Integrated Context Management for Multi-domain Pervasive Environments

Rui José, Filipe Meneses, Adriano Moreira

Information Systems Department, Universidade do Minho, Campus de Azurém,
4800-058 Guimarães, Portugal
{rui,meneses,adriano}@dsi.uminho.pt

Abstract. An important part of the value of ubiquitous computing environments is in their ability to interact with external domains. This paper addresses the issue of cross-domain context management in a scenario of seamless integration between a user home domain and a ubiquitous computing environment. The work is based on the broader concept of Value ADded Environment (Vade), where multiple integration possibilities are explored. This paper is focused on location context management issues and describes two integration paths for sharing context data: the possibility to provide local applications with access to the context of a visiting user, and the possibility to provide global applications with access to context sources that are specific to the visited location.

Introduction

Context awareness is central to the vision of ubiquitous computing and the availability of a vast range of sensors and inference techniques is providing ever increasing opportunities for generating rich context information. However, sophisticated context awareness is usually only possible in systems aimed at small and specific physical infrastructures, where all the sensors, context-management components and applications are carefully designed to work together. Following an alternative path, Mobile Network Operators (MNO) have been launching services that use the position of mobile terminals to provide users with information that is associated with their particular location. These services have the advantage that they can be useful across a potentially very vast range of locations, but they lack the integration with local infrastructures, without which they cannot reach such strong association with user context as the applications created for a specific environment.

The research described in this paper aims to combine the two previous approaches, and is specifically targeted at the issues of interoperability and seamless integration between a user home environment (e.g. that of a MNO) and the local environments visited by a user. The concept of local environment considered in this paper is that of a ubiquitous computing environment, such as shopping centres or airport terminals, that integrates a set of local facilities such as communications networks, sensors, public displays, network services and local applications particularly tailored for that environment. It is also considered to be administratively independent from the MNO.

Our approach for interoperability and integration is based on the concept of Value ADded Environment (Vade) as an administrative and physical domain where the locally available ubiquitous computing facilities can be combined with the home environment of visiting users. As a result of this integration, the user should be provided with an enhanced and extended functionality corresponding to the seamless combination of the global and local functionality. The global elements provide the common ground necessary for giving users a familiar application environment, whereas local elements provide the necessary link with the specificities and services from the local physical environment. The Vade approach explores the Boundary Principle [1], which states that ubicomp environments should be designed with boundaries that demarcate their content and scope but do not necessarily limit their interoperability.

Context management is clearly one of the areas in which the potential for integration is stronger. We see context management as a collaborative task between the Vade infrastructure and the external home environment of a visiting user, and we have identified two complementary approaches for this integration. The first is to allow the home environment of the user to complement its context with information obtained from the local environment. The other is to allow the local applications to access context information from the home environment of the user.

In our research we have developed a prototype of a Vade that supports both forms of integration, and explored how application on both domains can take advantage of this integration. In the remainder of this paper, we start by describing the key aspects of the Vade approach and how the process of context management is supported in a Vade system. We then briefly describe our experience of deploying a prototype Vade for our university campus and how we enabled the two forms of context integration mentioned above. The paper ends with an analysis of related work and conclusions.

The Vade Approach

In order to clarify the overall background of our approach we will now provide a brief overview of the Vade system and describe how it can support the integration between a local environment and the home environment of a visiting user. We have considered two distinct models for this integration. The first model, called *Vade-aware global application*, corresponds to a global application that, as the user moves, is continuously searching for resources that are specific to the current location of that user. A global application is not aimed at any Vade in particular and may be used regularly outside any Vade, but entering a Vade is a major opportunity for finding new, probably much more specific, resources and context information, which the application can then use to improve the value of the functionality provided. For example, a global map application, can present a map of the area where a shopping centre is located, but, as the user is detected as being inside the shopping centre (a Vade), a detailed plan of the shopping can be offered to the user. The other integration model, called *Vade local application* provides functionality that is specifically associated with a particular Vade. When a user enters a Vade, these applications should seamlessly integrate the user's mobile information space, i.e. the user should

not need to look for them, install them or configure them. The main advantage of this application model over the Vade-aware model is that, since these are local applications, they can be tailored for the particular environment for which they have been produced.

We have grouped the collaborating entities into three separate domains, more specifically: the Personal, the Vade, and the Home domains. The Personal Domain corresponds to the physical and logical elements carried by a person. It includes the mobile devices carried by the user and the corresponding computing capabilities. The Vade Domain is the local infrastructure that supports the Vade functionality, including its services and local devices. Finally, the Home Domain corresponds to the external elements that are associated with the user. We have been particularly interested in scenarios in which this Home Domain is managed by a Mobile Network Operator (MNO) or Wi-Fi operator, as this would enable interaction between Vades and key horizontal services, such as authentication, user profiles and terminal location. Fig. 1 presents the main elements of each domain and the respective dependencies.
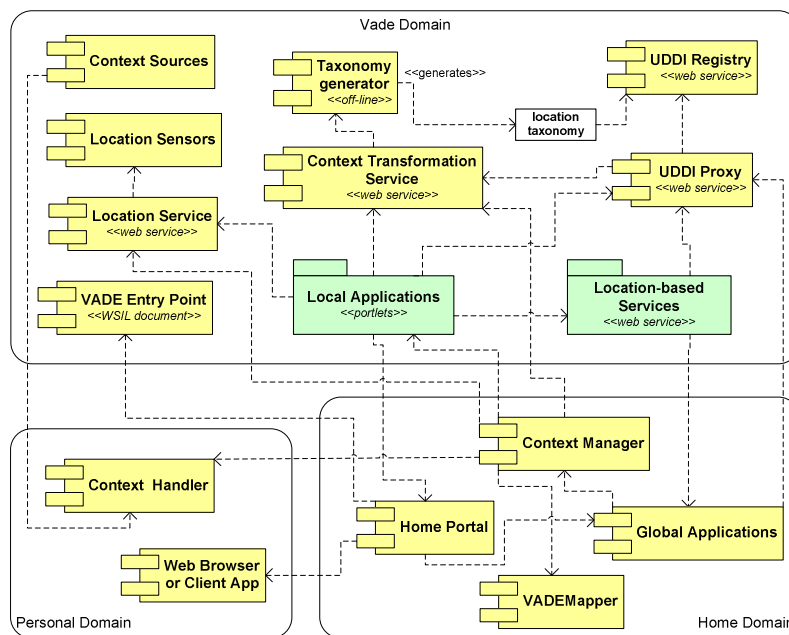
Fig. 1. Architecture of the Vade System

We assume that the Home Domain includes some form of mobile information space (e.g. a portal) where users may activate applications that match their specific needs. These global applications provide global functionality based on the context information available at the Context Manager. When the user enters a Vade, the functionality provided will be a dynamic combination of pre-defined user preferences, currently active applications, current user context and locally available services and applications. This integration is accomplished through two types of resource:

Location-based services and Local Applications. Location-based services will be used by end-user applications, global or local, as a means to enhance their association with the current physical environment of the user. In order to support dynamic finding and invocation, these services should be self-describable and meet strong interoperability and openness requirements. On the other hand, Local applications are user-facing applications that are specific to the Vade.

Context management in a Vade scenario

The central element in the context management process is the Context Manager in the Home Domain, which supports the collection of data from multiple context sources, possibly maintained by multiple administrative entities, enhances and aggregates that data, and finally provides an interface where authorised applications can access context data about a particular user. While we try to maintain an open view on what a Context Manager should be, we have assumed that user devices can collect context information and transmit it to the Context Manager. The Context Manager can also query the Location Service of the Vade where the user is currently located. Additionally, the Context Manager may calculate or estimate the value of other attributes based on raw data provided by sensors, notify one or more applications that new contextual data is available, and obtain, from Space Models servers, extra contextual data related to that raw data. When the user enters into a Vade, this Context Manager will try to combine the global context information and user preferences with the context data provided by the local infrastructure.

The Vade Location Service manages the location information collected by the local infrastructure, acting as a mediator between diverse local positioning mechanisms and consumers of position information, e.g. the Context Manager of a visiting user. Alternatively, position information can also be collected by Context Handlers on user devices and transmitted directly to the Context Manager of that user. A distinguishing characteristic of the Location Service is that it must be optimised for allowing the information it manages to be combined with external sources of information. Unlike other location services whose main goal is to produce the current position of a user, possibly fusing data from multiple sensors, the Vade location service should also be able to provide external entities with the data collected from multiple positioning mechanisms, thus allowing the fusion of the results to be made by the external environment where more complete information may exist.

The Context Transformation Service enhances the use of context within a Vade by supporting the mapping between different forms of context (very similar to the notion of context interpretation as described in [2]). The Context Transformation Service is supported by a Space Models server which contains information about what locations are bounded to each positioning data and information about the containing relationships for the locations. It supports three types of transformation, all of them based on location. We have decided to use location contexts, together with a space model of the Vade, as the backbone for all context transformations, because of its static structure and its ability to leverage other forms of mapping. The first transformation is from position information (raw data as provided by sensors) into a location context from the Vade space model. The second transformation explores the

topological relationships in the Vade space model to further enhance the determination of context. Since the Vade space model represents each location context and the respective contained relationships, it is possible to obtain from a location context all of its containing locations. This can be used to infer that, for example, a user at a particular room is also at a particular building. Finally, the other type of transformation supports the association between a location context and other forms of context associated with places. For example, "*If in room A then probably participating in meeting X*".

Another central issue in cross-domain context management is the definition of data formats for representing context information. Given the current lack of standards or any other widely accepted formats for representing context information, existing systems are necessarily based on proprietary data formats that are only valid within the narrow scope of the specific systems for which they have been created. Also, given the varied dimensions of context and the multiple application requirements, a context representation should be simultaneously, usable by applications, independent from the context sources, and also flexible enough to accommodate future needs. Taking these issues into account, we have used the data format for context representation described in [3].

Determining Vade containment and association

Presence in a Vade is also context information, and is essential for triggering the association with a Vade and for finishing that association. In order to begin an association, the Home Domain needs an initial reference for a Vade, in the form of a domain name. With this information, the Home Domain can access a resource called Vade Entry Point, that is essentially a Web Services Inspection Language (WSIL) document at a standard location (domain.name/inspection.wsil) and obtain the references for all Vade core services. The challenge is thus how to provide the Home Domain with the domain name for the Vade where the user is physically present.

One approach is to assume that the user's device is capable of collecting Vade references embedded in proximate context sources, such as Bluetooth beacons, or visual codes, and transmit them to the Context Manager on the Home Domain. However, this approach has some limitations when used as the single mechanism for managing associations. Firstly, many context sources, such as barcodes or low capacity RFID tags, while useful for determining location within a Vade, are very limited in the amount of information they can provide, and are thus more adequate for providing relative location references within an already known Vade. Secondly, unless there is full coverage within the Vade area, it becomes very difficult to determine when a person has left the Vade environment.

An alternative approach is to have an external Vade directory, called Vade Mapper, to support the mapping between geographical positions and Vade domains. This approach is particularly suited for cases in which the user is a subscriber of a MNO, as there is the possibility to obtain his or her position across very large areas and generate events based on the presence of the user in the geographical scope of a particular Vade. These events, captured by the Vade Mapper, can trigger a notification on the Context Manager to inform it that the user entered a Vade. These

two approaches complement each other very well in managing the association with Vade environments. The Vade Mapper, even with considerable overlap between environments and near-border situations, can reduce the options to a very few (if not only one) environment. Local mechanisms can solve any ambiguities and also provide some user control over the association process.

Campus prototype

We have developed a prototype of a Vade system in which one of our University Campi plays the role of a Vade. In our prototype, the Home Domain includes a web Portal, a Context Manager and a Vade Mapper. Our implementation of the Context Manager supports multiple types of Context Handler and various types of interface, such as SOAP, UDP and RMI, thus allowing Context Handlers to transmit their data using the most convenient technology [3]. It also includes global vade-aware applications, such as an environment browser that describes the current location of the user. The Vade Domain corresponds to the coverage area of a set of Wi-Fi access points in our campus. The Location Service provides the position of users based on the BSSID of the Wi-Fi cell to which the user is associated. Our Vade Domain also includes implementations of the Vade Entry Point, Context Transformation Service, UDDI Register and Proxy and a set of local applications, such as a virtual juke-box and a "find a person" application.

Given our focus on interoperability, we have decided to make all these resources available as web services and to use the Universal Description, Discovery, and Integration (UDDI) standards. In the case of local applications, the Web Services Remote Portlets (WSRP) specification is used for allowing them to be integrated into the Vade-aware portal of the Home environment. A private UDDI Registry is used as the main repository for the information about all local services and applications. In order to support the association between services and locations, and especially to support location-based queries, we have introduced specific location taxonomies and other enhancements into the use of this registry [4]. Instead of making queries directly to the registry, external applications place their queries to an UDDI proxy that uses its privileged knowledge about the local environment to transparently support and optimise location-based discovery. This approach has allowed the UDDI standard to be supported while enabling all applications to benefit from the specific mechanisms for location-based queries. For the User Domain we have considered mostly portable computers, Wi-Fi enabled, and running local Context Handlers that capture multicast UDP beacons. These beacons are used as a context source and provide a reference to the Vade Entry Point. All access to applications is based on web browsers.

Enabling cross-domain context management

Context management is always a very privacy-sensitive topic, and even more in situations of cross-domain context management such as this in which we expect the Home and Vade domains to exchange context information. We must therefore support

a balanced combination between integration possibilities and privacy guarantees, and we now describe our approach for guaranteeing these requirements for each of the two integration paths that we have considered, more specifically the ability of the home Context Manager to complement its context information with data obtained from the Vade Domain and the ability of Vade local applications to obtain context information from the Home Domain.

**External Access to location information in the Vade Domain**

One of the goals of the Vade approach is that the Home environment should be able to enhance its own context information about a particular user with information obtained from the Vade that is currently being visited by that user. The motivation is that by having access to an infrastructure that is co-located with the user, the Home environment can gather much more precise and dynamic information than what would be possible using only global mechanisms. The challenge, however, is how to enable this sharing of information without compromising user privacy. Ideally, a user should be able to access as much as possible of the Vade functionality while revealing as little as possible about its identity or any other personal information. In particular, the true identity of the user should not need to be revealed to local applications.

From the perspective of the Vade Domain, location information about a visiting user is managed by the Location Service, which should be accessible to the Context Manager of that user. This provision of location information to third party entities calls for secure mechanisms for access control. For example, a local positioning service can automatically identify a wireless device by its MAC address and track its position across the Vade. However, only applications authorised by the owner of that device should have access to that information. Furthermore, this access must be granted without any a priori information about visiting users, which leads to the problem of identifying the anonymous owner of the location information.

We have addressed this problem with a solution that combines the use of pseudonyms with a initialisation procedure for binding those pseudonyms with specific location identifiers, such as a Wi-Fi MAC address, an RFID tag or a Bluetooth identifier. As part of the creation of a new association with a Vade, the Home environment, more specifically the Context Manager, starts the procedure by generating a new pseudonym specifically for that particular Vade. That pseudonym is then registered in the Vade (Location Service) along with a set of location identifiers that are known to be associated with that user. This registration instructs the Location Service to dynamically create an association between that pseudonym and the location identifiers, thus overcoming the issue that the association between users and their identifiers is not known a priori by the Vade. However, for many identifiers it is a fairly simple process to obtain someone else's identifiers and then use them to track that person. To guarantee that only the owner of the location information can access it, an additional procedure is needed to verify that the registering entity is in fact the owner of that location identifier.

Unlike the registration process, which is common to all location identifiers, the verification process must be specific for each type of location identifier. In our prototype we have developed a solution for the specific case where the location

identifier is a Wi-Fi MAC address, as this was one of our main target scenarios. When a user, carrying a Wi-Fi enabled device, enters the Vade, the Context Manager generates a new pseudonym and provides it to the Vade Location Service along with the respective MAC address. The Location Service will then obtain the IP address of the device (using SNMP to query the ARP table in the Access Points) and, using that address, tries to confirm the association between the pseudonym and the MAC address by challenging the Context Handler on the mobile device with a random string. Upon receiving the challenge, the context handler will use the pseudonym as a key to encrypt the challenge string and return it to the location service. If the location service is able to decrypt the message, the association is created. This process is depicted in Fig. 2.
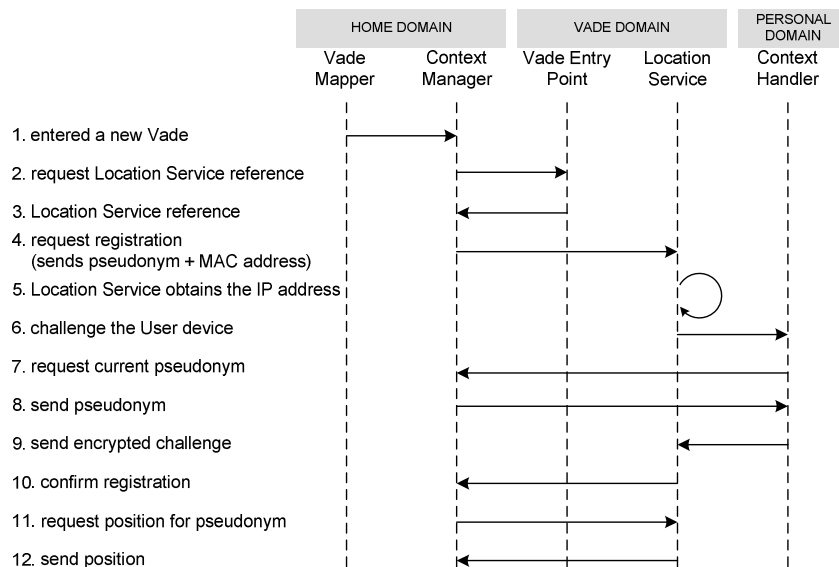
Fig. 2. Registration on the Vade location service

After the associations between the pseudonym and the MAC address has been created and verified, the location service will accept location queries based on the pseudonym. These queries will be made mainly by the Context Manager. Global applications, i.e., those that are deployed and maintained by the Home Domain, will not need to know about pseudonyms as they do not connect to the Vade directly. However, they will benefit from richer context information that will be provided by the Context Manager. Local applications are invoked from the Home portal which may provide them the pseudonym of their current user. They can then query the Location Service for information about the user position. The use of pseudonyms allows these applications to access location information about their user and provides a simple mechanism for defining access permissions. In order to reduce the possibility of unauthorized use of the pseudonyms, and thus increase the privacy guarantees, the pseudonym should be changed periodically and re-registered into the Location Service.

Secure and privacy-enabled access to the global context information

Another goal of the Vade approach is that local applications, when integrated into the Home Domain, should also be able to benefit from the services of the Context Manager. This poses considerable risks because there may be many Vade systems and the Context Manager cannot assume that they are trusted parties. Our approach for enabling this complementary integration path while guaranteeing user privacy and appropriate access control, extends the use of pseudonyms to the queries made to the Context Manager, but this time a temporary password is also used to provide additional security. The overall process is based on the following steps:

- When an association is created between a Home Domain and a Vade, the Vade local applications that integrate the information space of the user are given the respective pseudonym by the entity which is responsible for managing their integration. Additionally, they are also given a password that is associated with the pseudonym and has the same limited time validity.
- Access to the context of a user by a Vade local application is performed by directly querying the Context Manager with the appropriate pseudonym+password pair. The sub-set of context information that applications can access is defined by the user through the user interface of the Context Manager, but never includes his identity. Additionally, Vade local applications can also use that pseudonym for directly querying the location service of their Vade in case they need more specific information.
- When the user leaves the Vade, local applications are removed from the user information space and their passwords are invalidated.

This approach allows the user to specify different levels of trust depending on the Vade, on the application or simply on the requirements of the situation.

Related Work

Context-awareness has been a topic of intensive research in recent years and the issues of interoperability and high-level abstraction have been particularly focused. The Context toolkit[5] is a well-known framework for the development of context-aware applications that provide abstraction from specific sensors and applications, but assumes a single administrative domain. The Nexus project [6] has developed a global context management platform to manage the location of people and objects based on a set of distributed servers and a specific query language. The provision of unified context information in different administrative and application domains is specifically addressed in [7], but the approach is conceptually different in that it assumes a federated model. In [8], short-lived, disposable MAC addresses are proposed as a means for reducing the opportunities for location tracking. Even though we assumed that this is not a major issue in the Vade model, since there are multiple entities involved, we believe that our proposed approach would be compatible with short-lived addresses, which could be renovated together with pseudonyms in order to increase privacy guarantees.

Conclusions

An important part of the value of ubiquitous computing environments is in their ability to interact with external domains and provide an integrated functionality to users. We have been investigating the concept of a value-added environment as an integration approach in which the functionality of a ubiquitous computing infrastructure is combined with functionality from an external home environment of a visiting user. In this paper, we addressed in particular how context management can also benefit from that integration and described a prototype system in which we explored two complimentary approaches for cross-domain context management, more specifically the possibility to provide local application with access to the external Context Manager associated with a visiting user, and the possibility to provide global applications with access to a location service that is specific to the visited location. The proposed approaches enable both domains to benefit from the integration, while providing appropriate privacy guarantees to the user. Even though there are still many interoperability issues to be addressed by the research community, we believe that this cooperation models will be an important part of future context-aware applications.

References

[1]   T. Kindberg and A. Fox, "System software for ubiquitous computing", IEEE Pervasive Computing, vol. 1, pp. 70- 81, 2002.

[2]   A. K. Dey, "Providing Architectural Support for Building Context-Aware Applications", in College of Computing: Georgia Institute of Technology, 2000.

[3]   F. Meneses and A. Moreira, "A Flexible location-context representation", Proceedings of The 15th IEEE international Symposium on Personal, Indoor And Mobile Radio Communications (PIMRC 2004), Barcelona, Spain, 2004.

[4]   H. Pinto, R. José, and N. V. Boas, "Using a Private Uddi for Publishing Location-Based Information to Mobile Users", presented at 7th ICCC/IFIP International Conference on Electronic Publishing (ELPUB2003), Guimarães, Portugal, 2003.

[5]   D. Salber, A. K. Dey, and G. D. Abowd., "The Context Toolkit: Aiding the Development of Context-Enabled Applications", presented at Conference on Human Factors in Computing Systems (CHI '99), Pittsburgh, PA, 1999.

[6]   Nexus, "Nexus Project", http://www.nexus.uni-stuttgart.de/

[7]   O. Lehmann, M. Bauer, C. Becker, and D. Nicklas, "From home to world - supporting context-aware applications through world models", Proceedings of the Second IEEE Annual Conference on Pervasive Computing and Communications (PERCOM'04), Orlando, Florida, USA, 2004.

[8]   M. Gruteser and D. Grunwald, "Enhancing location privacy in wireless LAN through disposable interface identifiers: a quantitative analysis", Proceedings of the 1st ACM international workshop on Wireless mobile applications and services on WLAN hotspots (WMASH'03), San Diego, CA, USA, 2003.