, 8]. In the Context Toolkit, widgets act as wrappers that provide access to different sets of contextual information (e.g. user location, identity, time and activity), while insulating applications from context acquisition concerns. Each user (as well as other relevant entities such as physical objects or locations) has a context server that contains all the widgets relevant to it. This is similar to our notion of e-Wallet, which serves as a directory of all personal resources relevant to a given user (e.g. relevant location tracking functionality, relevant collections of preferences, access to one or more calendar systems, *etc.*). Our Semantic e-Wallet however goes one step beyond Dey's Context Toolkit. It makes it possible to leverage much richer models of personal resources - what personal information they give access to, when to access one rather than the other, how to go about accessing these resources. In addition, it includes access control and obfuscation functionality to enforce user privacy preferences. This richer model is key to supporting automated discovery and access of a user's personal resources by agents. In other words, while the Context Toolkit focuses mainly on facilitating the development of context-aware applications through off-line, re-use and integration of context-aware components (i.e. widgets), our architecture emphasizes real-time, on-the-fly queries of personal resources by context-aware agents. These queries are processed through several layers of functionality that support automated discovery and access of relevant personal resources subject to user-specified privacy preferences.

Our work builds directly on recent efforts aimed at moving the Web from an environment where information is primarily made available for human consumption to one where it is annotated with semantic markup that makes it understandable to computers (and intelligent agents). These efforts are part of a long-term vision generally

referred to as the *Semantic Web* [1, 13]. They have already resulted in a succession of semantic markup languages [4, 24] as well as early efforts to define Web Service ontologies and markup in the context of languages such as DAML-S [5]. In our work, we have relied on the use of DAML+OIL [4] and more recently OWL [24] to represent contextual information (e.g. location, calendar activities, social and organizational relationships, etc.) and privacy preferences and on Semantic Web service concepts to support the automated discovery and access of personal and public resources.

The notion of e-Wallet as introduced in systems such as Microsoft's .NET Passport is not new. However current implementations have been limited to storing a very small amount of information and offer very restricted control to the user when it comes to specifying what information can be made available to different services. For instance, in Passport, users can specify whether or not they are willing to share parts of their profiles with all participating sites but cannot distinguish between different participating sites. Our notion of Semantic e-Wallet lifts these restrictions and allows users to control access to any of their personal resources. It also allows for multiple sources of similar information (e.g. multiple calendars or multiple location tracking functionality) and for functionality that can dynamically select which of these resources to tap based on the context and the nature of the query at hand (e.g. using your car's GPS system when you are driving and your cell phone operator's location tracking functionality when you are not).

Last but not least, our notion of semantic e-Wallet extends recent efforts to develop rich languages for capturing user privacy preferences such as P3P's APPEL language [26]. It does so by making it possible to leverage any number of domain ontologies and by allowing for preferences that relate to any number of contextual attributes. In addition, it allows users to specify obfuscation rules through which they control the level of accuracy (or inaccuracy) at which their contextual information is disclosed to different parties under different conditions. This includes telling some people which room you are in, while simply telling others whether you are at work or not, or whether you are in town or not. It also includes scenarios where you might want to pretend you are in one place, while you are really elsewhere. Also, while the security community has developed powerful languages to capture access control privileges such as the Security Assertion Markup Language (SAML) [17] and the XML Access Control Markup Language (XACML) [18] for defining access control rules, these languages do not take advantage of Semantic Web concepts.

3 Overall System Architecture

We consider an environment, where, over time, users purchase (or subscribe to) different sets of task-specific agents. These agents are each intended to assist them in the context of different activities (e.g. scheduling meetings with colleagues, reminding them of purchases they need to make, arranging trips or filtering incoming messages). To function, each agent needs to access some information about its user as well as possibly other users. Access to a user's personal (or contextual) information is controlled by that user's e-Wallet subject to privacy enforcement rules. The e-Wallet Manager (or simply e-Wallet) serves as a repository of static knowledge about the user – just like .NET Passport, except that here knowledge is represented using OWL.

In addition, the e-Wallet contains knowledge about how to access more information about the user by invoking a variety of resources, each represented as a Web Service. This knowledge is stored in the form of rules that map different contextual attributes onto one or more possible service invocations, enabling the e-Wallet to automatically identify and activate the most relevant resources in response to queries about the user's context (e.g. accessing the user's calendar to find out about her availability, or consulting one or more location tracking applications in an attempt to find out about her current location). Privacy enforcing rules, also stored in the e-Wallet, ensure that information about the user is only disclosed to authorized parties, taking into account the context of the query. They further adjust the accuracy or inaccuracy of the information provided subject to obfuscation rules set by the user.




Fig. 1. myCampus architecture: a user's perspective - the smiley faces represent agents.

Figure 1 provides an overview of our Semantic Web environment. It illustrates a situation where access is from a PDA over a wireless network, as is the case in *my-Campus*, the environment in which we have instantiated our architecture. However, our architecture extends to fixed Internet scenarios and more generally to environments where users can connect to the infrastructure through a number of access channels and devices – information about the particular access device and channel can actually be treated as part of the user’s context and be made available through her e-Wallet.

As can be seen in Figure 1, other key elements of our architecture include:

- One or more *Platform Managers* that serve the dual role of Directory Facilitators and Agent Management Systems, as defined in FIPA [9]. They manage the agents running at their sites, and maintain white and yellow page directories of these agents and the services they provide.
- *User Interaction Managers* - typically one per user - that are responsible for interactions with the user. This includes managing login sessions as well as interactions with the user’s agents and her e-Wallet. Because different users interact with different sets of agents, this also includes the dynamic generation of interfaces for interacting with these agents and the customization of these interfaces to the current interaction context (e.g. particular access device). Communication with the User Interaction Manager typically takes place through a number of APIs, e.g. an Instant Messaging API, an HTTP/HTML API, *etc.*

Clearly, agents are not limited to accessing information about users in the environment. Instead, they also typically access public Web Services, Semantic Web annotations, public ontologies and other public resources. On CMU’s campus, where we have deployed *myCampus*, this includes access to a variety of services such as 23 restaurant web services or a public weather forecasting web service.

In the following sections, we focus on the e-Wallet functionality. Additional details on *myCampus* and some of the agents we have deployed can be found in [20].

4 A Semantic E-Wallet

The e-Wallet is a central element of our Semantic Web architecture for context-awareness and privacy. It provides a unified and secure semantic interface to all the user’s personal resources, enabling agents in the system, whether working for the owner of the e-Wallet or for other users, to access and, when appropriate, modify information about the user subject to that user’s privacy preferences (e.g. not just determining whether the user is available between 3 and 4pm but also, possibly, scheduling a meeting at that time). The e-Wallet is *not* a static information repository. While it does contain some static information about the user, it is an agent acting as clearinghouse and gatekeeper for a user’s personal resources. Its knowledge about the user, her personal resources and preferences falls into four categories:

1. *Static knowledge*. This context-independent knowledge typically includes the user’s name, her email address, employer, home address as well as context-independent preferences (e.g. “I like spicy vegetarian cuisine”). This knowledge, like all other in the e-Wallet, can be edited by the user via the User Interaction Manager.

2. *Dynamic knowledge*. This is context-sensitive knowledge about the user, often involving a variety of preferences such as “When driving, I don’t want to receive instant messages”.
3. *Service invocation rules*. These rules help leverage information resources external to the e-Wallet – both personal and public. They effectively turn the e-Wallet into a semantic directory of personal resources that can be automatically discovered and accessed to process incoming queries. Specifically, service invocation rules provide a mapping between contextual attributes and personal resources available to access these attributes, viewing each personal resource as a Semantic Web service. An example of one such mapping is a rule indicating that a query about the user’s current activity can be answered by accessing her Microsoft Outlook calendar. We have developed Web Service wrappers for a variety of personal resources such as Microsoft Outlook Calendar or location tracking functionality. Service invocation rules are not limited to providing a one-to-one mapping between contextual attributes and personal resources. Instead, they can leverage rich ontologies of personal resources, enabling the e-Wallet to select among a number of possible personal resources based on availability, accuracy and other relevant considerations. For instance, in response to a query about the user’s location, the rules can specify that, when the user is driving, the best method available is the GPS in her car. If she is at work and her wireless-enabled PDA is on, her location can be obtained using location tracking functionality running over the enterprise’s wireless LAN. If everything else fails, her calendar might have some information about her location.

Finally, it should be noted that to answer queries about the user, additional mapping rules that support automated discovery and access of public services may also be needed. For instance, a query like “Tell me whether Fabien is in a sunny place right now” will typically require accessing Fabien’s location as well as a public weather service.

4. *Privacy preferences*. These preferences encapsulate knowledge about what information about herself the user is willing to disclose to others under different conditions. These preferences themselves fall into two categories:
 - a. *Access control rules*. These rules simply express who has the right to see what information under different conditions. As other rules, they are expressed in OWL and can leverage a variety of domain-specific ontologies. An example is “My location should only be visible to members of my team during week days between 8am and 5pm”.
 - b. *Obfuscation rules*. Often user privacy preferences are not black-and-white but rather involve different levels of accuracy or inaccuracy.
 - *Obfuscation by abstraction* is about abstracting away some details about the user’s current context such as telling people whether or not you are in town but not your exact location.
 - *Obfuscation by falsification* aims at supporting scenarios where the user may not want to appear as if she is withholding information but would rather provide false information. Here again, you may not want to reveal to your customers that you are meeting with the competition or tell them you are at the hospital.

As indicated earlier, all the above knowledge is represented in OWL and leverages a number of relevant ontologies (e.g. ontologies about contextual attributes, personal resources, as well as more specific knowledge such as cuisine types and food preferences or message types and message filtering preferences).

Before delving deeper into the details of the e-Wallet, a scenario will help illustrate the key steps it goes through in processing incoming queries (Figure 2). For the sake of argument, we will assume a query submitted by a user (Norman) to the e-Wallet of a second user (Fabien) inquiring about that second user's current location. The main steps are as follows:

1. *Asserting the query's context*: As a first step, facts about the context of the query are asserted – namely they are loaded into the e-Wallet's inference engine for possible use as part of inferences to be made in processing the query. In our example, one such assertion is that “the sender of the query is Norman”.
2. *Asserting elementary information needs and the need to go through an authorization process*: Here the query is translated into an aggregate goal that includes (a) a combination of elementary information needs (or information retrieval goals) – in our example the need to find “Fabien's location”, along with (b) a requirement to go through an authorization process. The authorization process, which is distributed across some of the following steps, results in the request being either denied or cleared, the latter possibly following the application of obfuscation rules. In our example, the authorization goal requires checking that Norman is entitled to having access to Fabien's location and that the level of resolution at which the query is answered is compatible with Fabien's privacy preferences.
3. *Pre-checking whether the query is allowable*: A first check is performed to see whether the query is allowable based on access rights considerations. In our example, the e-Wallet checks whether Norman is allowed to inquire about Fabien's location. Fabien's e-Wallet might include a privacy preference specifying that his colleagues at work can see the building that he is in, when he is on campus, but that no one else should be given access to his location. In this first check, the e-Wallet might be able to determine that Norman is indeed a colleague of Fabien's – e.g. based on organizational knowledge stored in the static knowledge base of Fabien's e-Wallet. At this stage, because it has not yet determined whether Fabien is on campus or not, the e-Wallet has no ground for denying the request. Therefore, it continues processing it, as detailed below.
4. *Checking the e-Wallet's local knowledge base*: Some queries can be answered in whole or in part, using facts in the e-Wallet's local knowledge base, which, as we have seen in Section 3, contains both static (namely, context-independent) and dynamic (namely, context-sensitive) knowledge about the user. In our particular example, such knowledge is not particularly helpful and the e-Wallet needs to turn to outside sources of personal information to answer the query (see next step).
5. *Invoking personal resources as Web services*: When local knowledge is not sufficient to answer a query, the e-Wallet turns to its service invocation rules to identify external resources that might help answer it. This may involve accessing one or more of the user's personal resources such as his calendar and/or one or more trusted public services. In our example, the campus where Fabien works has a wireless LAN that supports location tracking. This functionality can be invoked by

the e-Wallet to obtain Fabien’s location. The actual invocation takes place through the web service invocation toolkit already introduced in Figure 1.

6. *Post-checking whether the query is allowable.* Armed with additional knowledge obtained by invoking one or more external resources, the e-Wallet is now in a better position to check whether the query is allowable. In our particular example, colleagues of Fabien’s are only allowed to see his location when he is on campus. Assuming that Fabien is on campus, the request is now deemed allowable. This does not mean however that the authorization process required as part of the goals set in step 2 has been fully completed. Obfuscation rules may still need to be applied.
7. *Application of Obfuscation Rules.* Suppose that the location tracking functionality used to answer our query about Fabien’s location returned the specific room he is in, while Fabien is only willing to disclose the buildings that he is in. This latter requirement is captured by the e-Wallet in the form of an obfuscation rule that returns the building in which Fabien is rather than the exact room. Application of this rule will typically involve accessing ontologies about rooms and buildings as well as annotations about the campus where Fabien works.
8. The query has now been fully processed and an acceptable answer generated. This answer (e.g. “Fabien is in Wean Hall”) can be returned to Norman.

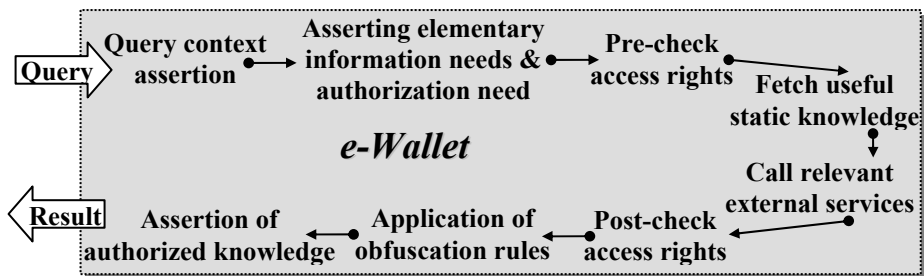


Fig. 2. Main steps involved in processing a query submitted to an e-Wallet.

5 A Three-Layer e-Wallet Implementation

As shown in Figure 3, we developed a three-layer implementation of our e-Wallet:

- *Core Layer:* At the most basic level, the e-Wallet’s knowledge includes an OWL meta-model – required to interpret OWL statements. In addition, it maintains both static (context-independent) and dynamic (context-dependent) knowledge about the user. This knowledge is obtained by loading available annotations about the user along with relevant ontologies and is currently completed using forward chaining reasoning – to avoid having to infer the same facts over and over again. Knowledge in this layer is represented using a (core) triple template:

(predicate, subject, object)_{TRIPLE}

- The *Service Layer* completes the e-Wallet’s core knowledge with invocation rules that map information retrieval goals about contextual attributes onto external service invocations. These are modeled as backward chaining rules. Given an information retrieval goal such as “Give me Fabien’s location”, they help identify and invoke one or more relevant information resources, each modeled as a Web service, as already discussed in Section 4. Knowledge in this layer is represented using a special type of triple called “service triple” denoted:

(predicate, subject, object)_{SERVICE TRIPLE}

Service triples reside in the service layer and are created in either of two ways. They can result from the migration of a triple from the core layer or from the activation of an invocation rule (e.g. an assertion about Fabien’s location as returned by a call to a location tracking service). Migration between the core layer and the service layer is implemented by rules specifying that any (core) triple can be used to generate an equivalent service triple.

- The outer layer is referred to as the *Privacy Layer*, as this is where privacy enforcing rules are applied. Assertions in this layer are represented as another special type of triple called “authorized triple”:

(predicate, subject, object)_{AUTHORIZED TRIPLE}

Only authorized triples can be sent in response to queries. Authorized triples are generated by applying privacy enforcing rules to service triples, thereby ensuring that information about the user is only disclosed to authorized parties and in accordance with relevant obfuscation rules. Privacy enforcing rules are encoded as backward chaining rules. These rules map needs for authorized triples onto needs for service triples to be post-processed subject to the privacy enforcing rules. Upon receiving an incoming query, the e-Wallet generates a need for one or more authorized triples. This need in turn typically triggers needs for service triples and core triples, eventually resulting either (a) in the generation of authorized triples that can be returned in response to the query or (b) in an exception, if the query is found unallowable (e.g. an unauthorized party requesting your location or trying to schedule a meeting in your calendar).




Fig. 3. e-Wallet 3-layer implementation

6 Additional Implementation Considerations

The current implementation of our e-Wallet is based on JESS, a high-performance Java-based rule engine that supports both forward and backward chaining – the latter by reifying "needs for facts" as facts themselves, which in turn trigger forward chaining rules. The e-Wallet's knowledge base is initialized with:

- a model of RDF [28] triples as a template for unordered facts,
- a model of specialized triples used in our three layers (core triples, service triples and authorized triples) along with associated migration rules between the layers, and
- an OWL meta-model.

Additional knowledge is loaded into the e-Wallet by translating OWL input files into JESS assertions and rules, using a set of XSLT stylesheets (figure 4). The OWL input files include ontologies and annotations that are transformed into (core) triple assertions, forward chaining rules (used to complete knowledge at the core layer) as well as service invocation rules and privacy enforcing rules – both represented as backward chaining rules. The XSLT templates act as meta-rules that generate the body, the head and typing (e.g. query transformation stylesheet in figure 5).

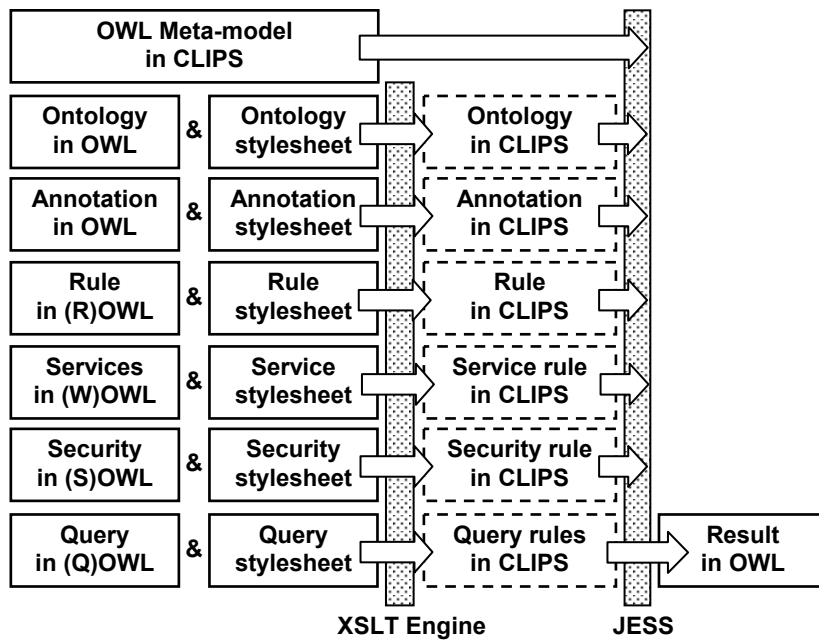


Fig. 4. High-level flows and processes in the e-Wallet

```

(...)
<xsl:template match="/rdf:RDF">
  (defrule query (declare (salience 0))
    <xsl:for-each select="*[not(self::qowl:Query)]">
      <xsl:call-template name="process-class-instance"/>
    </xsl:for-each>
    =>
    (store-result<xsl:call-template name="variable-list"/>)
  )
</xsl:template>

<xsl:template name="process-class-instance" >
  (authorized triple
  (predicate "&rdf;#type")
  (subject <xsl:call-template name="local-ID">
    <xsl:with-param name="id">
      <xsl:value-of select="@rdf:ID"/>
    </xsl:with-param>
    </xsl:call-template>)
  (object <xsl:value-of select="concat(namespace-uri(.),local-name(.))"/>)
  )
  <xsl:for-each select="*">
    <xsl:call-template name="process-property-instance"/>
  </xsl:for-each>
</xsl:template>

<xsl:template name="process-property-instance" >
  <xsl:choose>
    <xsl:when test='count(*)=1'> <!-- has an element for child -->
      <xsl:call-template name="process-objectproperty-instance"/>
    </xsl:when>
    <xsl:when test='count(text())=1'> <!-- has an text for child -->
      <xsl:call-template name="process-dataproperty-instance"/>
    </xsl:when>
    <xsl:when test='@rdf:resource'> <!-- has a reference -->
      <xsl:call-template name="process-referenceproperty-instance"/>
    </xsl:when>
  </xsl:choose>
</xsl:template>
(...)

```

Fig. 5. Extract of the query transformation stylesheet

Once all this knowledge has been loaded and the forward chaining rules have been applied to complete the core knowledge base, the e-Wallet is ready to process incoming queries. A query is transformed into the need for an authorized triple. This in turn triggers privacy enforcing rules and generates needs for service triples. The service triples are generated by either migrating core triples or activating service invocation rules or a combination of both. This is further detailed below:

1. Queries have two components (see Figure 4): (a) an annotation about the query providing its context (e.g., who the sender of the query is), and (b) the query itself in the form of a pattern using a special namespace to identify variables. The context of a query is asserted for the time it takes to process it – later, clean up rules take care of removing all assertions created while processing it. We assume that security protocols (e.g., using digital signatures) are used to verify assertions about the query’s context (e.g. verifying the identity of the sender). The query itself is transformed into a set of authorized triples in the privacy layer. These authorized triples form the body of a backward-chaining rule, whose head is a function that stores the results (in the form of variable instantiations) each time the rule is triggered and that generates RDF results in XML syntax (“pretty printing”).

2. The need of the query for authorized triples triggers privacy enforcing rules. As illustrated in Figure 8, these rules have two roles. First, they check that the sender of the query has the required access rights. In addition, they also apply obfuscation rules to triples to ensure that the level of accuracy or inaccuracy provided in answers to queries is compatible with the user's privacy preferences. The need for authorized triples in combination with privacy enforcing rules generates a need for service triples.
3. The need for service triples in turn triggers service rules. First a generic service rule is applied that checks whether the needed service triple is not already available as a core triple. If this is the case an equivalent service triple is simply created. If there is no equivalent core triple, the e-Wallet looks for matching rules that trigger internal function calls (e.g. getting the current time and date). If that fails too, it looks for matching (external) Web Services. To support this, we have extended the Jess library with internal functions (e.g. time) and functions to call external services. An example of a service invocation rule is given in Figure 9.

The general algorithm can be described in pseudo-code as follows:

```

//Load CLIPS model of RDF/S and OWL
{ $\forall$  triple( $T_i$ ) |  $T_i \in \text{OWL}_{\text{model}}$  } assert ( $T_i$ ) in JESS
{ $\forall$  rule( $R_i$ ) |  $R_i \in \text{OWL}_{\text{model}}$  } defrule ( $R_i$ ) in JESS
//Load ontologies
{ $\forall$   $O_{\text{OWLi}}$  | Ontology ( $O_{\text{OWLi}}$ ) }
   $O_{\text{CLIPSi}} := \text{Ontology stylesheet } (O_{\text{OWLi}})$  // extract ontology triples
  { $\forall$  triple( $T_i$ ) |  $T_i \in O_{\text{CLIPSi}}$  } assert( $T_i$ ) in JESS
//Load annotations
{ $\forall$   $A_{\text{OWLi}}$  | Annotation ( $A_{\text{OWLi}}$ ) }
   $A_{\text{CLIPSi}} := \text{Annotation stylesheet } (A_{\text{OWLi}})$  // extract annotation triples
  { $\forall$  triple( $T_i$ ) |  $T_i \in A_{\text{CLIPSi}}$  } assert( $T_i$ ) in JESS
//Load rules
{ $\forall$   $R_{\text{ROWLi}}$  | Rule ( $R_{\text{ROWLi}}$ ) }
   $R_{\text{CLIPSi}} := \text{Rule stylesheet } (R_{\text{ROWLi}})$  //makes forward rules producing triples
  { $\forall$  rule( $R_i$ ) |  $R_i \in R_{\text{CLIPSi}}$  } defrule ( $R_i$ ) in JESS
//Load service rules
{ $\forall$   $S_{\text{ROWLi}}$  | service description ( $S_{\text{ROWLi}}$ ) }
   $S_{\text{CLIPSi}} := \text{Service stylesheet } (S_{\text{ROWLi}})$ 
  //makes backward rules producing dynamic triples
  { $\forall$  rule( $R_i$ ) |  $R_i \in S_{\text{CLIPSi}}$  } defrule ( $R_i$ ) in JESS
//Load privacy rules
{ $\forall$   $P_{\text{ROWLi}}$  | Security rule ( $P_{\text{ROWLi}}$ ) }
   $P_{\text{CLIPSi}} := \text{Privacy stylesheet } (P_{\text{ROWLi}})$ 
  // makes backward rules producing authorized triples
  { $\forall$  rule( $R_i$ ) |  $R_i \in P_{\text{CLIPSi}}$  } defrule ( $R_i$ ) in JESS
//Load query
 $C_{\text{CLIPSi}} := \text{Query stylesheet } (A_{\text{OWLi}})$  // extract context triples e.g.: sender
{ $\forall$  triple( $T_i$ ) |  $T_i \in C_{\text{CLIPSi}}$  } assert( $T_i$ ) in JESS
 $Q_{\text{CLIPSi}} := \text{Query stylesheet } (Q_{\text{POWLi}})$ 
  // makes one backward rule requiring authorized triples
defrule ( $Q_{\text{CLIPSi}}$ ) in JESS
(Run the RETE Algorithm)

```

Fig. 6. Generic algorithm of the e-Wallet

The system has been designed with efficiency in mind. For instance, query processing stops as soon as a security check violation is identified. Also, at the service layer, rules ensure that the system first checks for available core triples before attempting to invoke external resources.

As a final note, we currently use RDF-S/OWL to represent rules. We do not reify the role of the relation and its arguments. We simply use triples to represent rules and take advantage of the typing mechanism of the XML syntax. We use a special namespace to identify variables. While we are following ongoing developments in RuleML [2] and DAML Rules, our current focus is on the use of rules that apply to OWL assertions. Later we could easily extend our system, for instance using XSLT stylesheets [27] to translate between our representation of rules in OWL and RuleML representations.

As shown in Figure 8, *privacy enforcing rules* are defined using three tags: the content of the *target tag* describes the piece of knowledge to which this rule applies; the content of the *check tag* describes the conditions under which the read access is granted; the content of the *revision tag* describes the obfuscation to be applied before migrating triples to the authorized layer.

As shown in Figure 9 the *service rules* have three children tags: the content of the *output tag* describes the piece of knowledge that this rule can produce; the content of the *precondition tag* describes the knowledge needed for calling the service; the content of the *call tag* describes the function to trigger and its parameters.

```
<qowl:Query rdf:ID="">
  <qowl:sender rdf:resource="http://cs.cmu.edu/~nsadeh"/>
</qowl:Query>
<mc:Person rdf:ID="http://cs.cmu.edu/~fgandon">
  <mc:location rdf:resource="http://sadehlab.cs.cmu.edu/Variable#location" />
</mc:Person>
```

Fig. 7. Query issued by the user 'nsadeh' requesting the location of user 'fgandon'

```
<sowl:ReadAccessRule>
  <rdf:label>people can only know whether or not I am on campus</rdf:label>
  <sowl:target>
    <mc:Person rdf:ID="&variable;#owner">
      <mc:location rdf:resource="&variable;#location"/>
    </mc:Person>
  </sowl:target>
  <sowl:check>
    <rowl:And>
      <rowl:condition>
        <mc:E-Wallet rdf:ID="&variable;#e-Wallet">
          <mc:owner>
            <mc:Person rdf:ID="&variable;#owner"/>
          </mc:owner>
        </mc:E-Wallet>
      </rowl:condition>
      <rowl:condition>
        <mc:Place rdf:ID="http://www.cmu.edu">
          <mc:include rdf:resource="&variable;#location" />
        </mc:Place>
      </rowl:condition>
      <rowl:not-condition>
        <qowl:Query rdf:ID="&variable;#query">
          <qowl:sender rdf:resource="&variable;#owner" />
        </qowl:Query>
      </rowl:not-condition>
    </rowl:And>
  </sowl:check>
  <sowl:revision>
    <mc:Person rdf:ID="&variable;#owner">
      <mc:location rdf:resource="http://www.cmu.edu"/>
    </mc:Person>
  </sowl:revision>
</sowl:ReadAccessRule>
```

Fig. 8. Privacy rule obfuscating the location of the owner

```

<owl:ServiceRule owl:saliience="50">
  <rdf:label>provide location for IP Address</rdf:label>
  <owl:output>
    <mc:Entity rdf:ID="&variable;#entity">
      <mc:location rdf:resource="&variable;#location" />
    </mc:Entity>
  </owl:output>
  <owl:precondition>
    <mc:Entity rdf:ID="&variable;#entity">
      <mc:ip>&variable;#ip</mc:ip>
    </mc:Entity>
  </owl:precondition>
  <owl:call>
    <owl:Service owl:name="call-web-service">
      <owl:parameter>http://cmu.edu/location_tracking#</owl:parameter>
      <owl:parameter>&variable;#ip</owl:parameter>
    </owl:Service>
  </owl:call>
</owl:ServiceRule>

```

Fig. 9. Service rule for location-tracking invocation

7 Static Knowledge and Domain-Specific Rules

As indicated earlier, the RDF triple meta-model is defined as a template used in forward chaining rules. The OWL meta-model is asserted as a list of unordered facts such as the one shown in Figure 10. The semantics attached to properties is translated into rules such as the one shown in Figure 11.

```

(triple
  (predicate "http://www.w3.org/1999/02/22-rdf-syntax-ns#type")
  (subject "http://www.w3.org/2002/07/owl#equivalentProperty")
  (object "http://www.w3.org/2002/07/owl#SymmetricProperty")
)

```

Fig. 10. Declare property equivalence as a symmetric property

```

(defrule equivalent-property (declare (saliience 100))
  (triple
    (predicate "http://www.w3.org/2002/07/owl#equivalentProperty")
    (subject ?p1)
    (object ?p2))
  (triple (predicate p1?) (subject ?s) (object ?o))
  (not (triple (predicate p2?) (subject ?s) (object ?o)))
  =>
  (assert (triple (predicate p2?) (subject ?s) (object ?o)))
)

```

Fig. 11. Rule for forward-chaining completion of property equivalence

As far as the OWL meta-model is concerned, we are focusing on those aspects of OWL-Lite relevant to our application scenarios. More precisely the current system handles: Resource, Class, Property, type, subclassOf, subPropertyOf, ObjectProperty, TransitiveProperty, SymmetricProperty, inverseProperty, equivalentProperty, equivalentClass, sameIndividualAs, DatatypeProperty, FunctionalProperty, InverseFunctionalProperty.

Likewise, triples in the ontologies and annotations loaded into the e-Wallet are asserted as unordered facts. Often, additional domain-dependent rules are needed such as the one in Figure 12. Such rules can help capture context-sensitive knowledge such as ‘Members of my team can see my location when I am at work’.

```

<rowl:Rule direction="forward">
  <rdf:label>A manager is included in the team she manages</rdf:label>
  <rowl:head>
    <mc:Team rdf:ID="&variable;#group">
      <mc:include rdf:resource="&variable;#manager"/>
    </mc:Team>
  </rowl:head>
  <rowl:body>
    <mc:Person rdf:ID="&variable;#manager">
      <mc:manage rdf:resource="&variable;#group"/>
    </mc:Person>
  </rowl:body>
</rowl:Rule>

(defrule A-manager-is-included-in-the-team-it-manages
  (triple
   (predicate "http://www.w3.org/1999/02/22-rdf-syntax-ns#type")
   (subject ?manager)
   (object "http://sadehlab.cs.cmu.edu/mycampus#Person")
  )
  (triple
   (predicate "http://sadehlab.cs.cmu.edu/mycampus#manage")
   (subject ?manager)
   (object ?group)
  )
  (triple
   (predicate "http://www.w3.org/1999/02/22-rdf-syntax-ns#type")
   (subject ?group)
   (object "http://sadehlab.cs.cmu.edu/mycampus#Team")
  )
  =>
  (assert
   (triple
    (predicate "http://sadehlab.cs.cmu.edu/mycampus#include")
    (subject ?group)
    (object ?manager)
   )
  )
)
)

```

Fig. 12. Rule ensuring that managers are included as members of their teams

The inference engine is used to complete the base applying all the rules, thus saving time during the query solving process and providing a rollback point if needed.

8 Service Layer Processing

As indicated earlier, needs for service triples can be satisfied by either migrating a matching core triple or by activating a matching service invocation rule. For obvious efficiency reasons, it makes sense to always look for core triples first. This can be enforced by assigning a high priority (salience in JESS terminology) to rules that look for matching core triples and lower priority to service invocation rules. Service invocation rules are themselves given different priorities, based on the nature of the resource they invoke. This is further detailed below:

- If the needed service triple can be obtained by invoking an internal function (e.g. getting the current time), that function will be activated;
- If no internal service can provide the triple (or if internal calls have failed) but there is a personal service that can possibly provide the needed triple (e.g., obtaining the user's current activity from her personal calendar), the corresponding backward invocation rule is fired, calling that personal resource's Web Service wrapper;

- If no personal resource can provide the needed triple or if the calls failed, the engine looks for invocation rules involving public web services. This can include invoking public Semantic Web search engines (e.g. CORESE [3] or JENA [15] or distributed search architecture such as TAP [12]) or public matchmaking services such as the one in [19] – this step is not currently implemented.
- If everything fails, the query is considered to have failed.

In summary, the body of each rule requires a need for a particular piece of information or triple (e.g. Fabien’s location) along with the availability of a specific set of arguments (e.g. knowledge of the IP address of Fabien’s PDA). When these conditions are matched, the rule fires and calls the service (figure 13 illustrates the semantic web service for location-tracking through CMU’s wireless network).

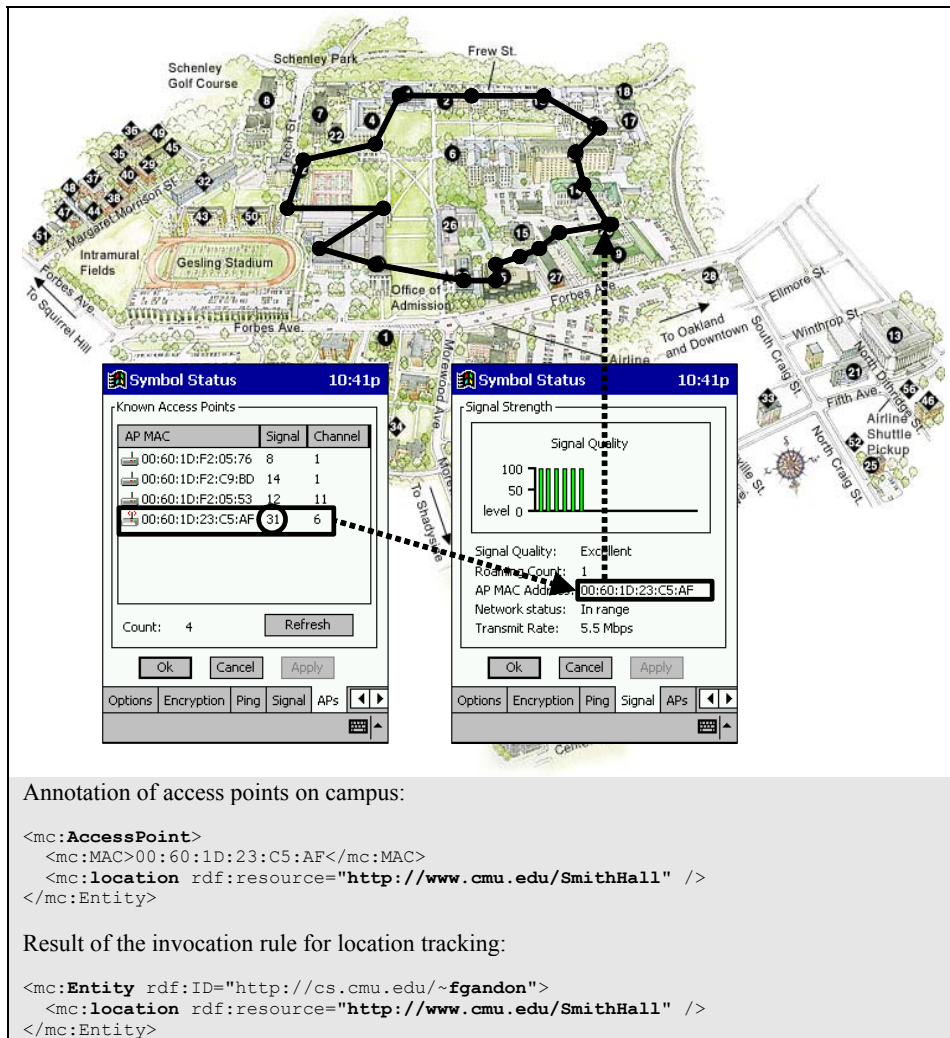


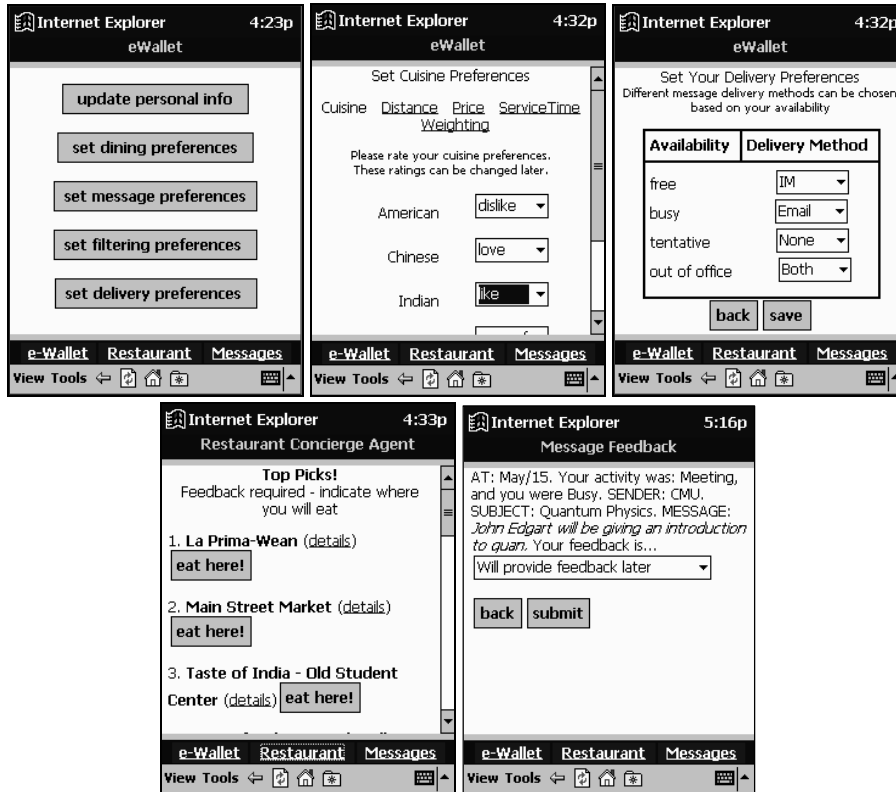
Fig. 13. Semantic web service for location-tracking through CMU’s wireless network

The salience of the rules is used to decide the order to follow when trying to obtain the information (e.g. if multiple sources of location information are available). In general, we envision having a set of rules, where, should everything else fail, the e-Wallet reverts to a low salience rule that invokes one or more semantic search engines (e.g. CORESE [3] or JENA [15] or distributed search architecture such as TAP [12]) and/or one or more public matchmaking service such as [19]. Clearly, as users acquire new personal resources (e.g. a new calendar), they will have to register them with their e-Wallets (e.g. using predefined service profiles that are provided with the resource itself).

9 Summary and Concluding Remarks

In this paper, we have presented a Semantic Web architecture for context-awareness and privacy. A key element of our architecture is its e-wallet, which supports the automated discovery and access of a user's personal resources subject to user-specified privacy preferences. Personal and public resources are represented as web services. Service invocation rules along with service ontologies and service profiles enable the e-Wallet to dynamically identify the most promising resources available to answer a query. When one resource is unavailable, service invocation rules can help identify the next most relevant resource (e.g. using a calendar resource instead of a location tracking resource to estimate the user's current location). Public matchmaking services (e.g. [19]) and semantic search engine functionality (e.g. [3], [15] or [12]) can also be leveraged through low salience rules that amount to reverting to these services when everything else has failed. Another innovation introduced in our e-Wallet is its support for a rich set of privacy preferences, including obfuscation rules that enable users to selectively adjust the accuracy or inaccuracy of responses they provide depending on the context of each query. We have described a three-layer implementation of our e-Wallet using JESS, OWL-Lite and XSLT stylesheets. A query to the eWallet successively results in the creation of needs for authorization triples and service triples. The latter can be satisfied through the identification of matching core triples and/or the activation of service invocation rules.

An early version of our architecture has been validated in *myCampus*, a context-aware environment aimed at enhancing campus life at CMU. The environment is accessible to members of the campus community from their PDAs over the university's wireless LAN. An example of a *myCampus* agent we have developed is a "restaurant concierge" that gives users suggestions on where to have lunch, depending on their food preferences, their location on campus and the weather. For instance, when it rains, the concierge might look for places that do not require walking outside – depending on how the user sets her preferences. Another task-specific agent that has proved particularly popular among students during the course of a 3-day study is a context-aware message filtering agent. The agent filters incoming alerts, taking into account a profile of topics a user is interested in as well as contextual attributes such as the user's current activities (e.g. "when in class, only show me emergency alerts"). Screenshots of both agents are shown in figure 14.



Ex. 14. Screenshot of the e-Wallet (top), a recommendation of the Restaurant Concierge Agent (bottom left) and the feedback for the Message Filtering Agent (bottom right)

Experiments with *myCampus* indicate that different students are interested in different task-specific agents and that, to be effective, many agents require access to a great variety of contextual resources. Our experiments also confirmed that users are concerned about protecting access to their personal information. The need for leveraging a variety of contextual attributes and the students' demand for privacy strongly argue for Semantic e-Wallets such as the one presented here. A key challenge however remains to reconcile the power of these Semantic environments with all important usability requirements that demand systems that are flexible, yet easy to configure. We are experimenting with different approaches to editing and learning user profiles, which we hope will help alleviate this problem.

Acknowledgments. This material is based on research conducted as part of the DAML initiative and has been sponsored by the Air Force Research Laboratory under contract F30602-02-2-0035 and, earlier, by the Defense Advanced Research Project Agency under contract F30602-98-2-0135. The US Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. This work was also in part supported by grants from IBM, HP, Symbol, Boeing and the IST Program (SWAP project).

References

1. Berners-Lee, T., Hendler, J., Lassila, O.: The Semantic Web, *Scientific American*, May (2001)
2. Boley, H., Grosz, B., Tabet, S., Wagner, G.: RuleML DTDs, The Rule Markup Initiative RuleML, <http://www.dfki.uni-kl.de/ruleml/indtd0.8.html>
3. Corby, O., Faron-Zucker, C., Corese: A Corporate Semantic Web Engine, Workshop on Real World RDF and Semantic Web Applications, WWW Conference, Hawaii, (2002)
4. DAML Joint Committee: DAML+OIL language, 27 March 2001, <http://www.daml.org/2001/03/daml+oil-index.html>
5. DAML Services Coalition: DAML-S: Web Service Description for the Semantic Web, First International Semantic Web Conference, ISWC'02, Sardinia, Italy, LNCS 2342, (2002) 348-363
6. Dertouzos, M.: The Future of Computing, *Scientific American*, August (1999)
7. Dey, A.K., Abowd, G.D.: Toward a Better Understanding of Context and Context-Awareness, GVU Technical Report GIT-GVU-99-22. College of Computing, Georgia Institute of Technology, (1999)
8. Dey, A., Salber, D., Futakawa, M., Abowd, G.: An Architecture to Support Context Aware Computing, Technical Report GIT-GVU-99-23. College Computing, Georgia Institute of Technology, Nov. (2000)
9. FIPA, Specifications (2002) <http://www.fipa.org/repository/fipa2000.html>
10. Friedman-Hill, E.: Jess in Action: Java Rule-based Systems, Manning Publications Com-pany, June 2003, ISBN 1930110898, <http://herzberg.ca.sandia.gov/jess/>
11. Garlan, D., Siewiorek, D., Smailagic, A., Steenkiste P.: Project Aura: Towards Distraction-Free Pervasive Computing, *IEEE Pervasive Computing*, Special Issue on Integrated Pervasive Computing Environments, Vol. 1, Number 2, April-June (2002) 22-31
12. Guha, R., McCool, R., Miller E.: Semantic Search, To be published in proceedings of WWW Conference, Budapest (2003)
13. Hendler, J.: Agents on the Web, *IEEE Intelligent Systems*, Special Issue on the Semantic Web, Vol. 16, No. 2, pp. 30-37, March/April, (2001)
14. Hong, J., Llanday, J.: A Context/Communication Information Agent, in *Personal and Ubiquitous Computing Special Issue Situated Interaction and Context-Aware Computing*, Vol. 5(1) (2001) 78-81
15. McBride, B.: Jena: Implementing the RDF Model and Syntax Specification, in *Proc. International Workshop Semantic Web*, (2001) <http://www-uk.hpl.hp.com/people/bwm/papers/20001221-paper/>
16. OASIS: Universal Description, Discovery and Integration standard, <http://www.uddi.org/>
17. OASIS: Security Assertion Markup Language (SAML), Technology Reports, April 14 (2003) <http://xml.coverpages.org/saml.html>
18. OASIS: Extensible Access Control Markup Language (XACML), Technology Reports, March 28 (2003) <http://xml.coverpages.org/xacml.html>
19. Paolucci, M., Kawamura, T., Payne, T., and Sycara, K.: Semantic Matching of Web Service Capabilities, *Proceedings of the 1st International Semantic Web Conference*, Eds. I. Horrocks and J. Hendler, LNCS 2342, Springer Verlag, 2002.
20. Sadeh, N., Gandon, F., Chan, T., Van, L., Kwon O., Takizawa, K.: Context-Aware m-Commerce: A Semantic Web Approach, *Proceedings of the 5th International Conference on Electronic Commerce*, Pittsburgh, PA, (2003)
21. Schilit, W.: A System Architecture for Context-Aware Mobile Computing, Ph.D. Thesis, Columbia University, 1995.
22. Schilit, B., Adams, N., Want, R.: Context-Aware Computing Applications. *Proc. of the Workshop on Mobile Computing Systems and Applications*, IEEE Computer Society, Santa Cruz, CA, (1994) 85-90
23. Want, R., Hopper, A., Falcao, V., Gibbons, J.: The Active Badge Location System, *ACM Transactions on Information Systems* 10(1) (1992) 91-102.
24. W3C: OWL Web Ontology Language Reference, Working Draft 31 March 2003, <http://www.w3.org/TR/owl-ref/>
25. W3C: Web Services Description Language (WSDL) 1.1, Note 15 March 2001 <http://www.w3.org/TR/wsdl>
26. W3C: The Platform for Privacy Preferences 1.0 (P3P1.0) Specification, Recommendation 16 April 2002, <http://www.w3.org/TR/P3P/>
27. W3C: XSL Transformations (XSLT) Version 1.0, Recommendation 16 November 1999, <http://www.w3.org/TR/xslt>
28. W3C: RDF Vocabulary Description Language 1.0: RDF Schema, Working Draft 23 January (2003) <http://www.w3.org/TR/rdf-schema/>