

Semantically Driven Service Interoperability for Pervasive Computing

Declan O'Sullivan
Knowledge and Data Engineering Group
Trinity College, Dublin
Ireland
+353 1 6083754
Declan.OSullivan@cs.tcd.ie

David Lewis
Knowledge and Data Engineering Group
Trinity College, Dublin
Ireland
+353 1 6082158
Dave.Lewis@cs.tcd.ie

ABSTRACT

The common vision of pervasive computing environments requires a very large range of devices and software components to interoperate seamlessly. From the assumption that these devices and associated software permeate the fabric of everyday life, a massive increase looms in the number of software developers deploying functionality into pervasive computing environments. This poses a very large interoperability problem for which solutions reliant solely on interoperability standards will not scale. An interoperability problem of a similar scale is presented by the desire for a Semantic Web supporting autonomous machine communication over the WWW. Here, solutions based on service-oriented architectures and ontologies are being actively researched, and we examine how such an approach could be used to address pervasive computing's interoperability problem. The paper outlines the potential role that semantic techniques offer in solving some key challenges, including candidate service discovery, intelligent matching, service adaptation and service composition. In particular the paper addresses the resulting requirement of semantic interoperability outlining initial results in dynamic gateway generation. In addition the paper proposes a roadmap identifying the different scenarios in which semantic techniques will contribute to the engineering and operation of pervasive computing systems.

Categories and Subject Descriptors

D.2.12 [Software Engineering]: Interoperability – *data mapping, distributed objects, interface definition languages.*

General Terms

Algorithms, Design, Experimentation, Standardization, Languages, Theory

Keywords

pervasive computing, service composition, semantic interoperability, topic maps, DAML-S.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiDE'03, September 19, 2003, San Diego, California, USA.

Copyright 2003 ACM 1-58113-767-2/03/0009...\$5.00.

1. INTRODUCTION

Weiser's original vision of pervasive computing was driven by the move towards increasingly cheap and miniaturized computers becoming embedded into our physical environments [24]. The benefits are perceived as ubiquitous access to information, untethered communications based on wireless technologies and computer mediated interaction with the environment through sensing, actuating, displaying. This would provide 'intelligent' support for things people want to do through interpreting their needs and task context in order to support their interaction with software and with the immediate physical world.

This vision, however, presents many technical challenges, which are actively being addressed in the research community, including:

- A massive scaling-up of connectivity requirements, which is prompting new network and distributed system architectures – e.g. peer-to-peer.
- A heterogeneity of processor forms, such that Windows and Linux address a shrinking 'PC' niche and must compete with small footprint OSs, e.g. PalmOS.
- Poor application portability over 'embedded' processors, with the larger range of processor types and operating systems increasing software platform dependency and presenting real challenges to Java as a potential solution.
- A heterogeneity of access networks, driven by rapid innovation in cellular and wireless LAN and PAN technologies. Multiple wireless standards are therefore emerging, but cheap terminal solutions for simultaneous access may obviate this problem.

However, the source of the most serious challenges to deploying the pervasive computing vision are not technological but structural. Embedding processors, sensors and actuators in everyday products implies an explosion in the number and type of organizations that need to be involved in achieving the seamless interoperability implied by the pervasive computing vision. Many of the network interoperability problems can be addressed by the inter-networking approach of the Internet's network and transport protocols. However, the potential for debilitating heterogeneity in application level interoperability remains. Consider the complexity involved in reaching agreements on and enforcing conformance to interoperability standards when the players involved expand from the computer and software developers

(Microsoft, IBM, HP and the like) to all the organizations which will embed application into their products, e.g. windows, alarm clocks, pens, automobiles, coffee machines etc to name a few implied by Weiser's original scenario. It is therefore clear that the pervasive computing vision implies a massive increase in scale of the application interoperability problem.

The work presented here is motivated by the observation that we cannot, therefore, rely on shared a priori knowledge via common interoperability standards to solve the application interoperability problems on the scale needed for pervasive computing. Instead, application software must somehow adapt at deployment time and runtime to integrate their functionality and dynamically interoperate with other application software.

The following sections first outline two major evolving technologies that, when combined, show great promise in addressing the interoperability problem involved in engineering pervasive computing systems. The first is service-oriented architectures, which provide a path to dynamically integrating third-party functionality. The second is ontologies, which involve the explicit representation of the semantics of information or functionality in a machine processable form. The paper then goes on to propose how these fields are being integrated in the development of semantic web services. We then outline a roadmap identifying various development-time and run-time challenges in engineering pervasive computing applications and propose that they may be addressed effectively through dynamic semantically-driven interoperability within a service-oriented architecture. Issues related to the engineering of such semantic service are further explored before conclusions are given and future research directions are identified.

2. SERVICE-ORIENTED INTEROPERABILITY

Service-oriented architectures focus on application level interoperability. This is achieved through the definition of well-defined abstract interfaces to distributed software components to allow their rapid integration into the design of a new application. Service-oriented approaches support explicit bindings between the abstract interface definition and different communication protocols or other distributed interaction technologies. This allows the same abstract service definition to be used in several technology domains, often allowing the same core software functionality to be reused also. The combination of an abstract interface definition and its binding to a concrete interaction technology constitutes a service. Service-oriented architectures often support the concept of reflection. This allows an instance of a service to reveal at runtime its service definition, thus supporting run-time binding between services by any agent able to understand and make decisions based on the service definitions.

One of the main concrete examples of an open service-oriented architecture was the OMG's Object Management Architecture, which included CORBA and a number of general purpose and domain specific application services. The OMA supported reflection through its dynamic invocation interface. It also defined the General Inter-ORB Protocol which could be mapped to different transport protocols, the binding most widely used being the Internet Inter-ORB Protocol. CORBA suffered however, from a use of IP that could not be easily deployed over firewalls, thus preventing its effective deployment over the Internet.

The proven inter-domain capabilities of WWW protocols has led to the application of service-oriented architecture over the Web in the form of Web services. This application of a service-oriented architecture is attracting huge research and commercial interest as it is seen as a major key to seamless e-commerce over the Internet. This has resulted in a number of standardization efforts. Most notable is the World Wide Web Consortium's (W3C) standardization of the Web Service Description Language (WSDL)[5]. This is a relatively simple XML language for defining services in terms of operation names and their inputs and output. WSDL services, called port types, can be bound to any suitable communication protocol, though a binding to SOAP over HTTP has been standardizing and is proving the most popular for ecommerce applications over the WWW.

The standardization of WSDL, the strong support for generating WSDL-based services on many application server platforms and the emergence of WSDL-services on the WWW offers the possibility of rapidly developing new services through combinations of existing ones. The process of organizing a set of existing services into a new service is called service composition. Service composition is of interest where services from different organizations may be integrated into some value added service and also for enterprise application integration, where services from different, often distributed, organizational units need to be rapidly integrated to meet changing business needs. The service-oriented view of such integration enables it to be done initially at the abstract interface level and then implemented with a large degree of automation, for example using workflow execution techniques. A wide range of languages for defining such service composition have appeared, e.g. BPEL4WS, WSFL, BPML etc. However, as pointed out in [1], these offerings seem primarily bound to existing platforms and associated commercial interests and don't exploit the formal process semantics offered Petri nets and process algebras. Such languages may therefore prove insufficient for the automated service composition.

2.1 Challenges of Service Composition in Pervasive Computing

The vision of pervasive computing offers many of the same challenges as are already being addressed by dynamic composition of web services. Principally these are:

- Using a common interoperability mechanism to dynamically locate and integrate application functionality from a large number of disparate sources, in this case software embedded in our physical surrounding or running on mobile devices, rather than on the web.
- Defining application interface functionality without detailed prior knowledge of how it will be used.
- Separating the interoperability of application functionality from that of communication mechanisms.

We therefore propose that a service-oriented architecture be adopted for the integration of functionality present in a pervasive computing environment. In other words, all the processors, sensors, actuators, displays etc, present in pervasive computing environments can only interact through services. Therefore, developing the highly adaptive behavior we would like pervasive computing environments to exhibit becomes an ad hoc service composition problem.

This approach potentially benefits from being able to exploit the momentum that has already built behind web services, in the form of standards for services and service composition and their system support in various development tools. A further advantage is that we can seamlessly combine local services and services available over the WWW in supporting a user's requirements.

A service-oriented approach to interoperability in pervasive computing presents some additional challenges. For instance, many pervasive computing devices have limited processing and memory footprints and wireless communication protocols offer differing reliability and quality of service profiles when compared to the Internet, so much of the existing Web Services infrastructure may not be directly applicable to this area. However, for the purposes of this paper we will be focusing on the challenges of automated ad hoc service composition in the pervasive computing domain. In particular we address the problems of application level interoperability and the pressures placed on the service engineering process by the need for semantics. To address this we take the view that the service-oriented approach must be supplemented by application level semantics capable of supporting machine reasoning in interactive timeframes.

3. SEMANTIC INTEROPERABILITY AND ONTOLOGIES

Sheth classifies the types of interoperability problems that can occur during the interchange between information systems as follows [20]:

- System, heterogeneous hardware and operating systems;
- Syntactic, different representation languages and data formats;
- Structural, heterogeneous model representations;
- Semantic, different meaning of terms used in the interchange.

System and Syntactic interoperability problems are more easily dealt with. However, achieving Structural and Semantic interoperability (which together will be termed semantic interoperability in this paper) in information interchange between information systems of different parties continues to be a difficult problem. In order to achieve semantic interoperability in a heterogeneous environment, the meaning of the information that is interchanged has to be understood across the systems. A simple example would be where one party might call the invocation charge for a particular web service "Fee" with the currency type "US Dollars" and the other party calls the charge "Cost" and expects a "Euro" representation.

More generally, Ceri and Widom identify four categories of semantic conflicts [3]:

- Naming conflicts where different names used to represent the same concepts, either homonyms and synonyms .
- Domain conflicts occur when different reference systems are used to measure a value. Examples are different currencies.
- Structural conflicts occur when different systems use different data organization to represent the same concept.
- Metadata conflicts occur when concepts are represented as one type within the modeling type of one system and a

different type within the other system (e.g. as at a schema level in one database and as an instance in another.

In the past, such semantic conflicts have typically been dealt with at "design time": through careful schema design in distributed database solutions; through the hand crafting of interoperability gateways (with system integrator solutions); or by forcing each system to conform to a standard mechanism for interchange (e.g. ebXML). Although these traditional approaches have been successful in well understood/static interchange environments, each of these approaches are inappropriate for systems that want to interchange in dynamic environments [7]: the schema design solution fails due to the rapidly changing nature of the interchanges required; the handcrafting of gateways solution fails as it does not scale to large numbers of information systems; and the standards solution fails due to the lack of certainty as to whether there is a common interpretation of the standard.

Ontologies can be used to describe the semantics of information sources and make the content explicit, and thus can be used to discover semantic equivalence between information concepts. The use of ontologies as a possible solution to the semantic interoperability problem has been studied over the last six or seven years. Wache et al reviewed and categorized twenty five approaches that have been proposed over this period and concluded that reasonable results for integration can be achieved through ontology based approaches [22]. Like the traditional approaches most of the approaches proposed are "design time" solutions towards integration. However as ontologies represent semantics, then semantic interoperability can be achieved by runtime comparison of and inference about ontological information.

The role of ontologies is becoming important in order to realize the vision of the Semantic Web (www.semanticweb.org) and several XML based representation languages are emerging. DAML and DAML+OIL [8], and OWL [24] are based on W3C's Resource Description Framework (RDF) [23]. XML Topic Maps [26] are based on the ISO 13250 standard (which was originally defined to facilitate merging of indexing schemes) All of the above representation languages are seen as candidate web-based XML languages for expressing and interrelating ontologies. Obrst and Liu provide a good insight into the relationship between Knowledge Representation and Ontological Engineering, and the emergence of web based languages [15].

Of course in a dynamic environment such as the one that we envisage, people will have freedom of choice in terms of how they structure their semantics and the representation language that they use to do so. We will term this the ontology heterogeneity problem. Visser et al focuses on ontology heterogeneity and classifies mismatches that may occur under headings of Conceptualization (on how concepts are classified) and Explication (on how concepts are specified). They then go on to discuss how easy/hard it is to deal with each type [21]. More recently Klein has undertaken a similar analysis of the mismatch problems that can occur and examines how various projects propose solutions to these problems [12]. Several papers propose particular approaches or architectures for resolution of ontological heterogeneity, mainly at design time. For example Corcho and Gomez-Perez propose a system that will semi-automatically integrate ontologies of different types in the e-business domain

[6]. With respect to runtime solutions, Campbell et al outline a number of issues that typically need to be addressed [2]:

- How to find a common frame of reference in both ontologies around which algorithms can traverse?
- Identifying what is or is not a match? What constitutes an exact match or a "good enough" match and how is this expressed?
- What online resources that might help identify equivalent terms (e.g. WordNet) should be used and when?

In summary, whereas the use of ontologies has been shown to overcome the semantic interoperability problem between parties that know they want to interact a priori, the use of ontologies to achieve semantically interoperable interactions between parties at runtime in a very dynamic situation is still in its infancy and must overcome a number of challenges.

3.1 How will Ontologies help in Pervasive Computing?

Ontologies could be useful at design time to help pervasive computing engineering, in a similar way to how it has been found useful in the support of system integration at design time. With semantic markup and ontologies, knowledge about various pervasive computing parties (physical or software) can be encoded at creation time and reasoned upon later at reuse or selection time by designers of pervasive computing systems.

The pervasive computing system could also start to reason about relations between users' context and needs, available device capabilities and the physical environment at runtime. It has been argued that ontologies are "overkill" as a solution to support runtime interactions. However, Kim summarizes the situation nicely that in general XML DTDs will be good enough to support runtime interactions in systems where the pressing need is reduction of complexity and common agreement amongst various parties can be achieved, whilst if the pressing need is reduction of uncertainty then ontologies are necessary [11]. The authors would argue that because of the kind of diversity and dynamics that is expected in pervasive computing environments, it is unlikely that there will be common agreement on one representation language or one set of terminology for the markup of all application elements. Therefore each element destined for use in a pervasive computing environment must encode knowledge about itself using explicit reference to an appropriate ontology. Thus, when such elements are integrated at design time or even at runtime for a pervasive computing system, several ontologies may need to be navigated and analyzed in order to dynamically bridge understanding between the constituent elements.

4. SEMANTIC SERVICE INTEGRATION

As the open expression of semantics, in the form of ontologies, supports machine reasoning, there is increasing interest in exploring how they can play a role in supporting increased dynamicity in the composition of services.

The DARPA Agent Mark-up Language initiative is addressing the need for automated service composition by taking an ontology based approach to defining services and service compositions called DAML-S [27]. As with other service composition languages DAML-S is bound to WSDL. However it allows for the definition of inputs and outputs to service operations to be defined

in terms of ontology elements. This allows for a richer expression of the semantics of inputs and output information than is possible with the XML schema information representation used in WSDL. DAML-S also allows for the definition of logical expressions for the pre-conditions to the use of the service and for the effects of using the service, again in terms of ontology elements. In this way the semantics of the service can be expressed in terms of an ontological model of the 'real-world' context in which the service operates. In addition, DAML-S provides some common semantics related to the resources used by a service and how they may be shared. This is an important abstraction in managing the assurance of a service, since the quality of service it is able to deliver will depend on the sharing of such resources. DAML-S thus allows for automated reasoning about the compositions in which the service may be involved without the service designer having to explicitly address the requirements of that particular composition. This is needed in pervasive computing environments where composite service must be formed from several services with no a priori knowledge of each other [4].

DAML-S presents a promising candidate for an open language for the composition of semantic services since it is expressed in the DAML+OIL language which has heavily influenced the W3C standardization of OWL. Future revisions of DAML-S will therefore be built on OWL. However DAML-S is still very much at the research stage, and reasoning application are targeted more at design time tools than at dynamic runtime scenarios [18].

5. ROADMAP

This section outlines a roadmap identifying the evolution of different human roles caused by the impact of semantic service interoperability on the engineering and use of pervasive computing systems. It takes as a starting point the assertion that all the different functionality potentially found within a pervasive computing environment will be expressed as a semantic service in an ontology-based language such as DAML-S.

We consider first the users of pervasive computing. Currently though much software functionality is embedded around them, there is no interoperability leading to personalized, context aware services. In the near term we can expect to see more integrated service offerings, though these will be hand crafted and probably tailored to a particular pervasive computing environment, e.g. the inside of an automobile or an airport lounge. Personalization will be limited to situations where electronic context information is already available, e.g. adapting services to terminal capabilities expressed in CC/PP. In the medium term the spread of semantic service descriptions will make ad hoc service composition more feasible, such that composite services can be dynamically selected or generated to meet a user's specification for the task at hand in a particular pervasive computing environment, e.g. chairing a meeting. As specific ontological models related to tasks and context become better understood and as semantic mismatches are dynamically handled, composite services will be generated for the user through interpretation of gesture and spoken word supported by learnt user preferences and knowledge of the task at hand and of the social setting.

Next we consider the developers of the atomic functionality that is embedded in or accessible from the user's environment. These are the developers of sensors, actuators, processors and displays that will increasingly be incorporated into a wide number of products and the physical infrastructure. The development of such

software functionality and of the accompanying semantic service descriptions will need to be tightly integrated into the manufacturing process of product concerned, albeit a window pane, an automobile or a coffee machine. As the cost of electronics will continue to fall relative to other physical products we expect the core value to remain with the product and its development to remain within the competency of traditional manufacturers. In other word coffee machine manufactures will learn to develop pervasive devices, computer companies will not start developing coffee machines. Initially, exposing a product's pervasive computing capability as semantic serviced will therefore impose some overhead on those charged with programming embedded devices. Developers will need to master current knowledge management tools in order to import and understand the ontologies or to develop new ontological models to accompany their product where no suitable ones exist. This task will be hampered by the lack of integration between such knowledge management tools and their native CAD and CASE tools. This will provide in impetus to integrate knowledge management functionality into these development tools so that ontology browsing, matching and mapping is better integrated into the software engineering and product manufacturing process. Manufacturer will increasingly see the benefits of contributing to industry efforts to agree ontologies for semantic services, e.g. contributing a coffee machine ontology to the kitchen appliance ontology standardization process. Eventually that development and use of semantics becomes tightly integrated and natural part of the develop process.

Lastly we consider the role of the human designer of service compositions. We see this as a specialized role concerned either with designing custom compositions for specific pervasive computing environments, e.g. for corporate customers, or with contributing to open-source composition repositories. Such composition will be in the form of declarative models submitted to service execution environments that orchestrate the invocation of instances of the atomic services involved in the composition. Their differentiating skills of this role will lie in their understanding of semantic descriptions of atomic services and their experience in using knowledge management tools for planning and validating service compositions. They will often have to retro-fit semantic descriptions to required services that have not yet been marked-up. However, the spread of semantic services and semantic interoperability solutions will spur the development of automated service composition agents. Though some specialization of this role may persist in computational hard areas such as optimizing compositions for non-functional requirements, service composition as a human activity will become increasingly redundant. Service composition will either be performed automatically, driven by machine interpretation of user needs or explicit task requests, or will be a skill used by domain experts in fast evolving field where semantics have not yet been adequately captured.

Though elements of this roadmap may be highly speculative, it provides a useful tool in guiding research that spans the two fast moving areas of semantic service composition and pervasive computing. This paper now focuses on two relevant topics:

- Semantic interoperability, i.e. the matching of meaning and its role in automated mediation between run-time systems

- Semantic Service Engineering, i.e. the motivations, process, and challenges of developing atomic and composite semantic services for pervasive computing environments.

6. DYNAMIC AUTO-GENERATION OF SEMANTIC INTEROPERABILITY GATEWAYS

Our group has been investigating semantic web technologies to dynamically bridge between the terminology of two parties from different pervasive computing environments which we assume have been developed using different ontologies [17]. First the ontologies of the two parties are compared. The comparison of the ontologies is then used to derive rule-sets that transform the negotiation in one party's terminology to that of the other party's terminology and vice versa. This comparison and dynamic bridging is undertaken at runtime and requires no a priori knowledge of the parties. Key to the approach is that each ontology need only adhere to a minimum subset of common concepts, with the rest of the ontology being defined to reflect the terminology most natural for the particular pervasive computing environment. This minimum subset of common concepts approach is a variation on the "hybrid ontology approach" outlined in [22]. Of course, deriving what the minimum subset should be, is a major challenge.

An initial prototype has been developed in order to explore the issues involved in the proposed solution. As shown in Figure 1, the initial prototype uses Enterprise Java Beans and XML technologies for its implementation. The scenario chosen involved a person roaming from one pervasive computing environment to another, but where there was no Service Level Agreement in place for that person in the target environment and where the service terminology used across the environments were different. The aim was to develop a system that would allow the negotiation to take place at runtime and real-time such that seamless continuity of appropriate services were maintained. Thus requests in one domain would issue requests/responses in its own service and quality of service terminology. These requests/responses would get translated in the Semantic Interoperability Gateway into requests/responses of the receiving party's terminology and vice versa.

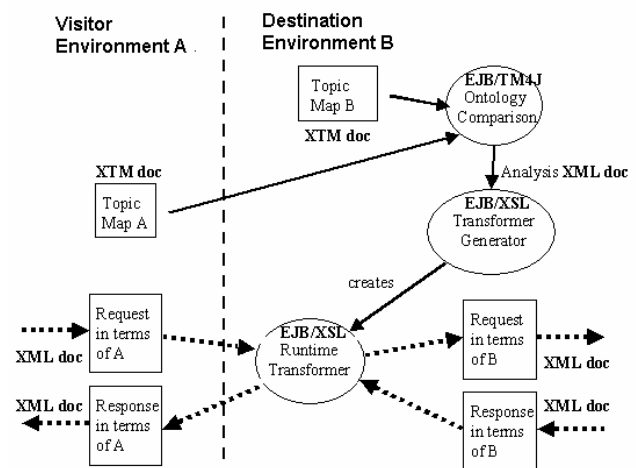


Figure 1. Prototype Auto Generation of SI gateways

The ontologies in the prototype are described using the XML Topic Map (XTM) representation (www.topicmaps.org). A subject, in its most generic sense, can be any ‘thing’ whatsoever — a person, an entity, a concept, really anything — regardless of whether it exists or has any other specific characteristics, about which anything whatsoever may be asserted by any means whatsoever. A topic is created within the topic map to ‘indicate’ (or refer to) this subject. A topic association asserts a relationship between two or more topics. A topic may be linked to one or more information resources that are deemed to be relevant to the topic in some way. Such resources are called occurrences of the topic.

A “Published Subject Indicators” (PSI) is any subject for which a subject indicator has been made available for public use and is accessible online via a URI, and is therefore any resource that has been published in order to provide a positive, unambiguous indication of the identify of a subject for the purpose of facilitating topic map interchange and mergeability. When two topics use the same PSI to indicate their subject, they are by definition “about” the same thing.

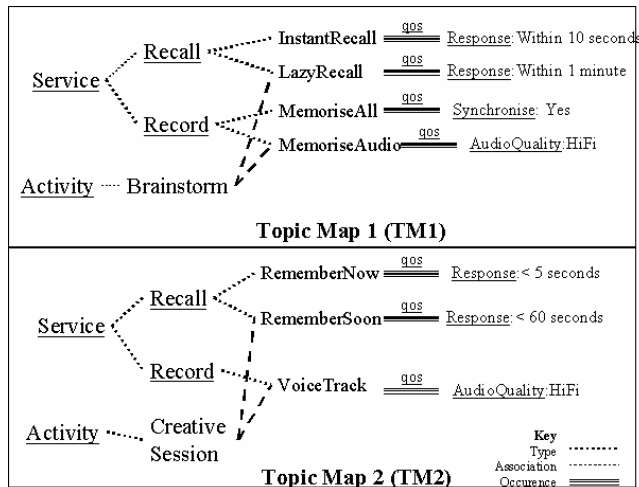


Figure 2. Example Topic Maps

Figure 2 provides a simplified representation of two example topic maps used in the prototype, each describing services of an independent pervasive computing environment. We are assuming that the terms which are underlined in the diagram (e.g. Service, Activity, Recall, Record, QoS etc.) will be subject to common agreement on their semantics. In this case the “Recall” service term would be commonly defined as being a service that “Allows pulling up and display of information of relevance to the activity underway” (e.g. the retrieval and display of minutes where a meeting activity is underway). When both parties use the common concepts indicated when representing activities and services in their topic maps, it is possible for a software tool to analyze the topic maps to check for equivalence. Thus the implementations of service “LazyRecall” in TM1 and service “RememberSoon” in TM2 could be considered equivalent as both are of type “Recall” service and have equivalent quality of service attributes. In addition, activity “Brainstorm” in TM1 and activity “Creative Session” in TM2 could be considered equivalent as both are supported by services of equivalent types.

Given two topic maps, the Ontology Comparison component will produce an analysis XML document that indicates which services and activities are exactly equivalent, partially equivalent or do not match at all. The Ontology Comparison component uses the facilities of a topic map platform (‘Topic Maps for Java’ from tm4j.org) to traverse both topic maps looking for matches. When a match is found it is documented in the analysis XML document. The Transformation Rules generator component will take this analysis XML document and create two XSL rule-sets (one for each direction of communication) that can be used dynamically at runtime to transform the terms found in service requests and service responses into terms understandable by the other party. For example, where it is found that term "LazyRecall" matches term "RememberSoon" in the analysis document, then an XSL rule is created that would substitute "RememberSoon" for "LazyRecall" (and vice versa depending on direction) in the raw request/response documents going through the gateway. How the Transformation Rules generator interprets the analysis document (i.e. how it handles exact/partial/no matches) is determined by an a master XSL rule-set. In this way the behavior of the component can be adapted easily. For example it might be stated in the master XSL rule set that any "no match" terms found by the gateway during translation would be ignored or for example it might be stated that where a "no match" term is found that the translation is aborted.

Although the prototype shows that the approach holds promise, a number of questions are being pursued in the evolution of this work:

- How realistic is the definition of a “minimum” of common concepts to be used in pervasive computing ontologies? If not, how can equivalence between concepts be deduced?
- What effect will different ontology representations have on the performance of the comparison analysis?
- What software mediation or software adaptivity approaches are most appropriate in the construction of the semantic interoperability gateways?

7. SEMANTIC SERVICE ENGINEERING

Though the above section demonstrates the potential power of ontologies in performing automated semantic interoperability at runtime, this is predicated on the availability of a large number of different services all willingly exposing their interaction semantics. To reach such a conducive environment will, as our roadmap indicates, require the cooperation of the developers of services, both atomic and composite, in accurately defining the service semantics for their products. For them to be motivated to undertake this perceived overhead, it must be both beneficial for them, by increasing the utility of their services, and easy to accomplish. Easing the process of semantic mark up for services requires integration of the software engineering process that produces the implementation of the service and the knowledge capture and maintenance process that results in usable ontologies. In a broad sense this requires an integration between the fields of knowledge engineering and software engineering. Some progress is being made towards such an integration. Based on the observation that ontology languages such as DAML+OIL share many characteristics with object-oriented model, the OMG is currently examining support for ontology modeling within the Unified Modeling Language [16]. Some object-oriented CASE

tools already offer support for exporting class models in ontology languages, albeit with proprietary mappings between the object-oriented and ontology meta-models.

Though support for ontology languages is a necessary step for the integration of software engineering and knowledge engineering processes for service-oriented systems, it is not sufficient. The location, navigation, manipulation and processing of ontologies need to be integrated into the development process for service-oriented software development. In particular, the potential for ontologies to support machine inference about the selection and testing of services and the verification and simulation of service compositions must be exploited in easing the semantic service development process. However, as the optimization of techniques for service-related machine reasoning are still at the research stage, semantic service development tools are still largely based on knowledge management tools [18] and are yet to be integrated into main stream engineering tools.

Machine reasoning has previously seen application to service engineering, principally in the telecommunications industry. Formal languages such as the Service Description Language (SDL) [9] were used to provide formal definition of services implemented in component-oriented telecommunication control systems, often referred to as intelligent networks. Reasoning with such formal service descriptions was used initially to detect unwanted feature interactions in the service creation process for intelligent networks, but research was also conducted into using them for user evaluation trials, design simulations and automation of test case generation [14]. SDL bears many similarities to DAML-S, as both express logical assertions about preconditions and post-conditions, i.e. effects, of service invocation, and can also express logical relationships and constraints between service inputs and output. However, the use of formal reasoning in intelligent network service creation was often seen as an addition to the service creation process, and therefore an overhead the benefits of which had to be directly assessed against the equivalent costs of traditional software engineering processes. Also, the logical concepts used in defining formal assertions for a particular intelligent network service was tied to that particular service and offered little benefit outside the development of that service. At best such formally defined objects were restricted to reuse in the intelligent network software market which was dominated by a few large players, and thus not open to the utility of formal definitions as a commodity or shared resource.

An ontology based approach to defining formal service semantics offers a more sustainable model for investing in the development of formal semantics than could ever have been the case for telecommunications service creation. The ubiquity of the WWW and the low entry costs to providing web services has provided a strong motivation of communities to collaborate in developing and freely publishing ontologies for their application domains in order to overcome the current semantic-based hurdle to cheap, widespread interoperability. Service developers are therefore likely to benefit from an ever growing reservoir of freely available ontologies providing logical concepts and related inference rules on a wide range of subjects. At the same time, the growing number of web services will force increased automation in their discovery, composition and invocation, such that services without well designed formal semantics will become increasingly less appealing to potential users. Service developers, therefore, will be increasingly motivated to contribute to the definition of

reusable semantics in the form of ontologies. Previously, in the context of component-oriented communication management software development, we have argued that explicit publication of the information content of services eases the service integration process and the need for interoperability promotes convergence of such models over time [13]. We expect to see a similar effect in the development of semantic services, where freely available models of information used in existing services, i.e. ontologies, will prompt their use in new services, easing the semantic interoperability problem over time. If our assertion that an interoperability problem similar to that addressed by semantic web services will face pervasive computing, then developers of pervasive computing devices will have the same benefits of ontology reuse and the same motivations in contributing to ontology development.

8. CONCLUSIONS AND FURTHER WORK

In this paper we have argued that the development of pervasive computing environments faces an interoperability challenge similar to that facing machine to machine interaction on the Web. We assert that using interoperability standards is not sufficiently scalable an approach to this problem, and that the semantic service approach being proposed for the Web is an appropriate route to a suitable adaptive solution. We outline a solution to dynamic interoperability based on ontologies, presenting some initial results using topic maps. We also discuss the potential challenges facing the developers of semantic services, including the need for integrated software and knowledge engineering tools and the harnessing of existing ontologies in developing semantic service descriptions and promoting semantic convergence.

Recent work has shown that expression of service capabilities in DAML-S facilitates service matching [19]. This approach compares the inputs/outputs of the required service with the inputs/outputs of the actual services available. It relies upon the availability of a common ontology for the inputs/outputs in order to facilitate matching. We plan to build upon this work, our experience of using Topic Maps, and the emergence of OWL to propose a solution that will enable service matching to take place in scenarios where heterogeneous ontologies are involved, which as stated earlier we believe will be the dominant situation. We hope to be able to assess if the DAML-S ontology provides the core common semantics needed for efficiency in dynamic semantic interoperability or to identify any further ontological concepts that are required. We intent to deploy a pervasive computing platform to investigate this by using semantics in: the adaptive composition of services, the policy-based management of these service compositions and the aggregation of context information from multiple sources. Our target platform is a meeting room application server, supporting control of data projectors and making available a number of office-style and unix command line applications as individual semantic services. These services will be complemented by proxies exposing semantic service representations of nearby services such as printing and file storage as well as web services, such as Yahoo's search engine service. User terminals, such as PDAs, will provide wireless access via Bluetooth or 802.11 to compositions of these services adapted to the user's preferences and context.

In support of this research, and to further explore the issues of semantic service engineering, we are planning a platform for the rapid design and evaluation of atomic and composite services.

This will be based on an existing reflective object-oriented application prototyping platform that will be extended to support ontological models and integrated with appropriate reasoning tools in an extensible manner. Overall, though we see semantic services supporting increasing automation of service composition that involve little or no human intervention, we recognize and actively address the need to support the development and reuse of ontologies in the human-led service development and composition process. Only such a dual approach will spur the widespread deployment of semantic services needed for automated composition while also using open ontologies to keep semantic divergence under control and thus aid automatic semantic interoperability.

9. ACKNOWLEDGMENTS

This work was funded by the Irish Higher Education Authority under the M-Zones programme.

10. REFERENCES

- [1] Van der Aalst, W.M.P. (2003), "Don't go with the Flow, Web services composition standards exposed", IEEE Intelligent Systems, vol18, iss 1., Jan/Feb.
- [2] Campbell et al (1995), Campbell A, Chalupsky H, and Shapiro S. "Ontological mediation: An analysis." Unpublished manuscript
- [3] Ceri and Widom (1993), Ceri S., Widom J., "Managing Semantic Heterogeneity with Production Rules and Persistent Queues", Proceedings of the 19th VLDB Conference, Dublin, Ireland, pp. 108-119.
- [4] Chakraborty, D., Perich, F., Joshi, A., Finin, T., Yesha, Y. (2002), 'A Reactive Service Composition Architecture for Pervasive Computing Environments', Technical Report TR-CS-02-03, Computer Science and Electrical Engineering, University of Maryland Baltimore County
- [5] Christensen, E., Curbera, F., Meredith, G., Weerawarana, S. (2002), "Web Services Description Language (WSDL) 1.1", W3C Note 15 March 2001, <http://www.w3.org/TR/wsdl>
- [6] Corcho and Gomez-Perez (2001), Corcho, O. and Gomez-Perez, A., "Solving Integration Problems of E-commerce Standards and Initiatives through Ontological Mappings", In: Proceedings of the Workshop on E-Business and Intelligent Web at the 17th International Joint Conference on Artificial Intelligence (IJCAI2001), Seattle, USA, August 5, 2001.
- [7] Cui (2002), Cui Z., Jones D. and O'Brien P., "Semantic B2B integration: issues in ontology-based approaches", ACM SIGMOD Record, vol 31 no 1, pages 43—48, ACM Press
- [8] DAML (2003), DARPA Markup Language, www.daml.org, Visited Apr 2003
- [9] Ellsberger, J., Hogrefe, D., Sarma, A. (1997), "SDL, Formal Object-oriented language for communicating systems", Prentice-Hall
- [10] Goh (1997), Cheng Hian Goh, "Representing and Reasoning about Semantic Conflicts in Heterogeneous Information Sources", PhD, MIT.
- [11] Kim (2002), Kim H., Predicting how ontologies for the semantic web will evolve, Communications of the ACM, vol 45 no 2, pages 48-54, ACM Press.
- [12] Klein (2001) Klein M. "Combining and relating ontologies: an analysis of problems and solutions". In IJCAI-2001 Workshop on Ontologies and Information Sharing, pages 53-62, Seattle, WA, 2001.
- [13] Lewis, D., Wade, V., Cullen, B. (2003) "Towards the Technology Neutral Modelling of Management Components", Journal of Network and Systems Management, vol 11, no. 1, March
- [14] Lodge, F., Kimbler, K., Hubert, M. (1999) "Alignment of the TOSCA and SCREEN Approaches to Service Creation" in Proceedings of the 6th International Conference on Intelligence in Services and Networks, Barcelona, Spain, Springer-Verlag, pp277-290
- [15] Obrst and Liu (2003), Obrst L. and Liu H., "Knowledge Representation, Ontological Engineering and Topic Maps", Chapter 7, XML Topic Maps, Jack Park (ed.), Addison-Wesley.
- [16] OMG (2003) "UML Profile for Web Ontology Definition, Request for Proposals", OMG document ontology/2003-12-03.
- [17] O'Sullivan D., Wade V., Lewis D., Conlan O. (2002), "Enabling Smart Space Service Mobility Negotiation by Ordinary Users", Proceedings of EURESCOM Summit 2002 Heidleberg, Germany
- [18] McIlraith, S.A., Son, T.C., Honglei Zeng, H. (2001), 'Semantic Web Services', IEEE Intelligent Systems, 16(2), March/April 2001
- [19] Paolucci M., Kawamura T., Payne T, and Sycara K (2002) ., "Semantic matching of web services capabilities", ISWC2002.
- [20] Sheth (1998) Sheth A.P. , "Changing Focus on Interoperability in Information Systems: From System. Syntax, Structure to Semantics", M. F. Goodchild, M. J. Egenhofer, R. Fegeas, and C. A. Kottman (eds.) Interoperating Geographic Information Systems, Kluwer.
- [21] Visser (1997), Visser, P.R.S., Jones, D.M., Bench-Capon, T.J.M., and Shave, M.J.R., "An analysis of ontology mismatches; heterogeneity versus interoperability." AAAI 1997 Spring Symposium on Ontology Engineering, 1997.
- [22] Wache (2001), H. Wache, T. Vögele, U. Visser, H. Stuckenschmidt, G. Schuster, H. Neumann, and S. Hübner, "Ontology-based Integration of Information - A Survey of Existing Approaches," In: Proceedings of IJCAI-01 Workshop: Ontologies and Information Sharing, Seattle, WA, 2001, Vol. pp. 108-117.
- [23] W3C (1999) World Wide Web Consortium, Resource Description Framework (RDF) Model and Syntax Specification, W3C Recommendation, <http://www.w3c.org/>
- [24] W3C (2003) Ontology Web Language, <http://www.w3.org/2001/sw/>, Visited Apr 2003.
- [25] Weiser, M. (1991), "The Computer of the 21st Century", Scientific American, vol. 265, no.3, pp 66-75.
- [26] XTM (2003) XTM Authoring Group, XTM: XML Topic Maps 1.0. www.topicmaps.org. Visited Apr 2003
- [27] 'DAML-S: Semantic Markup for Web Services', The DAML Service Coalition, <http://www.daml.org/services/>, October 2002