The Knowledge Engineering Review, Vol. 18:3, 209–220.
DOI: 10.1017/S0269888904000037 Printed in the United Kingdom

Use of ontologies in a pervasive computing environment

ANAND RANGANATHAN, ROBERT E McGRATH, ROY H CAMPBELL and M DENNIS MICKUNAS

Department of Computer Science, University of Illinois, Urbana-Champaign, USA; e-mail: ranganat@uiuc.edu, mcgrath@cs.uiuc.edu, rhc@uiuc.edu, mickunas@cs.uiuc.edu

Abstract

Ontologies are entering widespread use in many areas such as knowledge and content management, electronic commerce and the Semantic Web. In this paper we show how the use of ontologies has helped us overcome some important problems in the development of pervasive computing environments. We have integrated ontologies and Semantic Web technology into our pervasive computing infrastructure. Our investigations have shown that Semantic Web technology can be integrated into our CORBA-based infrastructure to augment several important services. This work suggests a number of requirements for future research in the development of ontologies, reasoners, languages and interfaces.

1 Introduction

There is growing evidence for the potential value of Semantic Web technology for Web services and other open, distributed systems (Goble *et al.*, 2002; Peer, 2002). This paper presents a case study of the use of Semantic Web technology in a pervasive (or ubiquitous) computing environment, GAIA (Roman *et al.*, 2002).

Pervasive (or ubiquitous) computing environments are physical environments saturated with computing and communication, yet gracefully integrated with human users (Lyytinen *et al.*, 2002). These environments involve the construction of massively distributed computing systems that feature a large number of autonomous entities (or agents). These entities could be devices, applications, services, databases, users or other kinds of agent. Various types of middleware (based on CORBA, Java RMI, SOAP, etc.) have been developed that enable communication between different entities. However, existing middleware has no facilities to ease semantic interoperability between the different entities.

Ontologies have been widely used in many areas such as knowledge and content management, electronic commerce and the semantic web. In this paper we show how the use of ontologies has helped us overcome some of the challenges in constructing and managing a pervasive computing environment. Of course, these problems are not unique to pervasive computing, but are faced by any multi-agent software system. We believe that our solutions to some of these issues can be extended to any multi-agent system.

This work has considered three major issues that confront the development and deployment of pervasive computing environments. These are:

- discovery and matchmaking,
- interoperability between different entities and
- context-awareness

This section briefly describes these three tasks in the domain of a pervasive computing environment.

In the future, we will extend this work to augment the configuration and management of multiple spaces, and to augment additional services, such as the quality of service infrastructure (Wichadakul *et al.*, 2002).

1.1 Discovery and matchmaking in a pervasive computing environment

Pervasive environments are open systems in which the entities are heterogeneous and autonomous. New entities can enter the environments at any time. In this environment, one of the key problems is *discovery* or *matchmaking* (Trastour *et al.*, 2001): entities must be able to locate services they require without prior knowledge of the configuration of the environment. Matchmaking seeks to find sets of candidates that might fulfil the required and desired service request, based on the request, available entities and services, and other criteria, such as quality of service. This is the first step to composing chains of services to accomplish a task.

In a massively distributed environment with a large number of autonomous entities, this is a significant challenge that is not met by conventional systems. Advertisers and requesters could have very different perspectives and knowledge about the same service and could use different terms to describe the same concept. Furthermore, the system is constantly evolving, so it is difficult to ensure that all entities have a complete and up-to-date ''dictionary'' of the system.

Conventional object registries and discovery protocols, such as the CORBA naming service, LDAP, UDDI, Salutation or JINI discovery service provide limited capability for object discovery. These protocols have several limitations, including:

- match based on syntax and string patterns,
- limited ability to add new types and classes,
- limited mechanisms for defining and checking constraints and
- limited ability to handle incomplete or uncertain information.

For example, consider an environment that contains a variety of network-attached devices, including laser printers, copiers and fax machines (some devices may have several or all of these capabilities). An application may search for ''printer''. Depending on what attributes have been registered for the devices, the search might or might not locate all the devices that could potentially serve the request. For instance, the fax machine may be able to perform printing operations, but it may have registered itself as a ''fax machine'' and not as a ''printer''. Hence it may not be discovered.

1.2 Interoperability between different entities

New entities may enter the environment at any time; these new entities have to interact with existing entities. The interaction must be based on common, well-defined concepts, so that there is no misunderstanding between the entities. The entities must have a common understanding of the various terms and concepts used in the interaction.

For autonomous entities to interact with one another, they need to know, beforehand, what kinds of interface they support and what protocols or commands they understand. In a truly distributed scenario, such as a pervasive computing environment, it may not be reasonable to assume that such agreement exists.

Similar mechanisms are needed for humans to interact with different entities. Humans need to understand what various entities do, and they need to understand the relationships between such entities. It is essential for humans to form an accurate conceptual model of the environment so that they can interact with the environment easily.

1.3 Context-awareness

Applications in pervasive and mobile environments need to be context-aware so that they can adapt themselves to rapidly changing situations. Applications in pervasive environments use different kinds of context (such as location of people, activities of individuals or groups, weather information etc.).

The various types of contextual information that can be used in the environment must be well defined so that different entities have a common understanding of context. Also, there need to be mechanisms for humans to specify how different applications and services should behave in different contexts. These mechanisms need to be based on well-defined structures of different types of context information.

1.4 Ontologies in pervasive computing environments

In order to tackle the problems described above, we apply the technologies from the emerging 'Semantic Web' (Berners-Lee *et al.*, 2001; W3C, 2003a). While the Semantic Web was designed to enhance Web search and agents, we show that it is well suited to some of the requirements of pervasive computing environments.

We have incorporated the use of ontologies in our prototype pervasive computing environment, GAIA (Roman *et al.*, 2002). The ontologies are written in DAML+OIL (daml.org, 2003), describing various parts of the GAIA environment. An ontology server manages a composite ontology that describes the entities of the system and performs operations on the ontologies. This composite ontology is built by composing different ontologies specified in DAML+OIL files. The ontologies are validated into a Knowledge Base (KB), built on the CORBA FaCT Server (Bechhofer *et al.*, 1999; Horrocks *et al.*, 1999).

Ontologies are used for describing various concepts in the GAIA pervasive computing environment. We have developed ontologies that describe the different kinds of entity and their properties. These ontologies define different kinds of application, service, device, user, data source and other entities. They also describe various relations between the different entities and establish axioms on the properties of these entities that must always be satisfied.

A second use of ontologies is to describe different types of contextual information in GAIA. The ontology defines standard descriptions for locations, activities, weather information and other information that may be used by context-aware applications.

The ontologies that describe the pervasive environment greatly help in the smooth operation of the environment. Some of the ways in which we use ontologies in our pervasive environment are:

- Checking to see if the descriptions of different entities are consistent with the axioms defined in the ontology. This also helps ensuring that certain security and safety constraints are met by the environment.
- Enabling semantic discovery of entities.
- Allowing users to gain a better understanding of the environment and how different pieces relate to each other.
- Allowing both humans and automated agents to perform searches on different components easily.
- Allowing both humans and automated agents to interact with different entities easily (say, by sending them various commands).
- Allowing both humans and automated agents to specify rules for context-sensitive behaviour of different entities easily.
- Enabling new entities (which follow different ontologies) to interact with the system easily.

In the following sections, we describe how ontologies are used within our pervasive computing environment, GAIA. We first introduce GAIA; we then describe how the ontology server fits into

the GAIA framework. We then briefly describe the kinds of ontology that have been developed and how they are used within GAIA. Finally, we evaluate our approach and suggest important areas for future research.

2 GAIA: a pervasive computing environment

GAIA is an infrastructure for smart spaces, which are pervasive computing environments that encompass physical spaces (Roman *et al.*, 2002). GAIA converts physical spaces and the devices they contain into a programmable computing system. It offers services to manage and program a space and its associated state. GAIA is similar to traditional operating systems in that it manages the tasks common to all applications built for physical spaces. Each Space is self-contained, but may interact with other Spaces. GAIA provides core services, including events, entity presence (devices, users and services), discovery and naming. GAIA uses CORBA to enable distributed entities to communicate with one another. GAIA has served as our test-bed for the use of ontologies in pervasive computing environments.

We have used GAIA to manage rooms in our computer science building. GAIA helps make these rooms smart and responsive to the needs of different users. There are a wide variety of devices that exist in these rooms. These include authentication devices like fingerprint sensors and smart-card readers, display devices like large plasma screens, video walls, handheld devices, wearable devices like smart watches and smart rings, various input devices such as touch screens and microphones, etc.. In addition, there are large number of applications and services like music-playing applications, presentation applications and drawing applications. Ontologies provide a very easy way to manage this diversity in our environments.

Unlike many uses of ontologies, we have tightly integrated Semantic Web services into the infrastructure. We have implemented an ontology server, which is a standard system service that provides a generic interface to manage DAML+OIL ontologies, a knowledge base and logic queries. Any entity in the system can use the ontology server.

We have created ontologies (in DAML+OIL) to classify and describe many concepts of the pervasive computing environment. The ontologies include entities of the system (not limited to software) and context information.

The ontologies and ontology server are used to augment system services, including:

- configuration management,
- discovery and matchmaking,
- human interfaces,
- interoperation of components and
- context-sensitive behaviour.

The ontologies and knowledge base are tightly integrated into the whole system.

2.1 Kinds of ontologies in GAIA

We use ontologies to describe various parts of our pervasive environment, GAIA. In particular, we have ontologies that have meta-data about the different kinds of entity in our environment. We also have ontologies to describe the different kinds of contextual information in our environment. The ontologies used in GAIA are described in more detail in McGrath *et al.* (2003).

2.1.1 Ontologies for different entities

We have developed ontologies that define the different kinds of entity, provide meta-data about them and describe how they relate to each other. The ontologies augment the informal and implicit taxonomy of the different kinds of entities in the pervasive computing environment.

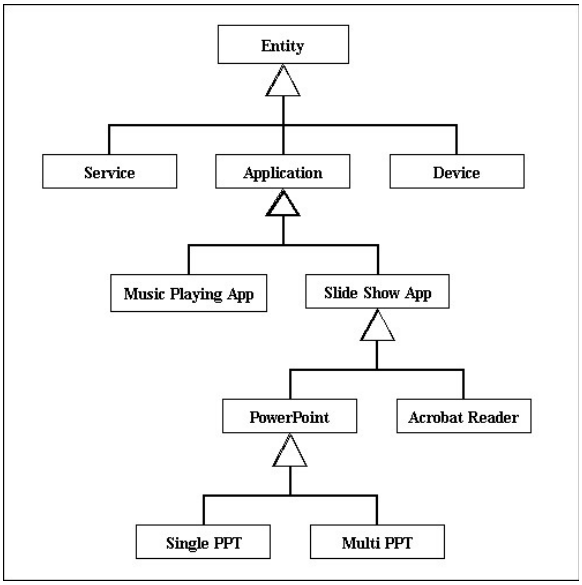

Figure 1 Part of the entity hierarchy of GAIA, with several categories of application

Each application in the space has a DAML+OIL file associated with it that describes its properties. These properties include the commands that are understood by the model of the application (and any parameters required for these commands), the search schema supported by the model (in case it does support searches), the data formats supported by the application, etc. The descriptions of these properties in the ontology ease the task of controlling the application to change its behaviour according to the context.

For example, the DAML+OIL description of an MP3-player application says that it understands commands like start playing a song, stop, pause and change volume. It also supports searches based on fields like artist name, genre etc. It understands and can execute ''.mp3'' files. Finally, it expresses the condition that the ''Presentation'' component of this application must execute in a machine with a speaker attached to it. There are also human-understandable descriptions of the properties and functionalities of the application.

Ontologies also help organising the different applications in hierarchies. Our hierarchy of applications is based on functionality; for example, one part of the hierarchy includes various applications that can play music and another part includes applications that can display slides. This hierarchy helps users easily discover the most appropriate application to their needs. Figure 1 shows a part of this hierarchy.

2.1.2 Ontologies for context information

GAIA has a context infrastructure (Ranganathan *et al.*, 2003) that enables applications to obtain and use different kinds of contexts. This infrastructure consists of sensors that sense various contexts, reasoners that infer new context information from sensed data and applications that make use of context to adapt the way they behave. We use ontologies to describe context information. This ensures that the different entities that use context have the same semantic understanding of contextual information. There are different types of context that can be used by applications. These include physical contexts (location and time), environmental contexts (weather, light and sound levels), informational contexts (stock quotes, sports scores), personal contexts (health, mood, schedule, activity), social contexts (group activity, social relationships, whom one is in a room with), application contexts (e-mail, websites visited) and system contexts (network traffic, status of printers).

We represent contexts as predicates. The use of predicates to describe context automatically lends itself to developing rules and performing reasoning in various forms of logic, for example by

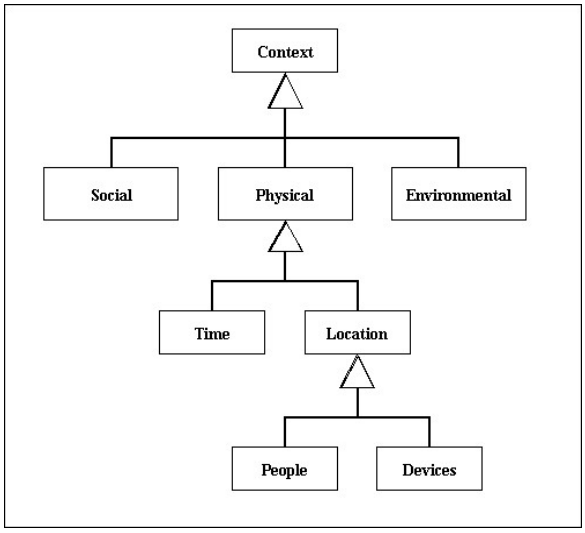

Figure 2 Part of the context hierarchy of GAIA

using Prolog or other reasoning engines. Example context predicates are:

- Location (chris, entering, room 3231)
- Temperature (room 3231 , ''='' , 98 F)
- Sister(venus , serena)
- StockQuote(msft , ''>'' , $60)
- PrinterStatus(srgalw1 printer queue , is , empty)
- Time(New York , ''<'' , 12:00 01/01/01)

The structures of different context predicates are specified in an ontology. Each context type corresponds to a class in the ontology. This ontology defines various context types as well as the arguments that the predicates must have.

For example, many context predicates are defined to have arguments in an SVO (Subject–Verb–Object) format. Thus the structure of these predicates is ContextType(<Subject>,<Verb>, <Object>). For instance, the ontology declares that the Location predicate must have a subject which belongs to the set of persons or things, a verb or preposition like ''inside'' or ''entering'' and a location, which may be a room or a building.

The ontology is used to check the validity of context predicates. It also makes it easier to write different context predicates in rules since we know what the structure of the predicate is and what kinds of value different arguments can take. Since the structure and the semantics of context predicates are well defined in the ontology, it allows different components in the system to have a common understanding of the semantics of different contexts.

Ontologies again help in organising the different types of context information in a hierarchy. Contexts are thus organised in various categories like physical contexts (location, time), environmental contexts (weather, light and sound levels), informational contexts (stock quotes, sports scores), personal contexts (schedule, activity), social contexts (group activity, social relationships, other people in a room), and so on. Figure 2 shows a part of the context hierarchy.

2.2 The ontology infrastructure in GAIA

Figure 3 shows the Ontology Infrastructure in GAIA. The key component in the ontology infrastructure is an ontology server. The ontology server is a CORBA service that maintains a single, cumulative ''current ontology'' for an active space. Each active space has one ontology server running in it. The ontology server implements algorithms to load and validate ontologies from DAML+OIL XML files, compose ontologies into a combined ontology for the entire system

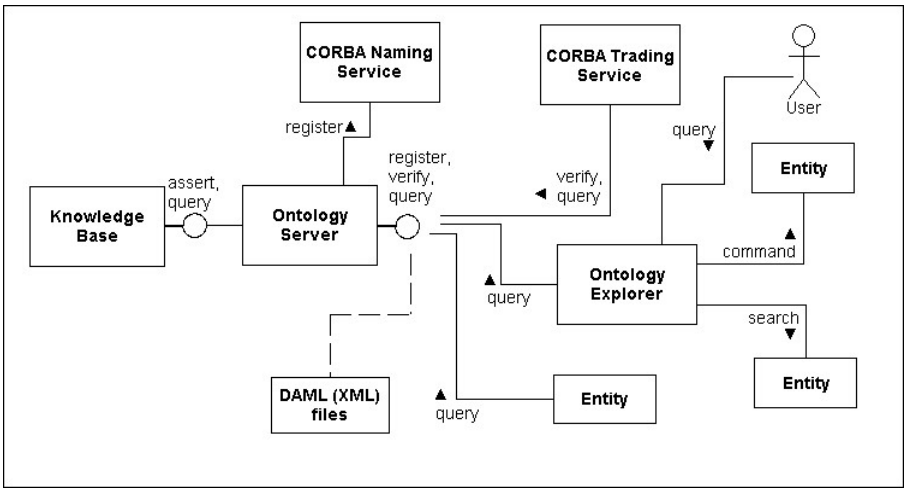

Figure 3 The ontology infrastructure of GAIA

and serve logical queries to a Knowledge Base (KB) representing the dynamically composed ontology base. It registers with the CORBA naming service so that it can be discovered by other entities in the environment. Other entities in GAIA contact the ontology server to retrieve descriptions of entities in the environment or definitions of various terms used in GAIA, or to perform semantic matchmaking.

The current implementation of the ontology server uses the CORBA FaCT server to answer logic queries (*satisfiability*, *subsumption*, and *equivalence*) and validate ontologies. The FaCT server is based on description logic. It is possible to use other reasoning packages as well (e.g. OntoBroker system, based on F-Logic, CLASSIC, OntoMerge or JTP).

A pervasive computing environment is very dynamic. New kinds of entity can be added to the environment at any time. The ontology server allows adding new classes and properties to the existing ontologies at any time. For this, a new ontology describing the new entities is first developed. This new ontology is then added to the shared ontology using bridge concepts that relate classes and properties in the new ontology to existing classes and properties in the shared ontology. These bridge concepts are typically subsumption relations that define the new entity to be a subclass of an existing class of entity. For example, if a new kind of fingerprint recogniser is added to the system, the bridge concept may state that it is a subclass of ''AuthenticationDevices''. We have no way at present of automatically generating these bridge concepts, although that would be very useful.

The KB managed by the ontology server only has class information: the KB has information about the types or classes of different entities or terms. The descriptions of actual instances of entities (i.e. the current real-time state of the system) are managed by other components of the system. For example, we have a component called the space repository that maintains information about the entities in the space at any time. Instances of context information are distributed among different sensors and other entities that use context. This separation of class and instance descriptions helps keep the system modular. The ontologies do not have to be modified when new instances are introduced into the environment (which happens quite often). These new instances, however, are consistent with the definitions of their classes, as described in the ontology.

One of the key benefits in using ontologies is that they aid interaction between users and the environment since they concisely describe the properties of the environment and the various concepts used in the environment. With that aim in mind, we have developed an ontology explorer, which is a graphical user interface that allows users to browse and search the ontologies in the space. The ontology explorer also allows users to interact with other entities in the space through it. This interaction with other entities is governed by their properties as defined in the ontology. The ontology explorer and ontology server are described in more detail in McGrath *et al.* (2003).

2.3 Uses of ontologies in a pervasive computing environment

The ontology server can be used by any application, component or service in the GAIA environment. The ontologies that describe entities and context information are used to enable different parts of the pervasive environment interact with each other easily. The ways in which ontologies are used in GAIA are described in more detail in McGrath *et al.* (2003).

2.3.1 Configuration management

A pervasive computing environment is very dynamic; the configuration must change as activities change and as people and devices enter and leave. Configuration management is very challenging, especially because

- new entities, never before seen, may enter,
- components need to automatically discover and collaborate with other components and
- entities and components are heterogeneous and autonomous.

Without ontologies, the GAIA environment is configured with scripts and ad hoc configuration files (Roman *et al.*, 2002). Ontologies can replace these mechanisms with a standard, formal XML language.

Each entity is associated with an XML file that describes its properties. When a new entity is introduced into the system, its description is checked against the existing ontology to see whether it is satisfiable. If the description is not consistent with the concepts described in the ontology, then either the description is faulty (in which case the owner of the entity or context has to develop a correct description of the entity or context), or there are safety or security issues with the new entity or context. For example, the ontology may dictate that all electrical and electronic devices that are to be introduced in an environment (like a smart room) must accept 110V AC power. In that case, if somebody tries to install a new television that is made for Europe and only takes 220V power, then the description of the new TV would be inconsistent with the ontology and a safety warning may be generated.

Formal ontologies also increase the capability to use descriptions from different, autonomous sources. The DAML+OIL ontologies can be published, to enable autonomous developers and service providers to describe their products with the correct vocabulary. Conversely, autonomous entities can specify the correct formal vocabulary to be used to interpret their descriptions by referring to the relevant DAML+OIL ontology. These actions require more than the URL; the formal semantics defined for DAML+OIL ensures that ontologies from different sources can be used together.

2.3.2 Semantic discovery and matchmaking

In our environment, the ontology server performs the tasks of semantic discovery and matchmaking. It poses logical queries involving subsumption and classification of concepts to the FaCT Server (Bechhofer *et al.*, 1999; Horrocks *et al.*, 1999), which has knowledge of all concepts used in the environment. Such queries are useful in finding appropriate matches. Other entities in the environment query the ontology server to discover classes of components that meet their requirements.

Matchmaking uses the information from the ontology to determine a set of concepts that fulfils the intersection of the requirements of two or more parties, such as a supplier and a consumer [27, 33]. We use a semantic matchmaking algorithm described by Trastour *et al.* (2001). The algorithm defines a match of a query (service request) with a service (advertisement) using logic operations on the two concepts (C1, C2). C1 matches C2 if:

- C1 is equivalent to C2, or
- C1 is a sub-concept of C2, or
- C1 is a super-concept of a concept subsumed by C2, or
- C1 is a sub-concept of a direct super-concept of C2 whose intersection with C2 is satisfiable

The result is a set of classes that are semantically similar to the query class (and not just syntactically similar, as most current discovery algorithms find).

2.3.3 Improved human interfaces

Ontologies can be used to make better user interfaces and allow these environments to interact with humans in a more intelligent way. Ontologies describe different parts of the system, the various terms used and how various parts interact with each other. All classes and properties in the ontology also have documentation that describes them in greater detail in user-understandable language. Ontologies enable semantic interoperability between users and the system.

For example, we have defined the term ''meeting'' as a subclass of ''GroupActivity''. A meeting is defined to have a location, a time, an agenda (optional) and a set of participants. It also has the following human-understandable comment:

''A meeting is an activity that is performed by a group of people. A meeting involves different people coming together at a particular time or place with a common purpose in mind''.

Thus both humans and automated entities in the environment can get a clear understanding of the term ''meeting'' by looking it up in the ontology.

We have developed a GUI called the ontology explorer that allows users to browse the ontology describing the environment. A user can search for different classes in the ontology. He can then browse the results – for example, he can get documentation about the classes returned, get properties of the class, etc. The ontology explorer also lets him interact with the entities by sending commands or performing searches as described in the following subsection. The ontology explorer is similar to a class browser, but it may be used to browse information about all concepts in the system (like context information, applications, services, terms), not just the software objects.

2.3.4 Improved interoperability between entities

The description of the properties of different classes of entity thus allows both users and other automated agents to interact with them more easily by performing searches on them or sending them various commands. This has proved to be one of the major advantages to using ontologies in a pervasive computing environment since it helps simplify the user's and the agent's interaction with such complex systems.

Entities that support searches have their schemas described in the ontology. The ontology also specifies which fields in the query are required and which are optional. Thus any other entity (including users) can browse the ontology to learn the schema and query formats supported by the searchable entity. They can then frame their query and get the results. For example, we have an MP3 server that exposes a query interface for searching for MP3 files. The schema for querying contains fields like artist name, length of song etc. This schema is described in the ontology. Other entities thus know how to query the MP3 server.

The same idea is used to let entities interact with one another, i.e. by sending commands. Different entities allow different types of action to be performed on them. For example, the MP3 server described above allows different commands to be sent to it – start, stop, pause, change volume etc. In our framework, entities specify the commands they support and the parameters of these commands in an ontology. Other entities can thus send commands to these entities.

These concepts are also used in allowing humans to interact easily with the different entities in an environment. The ontology explorer allows humans to send search queries or commands to various entities based on their description in the ontology.

2.3.5 Context-sensitive behaviour

An ontology can improve the robustness and portability of context-aware applications. It is not possible to design in all possible contexts – or even to know what contexts may be used. Ontologies

for context information are an important mechanism for adapting to environments. The application specifies rules for context-sensitive behaviour using a specific set of context concepts and events (a vocabulary). When the application moves to a new space, the context may be different. This might be due to different sensors, different versions of services or different localisations. If the differences are terminological, an ontology may allow the rules to be ''translated'' and then work correctly in the new environment.

Context-aware applications in GAIA have rules that describe what actions should be taken in different contexts. In order to write such a rule, an application developer must know the different kinds of context available as well as possible actions that can be taken by the application. We have ontologies that describe the different kinds of context information – location, time, temperature, activities of people; and also different applications and what commands can be sent to them.

These ontologies greatly simplify the task of writing rules. We have developed a tool that uses these ontologies to allow a developer to write rules easily. The tool allows him to construct conditions out of the various possible types of context available. It then allows him to choose the action to be performed at these contexts from the list of possible commands that can be sent to this application as described in the ontology. Developers can thus very quickly impart context-sensitive behaviour to applications by associating context expressions (involving context predicates) with actions. An example of such a rule for a context-sensitive application is:

IF Location(Manuel, Entering, Room 2401) AND Time(morning) THEN play a rock song.

3 Lessons learned

We have integrated ontologies and Semantic Web technology into our pervasive computing infrastructure. This work suggests a number of requirements for future research and development of ontologies.

Our investigations have shown that the Semantic Web technology can be used with CORBA-based infrastructure to solve some important problems for a pervasive computing environment. Our ontology server provides a standard interface to a knowledge base and logic engine. Ontologies for descriptions of entities and relationships are developed within a knowledge engineering environment and stored as DAML+OIL XML files. Components of the system use the CORBA-based infrastructure to update and query the ontology server. In the future, we will extend this work to augment the configuration and management of multiple spaces, and to augment additional services, such as quality of service infrastructure (Wichadakul *et al.*, 2002).

Our system has integrated semantic services to an unprecedented degree. This was enabled partly by the availability of the CORBA FaCT server (Bechhofer *et al.*, 1999) and the Java classes from OILed (Bechhofer *et al.*, 2001; OilEd, 2002). These packages are a model for what is needed in future software standards:

- A standard API for DAML+OIL (or, more likely, OWL (W3C, 2003b)). The DAML Query Language (DQL) seems to be a promising step in this direction
- A standard interface for generic knowledge base services

Alternative logic engines and knowledge bases should be wrapped in a generic interface, so they can be plugged into infrastructure services. For example, the Open Knowledge Base Connectivity (OKBC) (Chaudhri *et al.*, 1998) could be extended to support DAML (OWL). The Java Theorem Prover (JTP) (Fikes *et al.*, 2003) is a promising step in this direction.

Ontologies and semantic services will play a key role as we develop more sophisticated tools to construct and manage multiple spaces. It will be important to simplify the construction and maintenance of ontologies, perhaps with a repository of standard ontologies. It will also be important to integrate ontologies with the software generation and management, perhaps using

ontologies to semi-automatically generate interfaces. In short, it will be necessary to incorporate successful developments of knowledge engineering environments into the software engineering and configuration management tools of the pervasive computing environment.

A standard upper ontology for services based on, say, DAML-S (Ankolekar *et al.*, 2002) is a good first step. We foresee the need for standard ontologies for many aspects of the pervasive computing environment, including devices, software services, events, people, places and things.

Merging (composing) ontologies from multiple autonomous sources is critical for the pervasive computing environment. In particular, it is necessary to incorporate descriptions of new classes of entity (devices, services, components, and so on) and new types of context information (e.g. new sensors) as they are introduced. It should be possible to develop frameworks and editors to assure that the creators of new entities can create descriptions that can be easily merged into system ontologies. Successful research results in merging ontologies must be implemented in standard services and libraries, in order to be integrated into the infrastructure.

A pervasive computing environment poses significant challenges for the architecture and implementation of reasoners and query engines. DAML+OIL (and in the future OWL (W3C, 2003b)) has proven to be quite useful, especially in combination with a programming interface. However, it seems clear that the DAML and the Description Logic (DL) underlying DAML are necessary but not sufficient for ubiquitous computing applications. Specifically, description logics are not suited for some critical aspects of pervasive computing: DL does not deal well with quantitative concepts, including order, quantity, time or rates. It is also insufficient for spatial reasoning. These kinds of reasoning, however, are essential to certain aspects of ubiquitous computing, including quality of service management (Wichadakul *et al.*, 2002), resource scheduling, and location tracking. Ontologies for pervasive computing environments will require logical models that include spatial and temporal logic, geometry and other quantitative reasoning. It is unclear whether DAML+OIL should be extended to include additional logical concepts, or whether other kinds of markup language should be developed for expressing concepts involving these quantitative aspects. The DAML-Time ontology, currently under development, is an interesting experiment in this area.

The pervasive computing environment is a long-running, open, real-time system. Maintaining an ontology in real time as the system evolves presents important challenges for the design and implementation of ontologies and knowledge bases. In particular, the system needs to address the issues of

- scalability (many thousands of concepts and relations, many hundreds of services using the ontology and KB),
- incremental updates (add, delete or modify a few concepts in a large, active KB),
- persistence and fault-tolerance and
- federation of multiple local knowledge bases.

Another area that requires investigation is security, privacy and access control. The Semantic Web as a whole is largely conceived as a completely open system, in which everything is published for everyone to see. It is far from clear how access control could or should be applied, e.g., to the information in an ontology or a KB. Reasoning engines typically cannot enforce security policies, and the DAML language, for instance, has no facility to limit visibility of concepts or attributes. This topic must be addressed in future research.

Acknowledgements

This research is supported in part by the National Science Foundation grant NSF 98-70736, NSF 9970139 and NSF infrastructure grant NSF EIA 99-72884. Important aspects of this study used software from Iona (IONA Technologies Inc., 2001) and University of Manchester (Bechhofer *et al.*, 1999; OilEd, 2002).

References

Ankolekar, A, Burnstein, M, Hobbs, JR, Lassila, O, Martin, D, McDermott, D, McIlraith, SA, Narayanan, S, Paolucci, M, Payne, T and Sycara, K, 2002, ''DAML-S: web service description for the Semantic Web'' *First International Semantic Web Conference (ISWC)* 348–363.

Bechhofer, S, Horrocks, I, Gable, C and Stevens, R, 2001, ''OilEd: a reason-able ontology editor for the Semantic Web'' *KI2001: Advances in Artificial Intelligence* 396–408.

Bechhofer, S, Horrocks, I, Patel-Schneider, PF, and Tessaris, S, 2001, ''A proposal for a description logic interface'' *International Workshop on Description Logics (DL'99)*.

Berners-Lee, T, Hendler, J and Lassila, O, 2001, ''The Semantic Web'' *Scientific American* **284**(5) 35–43.

Chaudhri, VK, Farquhar, A, Fikes, R, Karp, PD and Rice, JP, 1998, *Open Knowledge Base Connectivity 2.0.3 – Proposed* Stanford Research Institute.

daml.org, 2003, ''The DARPA Agent Markup Language homepage'' available at `http://www.daml.org`.

Fikes, R, Frank, G and Jenkins, J, 2003. *JTP: A System Architecture and Component Library for Hybrid Reasoning* Knowledge Systems Lab KSL-03–01 Stanford University, available at `http://ksl.stanford.edu/KSL_Abstracts/KSL-03-01.html`.

Goble, C and De Roure, D, 2002, ''The grid: an application of the Semantic Web'' *ACM SIGMOD Report* **31**(4) 65–70.

Horrocks, I, Sattler, U and Tobias, S, 1999, ''Practical reasoning for expressive description logics'' *International Conference on Logic for Programming and Automated Reasoning (LPAR'99)* 161–180.

IONA Technologies Inc., 2001, ''ORBacus Trader, Version 2.0.0'' available at `http://www.iona.com/products/orbacus_trader.html`.

Lyytinen, K and Yoo, Y, 2002 ''Issues and challenges in ubiquitous computing'' *Communications of the ACM* **45**(12) 62–65.

McGrath, R, Ranganathan, A, Campbell, R and Mickunas, D, 2003, *Use of Ontologies in Pervasive Computing Environments* UIUCDCS-R-2003-2332 UILU-ENG-2003-1719 Department of Computer Science.

McGrath, RE, 2000, *Discovery and Its Discontents: Discovery Protocols for Ubiquitous Computing* UIUCDCS-R-99-2132 Department of Computer Science University of Illinois Urbana-Champaign.

OilEd, 2002, `http://oiled.man.ac.uk/`.

Peer, J, 2002, ''Bringing together Semantic Web and Web services'' *First International Semantic Web Conference (ISWC)* 279–291.

Ranganathan, A and Campbell, RH, 2003, ''A middleware for context-aware agents in ubiquitous computing environments'' ACM/IFIP/USENIX International Middleware Conference, Rio de Janeiro, Brazil 143–161.

Roman, M, Hess, CK, Cerqueira, R, Ranganathan, A, Campbell, RH and Nahrstedt, K, 2002, ''GAIA: a middleware infrastructure to enable active spaces'' *IEEE Pervasive Computing* **1**(4) 74–83.

Trastour, D, Bartolini, C and Gonzalez-Castillo, J, 2001, *A Semantic Web Approach to Service Description for Matchmaking of Services* HPL-2001-183 HP Laboratories Bristol, available at `http://www.hpl.hp.com/techreports`.

W3C, 2003, ''The Semantic Web'' available at `http://www.w3.org/2001/sw`.

W3C, 2003, *Web Ontology Language (OWL) Reference Version 1.0.* W3C working draft, available at `http://www.w3.org/TR/2002/WD-owl-ref-20021112/`.

Wichadakul, D, Gu, X and Nahrstedt, K, 2002, ''A programming framework for quality-aware ubiquitous multimedia applications'' *ACM Multimedia*, Juan Les Pins.