International Journal of Pervasive Computing and Communications

Deploying pervasive secure knowledge management infrastructures

Apostolos Malatras, George Pavlou, Petros Belsis, Stefanos Gritzalis, Christos Skourlas, Ioannis Chalaris,

Article information:

To cite this document:
Apostolos Malatras, George Pavlou, Petros Belsis, Stefanos Gritzalis, Christos Skourlas, Ioannis Chalaris, (2005) "Deploying pervasive secure knowledge management infrastructures", International Journal of Pervasive Computing and Communications, Vol. 1 Issue: 4, pp.265-276, https://doi.org/10.1108/17427370580000130
Permanent link to this document:
https://doi.org/10.1108/17427370580000130

Users who downloaded this article also downloaded:

For Authors

About Emerald www.emeraldinsight.com

Deploying Pervasive Secure Knowledge Management Infrastructures

APOSTOLOS MALATRAS, GEORGE PAVLOU
Department of Electronic Engineering, Centre for Communications Systems Research, University of Surrey, UK
Email: {a.malatras,g.pavlou}@surrey.ac.uk

PETROS BELSIS, STEFANOS GRITZALIS
Department of Information and Communication Systems Engineering University of the Aegean, Karlovasi, Samos, Greece
Email: {pbelsis,sgritz}@aegean.gr

CHRISTOS SKOURLAS, IOANNIS CHALARIS
Department of Informatics, Technological Education Institute, Athens, Greece
Email: {cskourlas,ixalaris}@teiath.gr

Received: October 24 2005; revised: November 30 2005

Abstract— **Pervasive environments are mostly based on the ad hoc networking paradigm and are characterized by ubiquity in both users and devices and artifacts. In these inherently unstable conditions and bearing in mind the resource's limitations that are attributed to participating devices, the deployment of Knowledge Management techniques is considered complicated due to the particular requirements. Security considerations are also very important since the distribution of knowledge information to multiple locations over a network, poses inherent problems and calls for advanced methods in order to mitigate node misbehaviour and in order to enforce authorized and authenticated access to this information. This paper addresses the issue of secure and distributed knowledge management applications in pervasive environments. We present a prototype implementation after having discussed detailed design principles as far as the communications and the application itself is regarded. Robustness and lightweight implementation are the cornerstones of the proposed solution. The approach we have undertaken makes use of overlay networks to achieve efficiency and performance optimization, exploiting ontologies. The work presented in this paper extends our initial work to tackle this problem, as this was described in [28].**

Index Terms— **Mobile ad hoc networks, pervasive environments, communication and information security and distributed knowledge management**

I. Introduction

The proliferation of mobile ad hoc networking solutions observed in the past few years and the high rates of user adoption regarding this technology combined with the continuously increasing number of mobile devices [7] [8], leads us to consider that there is an established paradigm shift from traditional, infrastructure networking towards a mobile, operator - free and with no fixed infrastructure type of networking, the one based in Mobile Ad Hoc Networks (MANETs) [1]. MANET based networks together with other emerging networking technologies such as sensor networks will constitute the foundations for future pervasive applications. The major strengths of this technology lie in the fact that it is easy to be deployed and has a very low cost, while allowing for user creativity through the lack of central, authoritative management [5] [6]. In the research area of mobile ad hoc networks, fundamental work was undertaken at the Terminodes project [4]. The notion of combining terminal and node capabilities in every mobile node is the driving force of this project. The notion of pervasive computing and ubiquity are strongly correlated to that of mobile ad hoc networking technologies that thus assist in reaching Weiser's innovative conceptualization of future computing [9].

MANETs undoubtedly are not a panacea for every networking problem and the emerging pervasive realm. Noteworthy drawbacks include their highly dynamic topology, since every node participating in a MANET is mobile and possibly very volatile. These constant topological variations will eventually lead to a continuous state of network instability, which in turn can extremely deteriorate the performance of applications and services on these networks. Access to knowledge scattered across a MANET may be hindered as a sequence of the volatility that characterizes the network topology. Another important issue is that typically MANET devices have limited resources as far as storage, energy and processing capabilities are concerned [2] [3]. We focus our field of networking in the ad hoc paradigm since this is the most commonly used one in the pervasive domain. It is obvious there is a need for techniques that mitigate these problems in order to be able to harness the vast range of advantages that MANETs have to offer.

Knowledge on the other hand, is probably the most important capital for an organization, constituting thus its management an issue of high significance. Modern organizations and user environments in general are characterized by a diversity and dispersion in location of both users and knowledge information, leading to pervasive scenarios like the ones described. Another case that is becoming increasingly popular is that of establishing coalitions amongst administrative domains in order to share information towards a common goal, i.e. two enterprises collaborating to promote a new service/product. Additional requirements regarding security mostly arise from such scenarios. In this paper we study the requirements needed to deploy Knowledge Management techniques in such an inherently unstable environment and propose a system architecture that enables KM operations in a distributed, robust and secure way. Main motivation for our work is the absence of significant related work and the established cross-level benefits that are to be gained. We propose a concrete solution to locate, authenticate, access and retrieve information in an environment like the one described.

The remainder of the paper is organized as follows. After the brief introduction and motivation for our research in Section I, related work and background literature is studied in Section II. Section III analyses the requirements placed on the system design and Section IV raises the issues related to the system's design. Section V describes in full the system architecture, the communication protocols amongst the entities in the pervasive domain and specific implementation details of our prototype. In Section VI we discuss authorization and access control models as they are adjusted to fit to our pervasive scenario. Section VII provides initial results regarding the performance of our proposed research approach and presents an experimental scenario setup based on our implemented prototype. The paper concludes with future work and concluding remarks in Section VIII.

II. Related Work

A Mobile Ad Hoc Network (MANET) is a distributed, wireless, mobile, multihop network architecture that relies on no pre-existing network infrastructure for its deployment. The only requirement for a MANET to be deployed is the existence of at least two mobile nodes (MN) in communication range with each other that will form the MANET. The nodes comprising the MANET are characterized by their dynamic nature, which means that they can move in and out of the network at any time and with no pre- or post-condition being met [1] [2]. Consequently, the network topology itself is not at all static, but it can change dynamically with time, being dependent on the high degree of mobility the mobile nodes exhibit [2] [3] [26]. In [6] a very extensive and up-to-date literature review on mobile ad hoc networking is presented. In the following we will present related work to that of knowledge management in pervasive environments.

In [18] an office work scheduling tool is presented for appliance to ubiquitous environments. This system emphasizes mainly on social aspects of KM and therefore it facilitates the socialization process as defined in Nonaka's definition about organizational knowledge transformation process [19], though there is not efficient support for other KM related activities or processes. In addition, there is no direct proof about the ubiquitous characteristics of the system.

Researchers from the University of Texas have been looking into a middleware framework for pervasive computing called PICO (pervasive information community organization) [27]. PICO's architecture comprises two elements, namely delegents (intelligent delegates) and camileuns (connected, adaptive, mobile, intelligent, learned, efficient, ubiquitous nodes). The notion of this approach is the creation of delegent communities that collaborate proactively to handle dynamic information, provide content delivery and facilitate application interfaces. Every intelligent device can be part of the PICO architecture. PICO is a promising proposal as far as pervasive and dense sensor networks are concerned, but it does not seem appealing to pure MANET environments, where the mobile nodes may be scattered and no other sensors are likely to be available.

ADAM [21] is a flexible and resilient infrastructure for secure distributed knowledge exchange that utilizes the notion of trust for authorization purposes. ADAM mainly collects knowledge related with a user's transaction history before authorizing her for transactions. ADAM handles scalability issues very effectively, though it mainly performs on environments characterized by total absence of well-defined organizational policy. This system is not a KM system with the broader meaning of the term as it does not encompass main KM related activities.

In [10] [11] [12] distributed service distribution on MANETs is discussed. We undertake a similar approach but we do not focus on the effects on the network performance, latency imposed, and throughput like their systems do. Our approach on distributing the lookup of knowledge sources will not deviate much but the focus will be on describing the system's efficiency in terms of the resources consumed on the device and its robustness as far as the dynamic nature of pervasive environments is considered. Network issues will be slightly ignored for now while they still reside on our future research interest agenda. Emphasis is being placed on providing a working and viable solution to enable knowledge management deployment on pervasive environments.

III. Requirement Analysis

The networking basis of ubiquitous and pervasive environments is that of mobile ad hoc networks. The term mobile does not necessarily mean that all the participants of the network are mobile, since any of them can be stationary for a small or large amount of time. In this environment knowledge management applications have to consider a series of issues:

- Limited resources in terms of processing power, memory capacity and energy usually characterize the devices participating in such environments. It should not be taken for granted that ”thick” devices with sufficient capabilities will exist in such networks. Sophisticated solutions proposed for more resource-powerful devices cannot be thus considered. Every possible solution has to take into account the significant energy considerations and consume the least acceptable memory and processing power. Solutions like cyber foraging where tasks from small devices migrate to powerful devices in the network should be examined. In addition, distribution of tasks

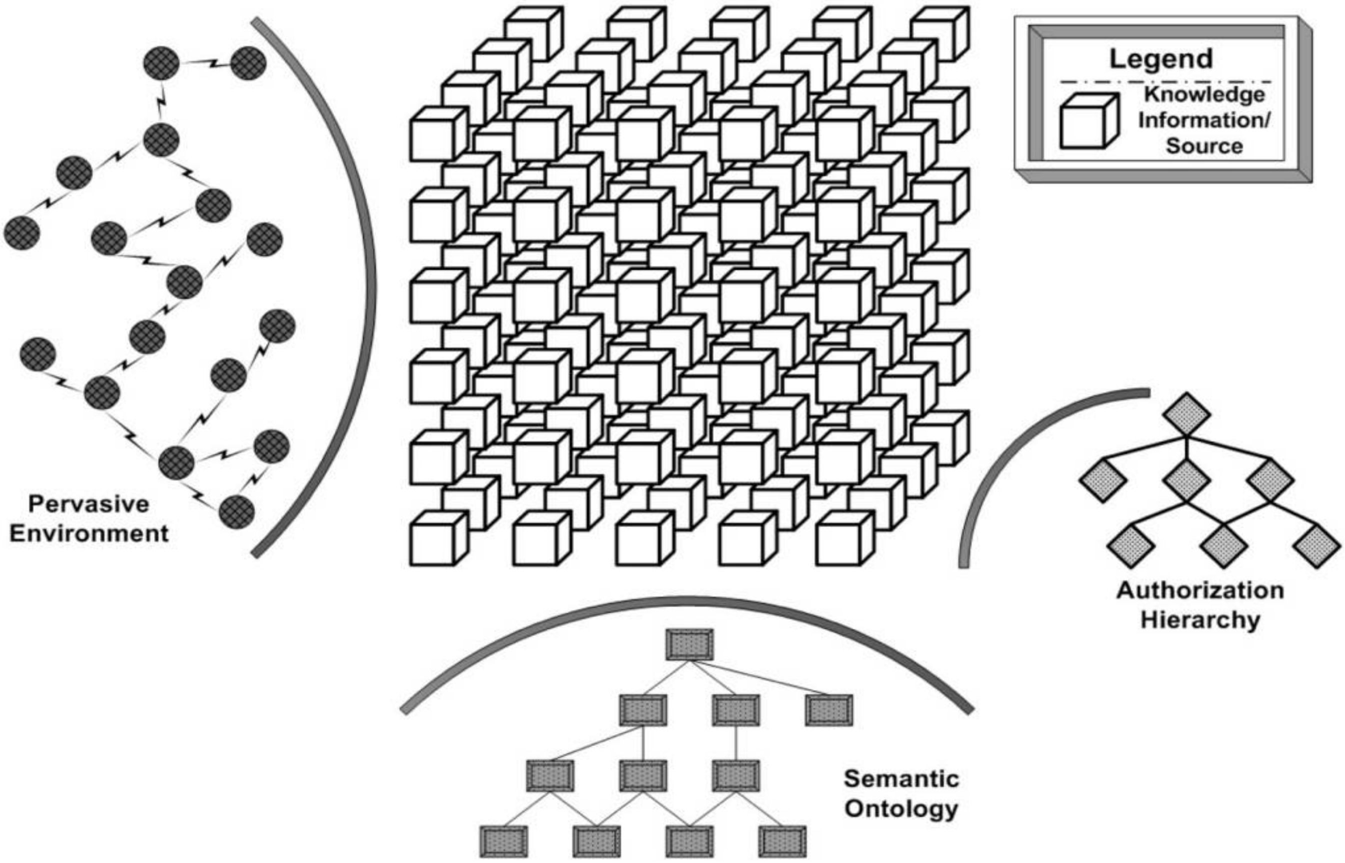

Fig. 1. Knowledge Information/Sources and the aspects that characterize it.

across the nodes of the MANET is another potential solution.

- Knowledge information is not limited in a closed environment as is the case in traditional networks, it spans on the contrary across multiple domains being characterized by a high degree of instability and ubiquity. Also, knowledge information might be derived and belong to different administrative domains placing further burden on he proposed system. Solutions to cater for the discovery of the appropriate knowledge information, from the appropriate source are necessary. Another issue to consider is the access to knowledge sources from different domains. Unified methods to perform this operation must be provided to the users with the principle of transparency being of prominent nature.
- Pervasive environments based on ad hoc networks inherit their dynamic nature. Communication links are unstable and prone on disconnections based on a variety of reasons with mobility of devices being the foremost. Every solution proposed for the secure knowledge management in pervasive environments has to be constructed in a fashion that enables robustness and resilience to these continuously changing conditions. Important issues to consider include the use of uninterrupted operation in case of nodes departing or joining the network. This can be achieved possibly by exploiting redundancy techniques.
- The mere notion of pervasive and ubiquitous computing adheres to the ”anybody, anywhere, anytime” concept of user access to information and services around the network. This concept though incorporates a significant degree of security concerns, since especially for knowledge management systems, access to information should be controlled by access policies. Special consideration is necessary when conflicting security policies from different administrative domains arise. In this case access to knowledge information will be the result of coordination and negotiation procedures that need to be defined in a concrete way.

IV. System Design Issues

This paper addresses the issue of knowledge management in ubiquitous environments by proposing a system that will act upon the pervasive realm and handle all the corresponding operations that need to be performed. Based on the nature of the underlying networks, which mostly follow the ad hoc paradigm, we cannot undertake the traditional approach for knowledge management systems, where - in the general case - knowledge sources would register themselves on an appropriate service and users would query the service to gain access to the sources. Security considerations are handled by the service, by means of authenticating users to access only the knowledge information that they have clearance to do. This is the most common approach and it is push-based. Alternatively, in a pull-based approach, users may request, by flooding the network, a particular knowledge source. When a matching knowledge source identifies a request from a user, it applies security policy procedures to establish the user's validity to access its information and proceeds accordingly.

Both of these approaches cannot be considered for the unstable ad hoc environment. The former approach undertakes a centralized notion, where a lookup server is responsible for handling all the knowledge information. This single point of failure and perhaps bottleneck might be suitable for other environments, if for example one would place the server on a powerful device, but this does not stand for pervasive environments, where frequent disconnections are a norm. The system design should therefore take into account the need for robust solutions that are not dependent on the reliability of a sole device, but also consider the need to distribute the reliability perhaps through the whole network. The latter approach requires significant network resources to be consumed, since flooding the network with requests for knowledge sources imposes a large network overhead. Bandwidth though is a scarce resource in pervasive environments, causing this solution inapplicable in such settings.

The system we propose in this paper addresses the issues of providing knowledge management services in pervasive, ad hoc environments.

Our first design principle is to distribute the management of knowledge information available in the domain amongst many devices that will collaboratively act like the lookup server did in the centralized approach. These devices will distribute amongst themselves the load of registering the available knowledge sources and presenting them to users upon queries placed by them. This configuration achieves the avoidance of single points of failure and lessens the consumption of resources on the devices forming the ad hoc network. This approach is considered as hybrid when the two other approaches are considered, since we have a distributed lookup server and devices do not flood the network but only part of it, until they come across a member of the distributed lookup server (DLS).

Our second design principle is to enhance the searching in the ad hoc environment, since the notion of ubiquity incorporates diversity in knowledge sources and their semantic meaning. For that reason a predefined ontology is used that maps knowledge information to its terms, promoting thus interoperability.

Security considerations form the third principle guiding our system design, with emphasis being placed on authenticating users to access the information they are allowed to. Security policies are exploited to cater for this goal.

Fig. 1 demonstrates in a graphical way the three different aspects that characterize the knowledge information and their respective sources in a pervasive environment, as the one MANETs constitute. Borrowing from the database design and warehousing field, we can view knowledge information as information being dependent on the three axes (the aforementioned principles) of a three dimensional cube: the topology of the MANET and the corresponding location of the information, the semantic notion of the information in accordance to the existing ontology and the security clearance and authorization levels. So, single knowledge information is not only characterized by its content, but also from these three diverse aspects

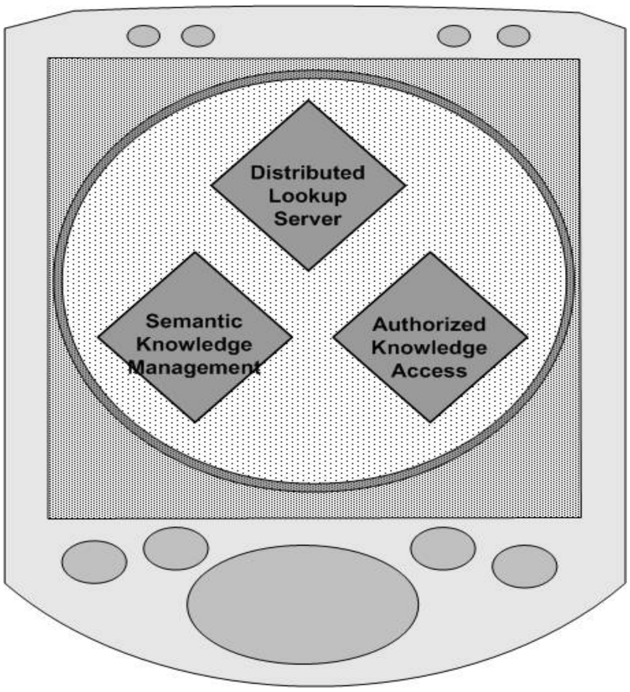
Distributed Lookup Server
Semantic Knowledge Management
Authorized Knowledge Access

Fig. 2. Components of the distributed, mobile-device oriented knowledge management system.

Of particular importance are also the protocols for the knowledge resources distribution upon a multitude of devices in the pervasive realm, namely the ones that form the DLS. There is a trade-off to consider in this case. On one hand one has to consider the special characteristics of ubiquity and act upon them by distributing the load and operations, on the other hand though the management protocols required for this distribution should not consume the scarce bandwidth and computational resources available for the devices in typical pervasive configurations. The next section discusses in detail our proposed approach and we underline the design and implementation decisions we undertook in respect to the system design issues raised in this section.

V. SYSTEM ARCHITECTURE AND IMPLEMENTATION

As we have already mentioned, the system we propose for knowledge management in ubiquitous environments is distributed and its three main parts are depicted in Fig. 2. In the following sections we delve into more details regarding the system components and provide justifications for our choices.

A. *Distributed Knowledge Information Lookup Server*

The concept behind the knowledge information lookup server distribution is to build an overlay management network on top of the ad hoc network and place upon the nodes that form that overlay the distributed lookup servers. These will communicate amongst them and exchange information regarding the knowledge information they handle. Each knowledge source on the pervasive environment will register its

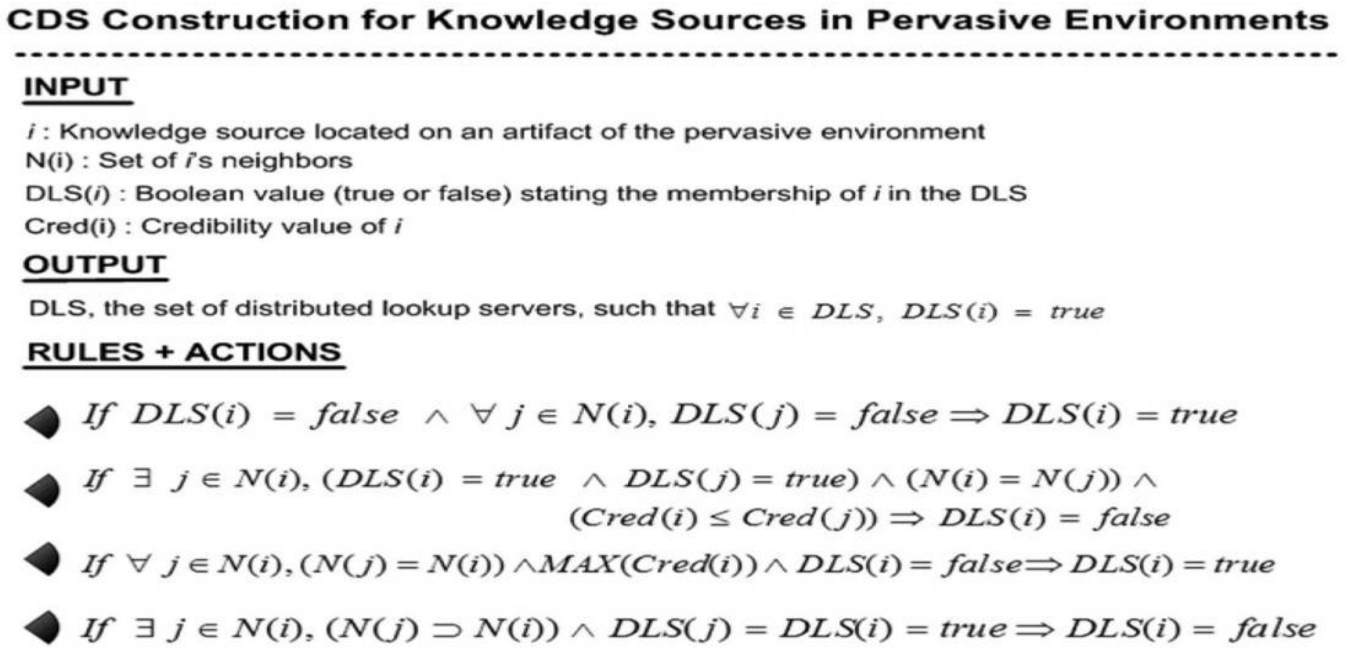

CDS Construction for Knowledge Sources in Pervasive Environments

INPUT

i : Knowledge source located on an artifact of the pervasive environment
N(i) : Set of i's neighbors
DLS(i) : Boolean value (true or false) stating the membership of i in the DLS
Cred(i) : Credibility value of i

OUTPUT

DLS, the set of distributed lookup servers, such that $\forall i \in DLS,\ DLS(i) = true$

RULES + ACTIONS

- *If* $DLS(i) = false \ \wedge\ \forall\ j \in N(i),\ DLS(j) = false \Rightarrow DLS(i) = true$
- *If* $\exists\ j \in N(i),\ (DLS(i) = true\ \wedge\ DLS(j) = true) \wedge (N(i) = N(j)) \wedge$ $(Cred(i) \leq Cred(j)) \Rightarrow DLS(i) = false$
- *If* $\forall\ j \in N(i), (N(j) = N(i)) \wedge MAX(Cred(i)) \wedge DLS(i) = false \Rightarrow DLS(i) = true$
- *If* $\exists\ j \in N(i),\ (N(j) \supset N(i)) \wedge DLS(j) = DLS(i) = true \Rightarrow DLS(i) = false$

Fig. 3. Construction of CDS based on [10].

information on one of these servers, establishing thus that all the knowledge information in the network is registered and can be reached by a member of the set of distributed servers. The idea of using a virtual backbone in the ad hoc network to be responsible for management decisions and even routing is not new. It has been extensively used as far as reliable routing in MANETs is considered. Using such a technique though for knowledge management has not so far been addressed though.

We can identify three basic steps to this procedure:

1) Selection of the most appropriate artifacts of the network to participate in the distributed servers set and construction and management of the set.
2) Registration of the knowledge information in the set of distributed servers and management of updates and other changes, like relocation of knowledge sources.
3) Attaching semantic meaning to the network and applying security controls to limit the access to the knowledge information scattered around the network (Sections V-B and VI).

We view that the pervasive domain can be mapped on a graph where the nodes are the devices and the links their interconnections. There have been several approaches in the literature on building distributed, collaborative management entities as virtual backbones of the graph [10] [11] [12]. The concept behind these approaches lies on the observation that centralized-management architectures are not suitable for MANETs. The need to distribute the load across the MANET is necessary for both resource constraints and reliability and robustness reasons. Solutions include connected dominating sets (CDS) of the graph, maximal independent sets (MIS) of the graph and even BFS trees. Extended work has been performed in the field and we do not wish to delve into this more, since it is not the focus of our research. In order to construct a CDS of the MANET graph we chose to discard computationally complex solutions based on approximation algorithms of the MCDS problem and solutions based on centralized approaches.

Before proceeding we have to provide some definitions in order to have a holistic understanding.

Definition 1. Every pervasive environment can be mapped on a Graph $G = (V, E)$, where V constitutes the set of artifacts comprising the environment and E is the set of links amongst these artifacts.

Definition 2. In a Graph $G = (V, E)$, a dominating set is a subset D of V such, that every vertex V is dominated by some vertex in D.

Definition 3. The domination number is the minimum size of a dominating set of G.

Definition 4. A graph $G = (V, E)$, is called connected when there exists a path of edges in E to connect every two vertices in V.

Freidman *et al.* have researched the construction of CDS in Mobile Ad Hoc networks. This is the reason we chose to adopt the connected dominating set approach and in particular the algorithm presented in [10] for the CDS construction. The details of the algorithm as adapted to the DLS context are given in Fig. 3. We do not essentially differentiate ourselves from their work rather we port it to the context of knowledge information scattered in the network and knowledge management in general. The ”goodness number” (used in [10]) as far as knowledge sources are concerned refers to the long-term credibility of the sources to deliver information to those who request it. This work on knowledge sources credibility is part of our ongoing research and is based on reputation schemes, we are not going to elaborate further though in this paper. What should be noted though is that this credibility value is also dependent on the capabilities of the mobile nodes (computational) to support the activation of the DLS service. It should not be expected that "thin" nodes would have this service deployed.

Our proposed capability function (CF) exploits the following attributes: memory requirements (MEM), processing power (PP), battery power (BP), mobility ratio (MR) and current load (CL). These 5 variables need to be combined in a single equation, the Capability Function (CF). MEM, PP and BP are obviously proportional to CF while MR and CL are inversely proportional. By assigning weights to these variables in accordance to their significance, we have an initial, simplistic CF equation (1):

$$CF(x) = \frac{(w_1 \times MEM(x)) \times (w_2 \times PP(x)) \times (w_3 \times BP(x))}{(w_4 \times MR(x)) \times (w_5 \times CL(x))} \qquad (1)$$

where $\sum w_i = 1$, $i = 1..5$, and x is the MN.

The main requirement for the CF is to lead to comparable results among mobile nodes. For this reason the various attributes must be demoted in common range values. Space limitations do not allow us to delve into more details on how to achieve this. Equation 1 is used to derive a value for every mobile node that is proportional to its capability of being part of the DLS. Our approach is based on building a relatively stable CDS with the "thickest" nodes according to the CF mentioned, but also takes into consideration the need for maintenance of the CDS due to the inherently unstable MANET nature.

When the CDS of the pervasive realm graph has been constructed the nodes that have been selected to be part of it and are consequently also members of the DLS activate the DLS service. The DLS service is installed in every artifact of the pervasive realm, but it is activated only on the members of the DLS. The next step is to register every knowledge source in the pervasive realm to the DLS, so as queries and knowledge management can take place. By definition and construction of the DLS every knowledge source will have at least one member of the DLS in its one-hop neighborhood. It is to that node that the knowledge source actually registers its existence.

Figure 4 depicts an example pervasive scenario and the corresponding DLS, indicating briefly its functionality. In this particular example all the nodes initially are considered as equals and none of them has activated the DLS service although everyone has a local copy of it installed. At a given time the whole process begins and every node calculates its own value based on the previously described capability function. In the scenario displayed in Fig. 4 nodes B, C, E and G are selected to form the DLS for the whole MANET based on their positioning and their capabilities. Each of these devices has certain local knowledge information to advertise to the other nodes. This occurs by registering these information, bound with its own identification information to the Distributed Lookup Server formed by the 4 nodes B, C, E and G. These nodes communicate to each other and exchange relevant information so that collectively they hold a complete understanding of which knowledge source/information exist in the MANET. Every node wishing to access such information can do so by querying the DLS service of its DLS neighbor. By definition there will always be a 1-hop neighbor that is an active member of the DLS.

The Link Expiration Time (LET) amongst two nodes can be in a probabilistic way calculated. Techniques to achieve that include sophisticated solutions incorporating intrinsic measurements from GPS receivers, accelerometers and synchronized clocks, as well as probabilistic methods based on the past history movement of the nodes. In our case we adopt the LET calculation based on the transmission power samples measured from packets received from a node's neighbors [13] [14] [15]. Every node performs this action and has therefore an understanding of when a link with one of its neighbors is about to break. Members of the DLS use this information when they realize they are going to be breaking the CDS. In this case, the node that identifies the prospective link expiration issues a flood message across the network stating this situation and requesting the reconstruction of the DLS, as described earlier.

Obviously, this solution would not scale since it would require significant bandwidth overhead to reconstruct the DLS every time a single link breaks. We thus chose to act in two phases. If the number of link breakages predicted is larger than a particular threshold (Max_Link_Num) then the aforementioned procedure takes place (in the prototype implementation of our system this value is determined as 10% of the total number of nodes. We have experimented with other threshold values as well, ending up with the value of 10% as the most suitable to reach equilibrium in the stability/performance tradeoff). When the link breakages are below the threshold value then another approach is undertaken that does involve network wide reconfigurations. The node that is about to leave finds one or more of its neighbors that collectively have the same set of neighbors and instructs them to join the DLS and activate the appropriate service. The remaining members of the DLS are also informed. Every new member of the DLS also inherits the registered knowledge sources list from the departing node.

B. Semantic Searching of Knowledge Information

To further enhance the knowledge management and produce more efficient results in terms of responding to users' queries on specific knowledge information we further propose maintaining a number of virtual overlay networks in respect to branches of a predefined ontology used for semantically enriching the pervasive environment. Each branch of the ontology will have knowledge information from various sources in the pervasive domain associated with it. The corresponding virtual overlay network will be built on top of the DLS and include only those artifacts that have the related knowledge information registered to them. In this way a user query will first be mapped on the ontology and then it will be flooded to a significantly smaller number of artifacts in the pervasive environment.

The user, holding a device mapped on the graph of the pervasive environment will query upon particular knowledge information. This query will be relayed to the next hop

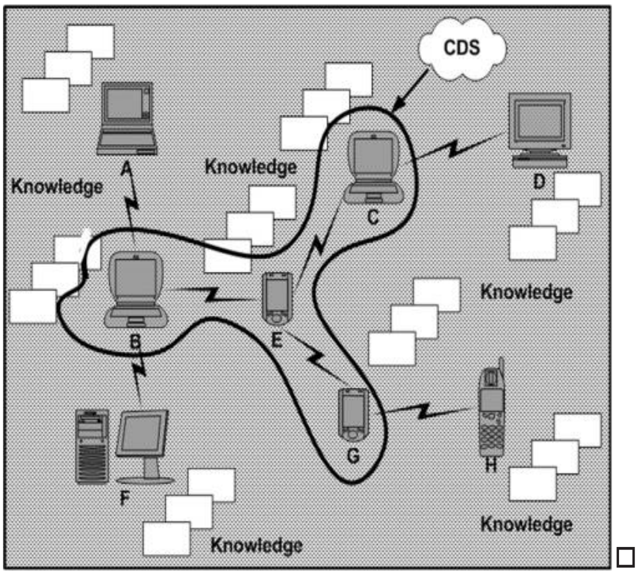

Fig. 4. Example pervasive scenario with respective CDS.

neighbor that is member of the DLS. The next step in the general case would be to flood the query among the nodes of the DLS and wait for the appropriate replies to forward back to the originating querying user. Pervasive environments can grow significantly in size, so that even the use of backbone management bodies like DLS cannot significantly assist in scaling down the increased complexity. The use of the ontology is of great significance. Major categories of the ontology are distributed between nodes of the DLS forming smaller virtual semantic sets of nodes from the nodes of the DLS. Figure 5 describes possible virtual semantic groups overlay on top of the DLS.

For example a subset can refer to information regarding e-government (Fig. 5). All the relevant information sources will be registered with nodes of this overlay network and all relevant queries will be forwarded to this group of nodes. Semantic groups can be overlapping in terms of participating nodes. All nodes of the DLS are aware of the semantic groups their neighbors belong to and where to locate the nearest member of a particular semantic group. This is accomplished by referencing a specially formatted list like the one depicted in Fig. 6 (coded in XML).

In order to create and maintain the virtual semantic groups we exploit an approach similar to the one described for the DLS itself. Various CDS of the DLS are used to place the related service. The semantic service is an extension to the DLS service that performs the semantic categorization of queries and their mapping on the used ontology. To avoid repetition we do no delve into much detail regarding the construction and maintenance of the semantic overlay of the pervasive environment.

VI. Access Control Enforcement

A. Access Control Models

Enforcing authorized access to knowledge assets dispersed around the pervasive environment is a difficult task, and many issues both from a design as well as from technical perspective have to be resolved. Pervasive environments are being characterized by extensive number of users, making in general security management a complex issue. Two types of systems can be distinguished according to the security model adopted:

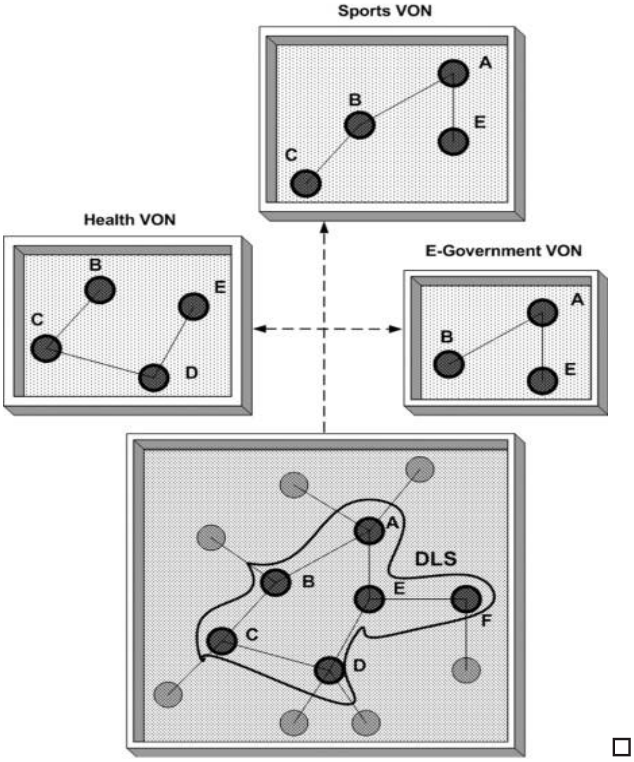

Fig. 5. The DLS of the distributed knowledge environment and some indicative Virtual Overlay Networks (VON) based on the ontology semantics.

1) Trust based systems. The notion of trust is introduced mainly in complex, non-hierarchical systems without explicitly stated organizational policy. The authorization of a transaction is based on the degree of trust associated to the requesting user, which can be obtained by questioning the user's previous activity.
2) Autonomous systems, with well formed security policy and well defined organizational structure. Under this framework, users are associated to roles and certain privileges are assigned to them.

We will confine ourselves to the second category and we will consider every VON as a separate domain maintaining its own set of users, roles and assignment of privileges to the domain's knowledge assets. More specifically we will comply with the basic principles of the Role Based Access Control Model (RBAC), while we assume that in accordance with the conceptual division of the overall infrastructure to different VONs, each domain (VON) can be considered as independent and maintain its own security policy. In addition, each policy may define a number of roles serving the domain's specific access control demands. The basic notions of the RBAC model are: users, roles and permissions. A user represents a human entity or an autonomous agent. A role is associated

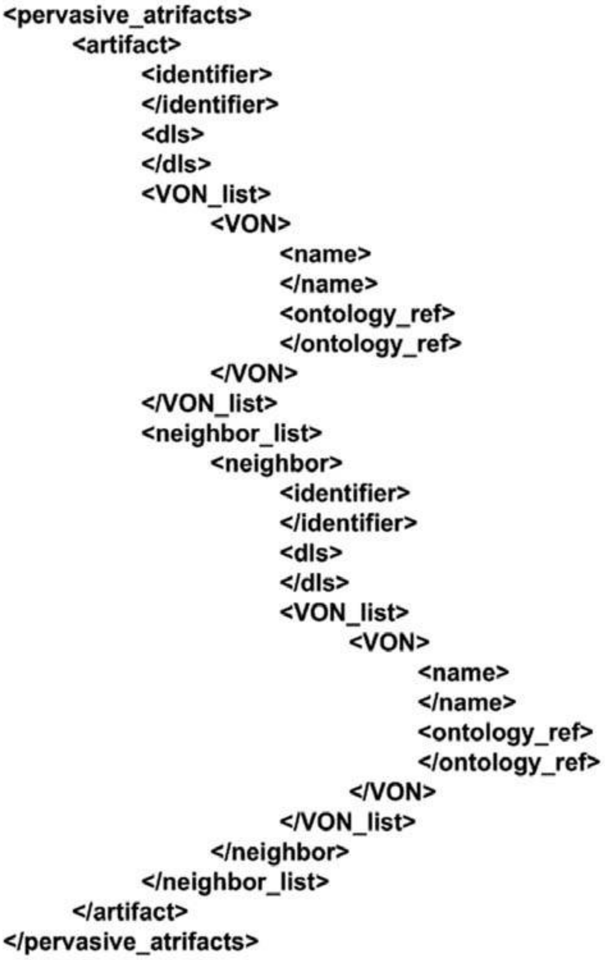

```
<pervasive_atrifacts>
    <artifact>
        <identifier>
        </identifier>
        <dls>
        </dls>
        <VON_list>
            <VON>
                <name>
                </name>
                <ontology_ref>
                </ontology_ref>
            </VON>
        </VON_list>
        <neighbor_list>
            <neighbor>
                <identifier>
                </identifier>
                <dls>
                </dls>
                <VON_list>
                    <VON>
                        <name>
                        </name>
                        <ontology_ref>
                        </ontology_ref>
                    </VON>
                </VON_list>
            </neighbor>
        </neighbor_list>
    </artifact>
</pervasive_atrifacts>
```

Fig. 6. XML representation of neighborhood information.

with the execution of a specific task, while a collection of permissions are assigned to each role, enabling the fulfillment of the obligations associated with such a task. To extend the support for the least privilege principle (that allows to a user the minimum privileges necessary to fulfill a task), sessions are introduced. A complete RBAC model includes the following variables and functions:

- The sets U (users), R (roles), P (permissions) and S (sessions)
- User to role assignment $UA \subseteq U \times R : U \to 2^R$
- Permission to role assignment $PA \subseteq P \times R : R \to 2^P$
- A mapping of sessions to a single user assignment $US : S \to U$
- A mapping from sessions to the set of roles associated with each session $S \to 2^R$
- A partial ordering $RH \subseteq R \times R$, represented by the symbol: $\geq$, which defines role hierarchy. $R_1 \geq R_2$ implies that R_1 inherits permissions from R_2.

RBAC has become a dominant security model due to its flexibility and due to the fact that it reflects organizational hierarchy; moreover, its parameters can be easily codified. For this purpose, several RBAC security policy representation languages have emerged, ranging from formal, graphically annotated to expressive full-scale policy management systems with software tools support. We do not intend to create a new policy representation language. Our work focuses on enabling the coalition of autonomous pervasive systems, where each one retains its own security policy. In fact, there is no restriction that all the domains should follow the same policy language; the only requirement being adherence to the RBAC principles.

B. Policy languages-extensions to the existing models

Several modifications and extensions to basic RBAC are necessary to adjust it to multi-domain environments. The nature of pervasive environments and the unstable character of wireless connections pose additional restrictions and demand flexible modifications, incorporated in our framework, that:

- Are pre-settled by the system administrator, and that the user must specify in order to be granted authorization to activate a role (parameters domain specific mainly, that correlate a user with a specific domain)
- Enable time periodicity, for example allow access within pre-specified time-intervals
- Enable context based role-assignment and role correlation for roles specified in different domains.

For this reason, at the (lower) access control enforcement level, authorization is handled through the utilization of the Extensible Access Control Markup Language (XACML), which enables context based decisions and supports role based authorizations [24]. A basic characteristic of XACML is that it enables role based policy definition as well as it supports time-restricted based activation. XACML is a policy language that supports prohibitions, obligations, and resolution of conflicts. Its expressiveness and XML (Extensible Markup Language) codification support allow its integration on a variety of environments, such as web-service based environments, distributed autonomous systems, and with modifications like the ones discussed in our approach it can also be applied to pervasive environments. Among XACML's strong points, are:

- It is standardized and it is open, allowing extensions that enable interoperation between various platforms.
- It is codified in (XML) which tends to dominate as codification standard and is operating system independent.
- It allows extensions as to support the needs for a variety of environments.
- It allows context based authorization, which is a big advantage.

The basic operational principles behind the XACML authorization process rely on the following entities (Fig. 7): The Policy Enforcement Point (PEP) which grants access to roles, the Policy Decision Point which reasons over a specific access request after evaluating the requestor's credentials and accordingly the request according to the available policy.

In more detail the overall operation of this pervasive-adjusted authorization framework, functions as follows:

The policy administrator is responsible for editing the policy and making it available for the domain, through the Policy Decision Point (PDP). The concrete policy is not necessary to be replicated in every node, but its sub-modules may be spread among the nodes which collaboratively act as a distributed PDP. When a request for a resource appears (Fig. 6, action

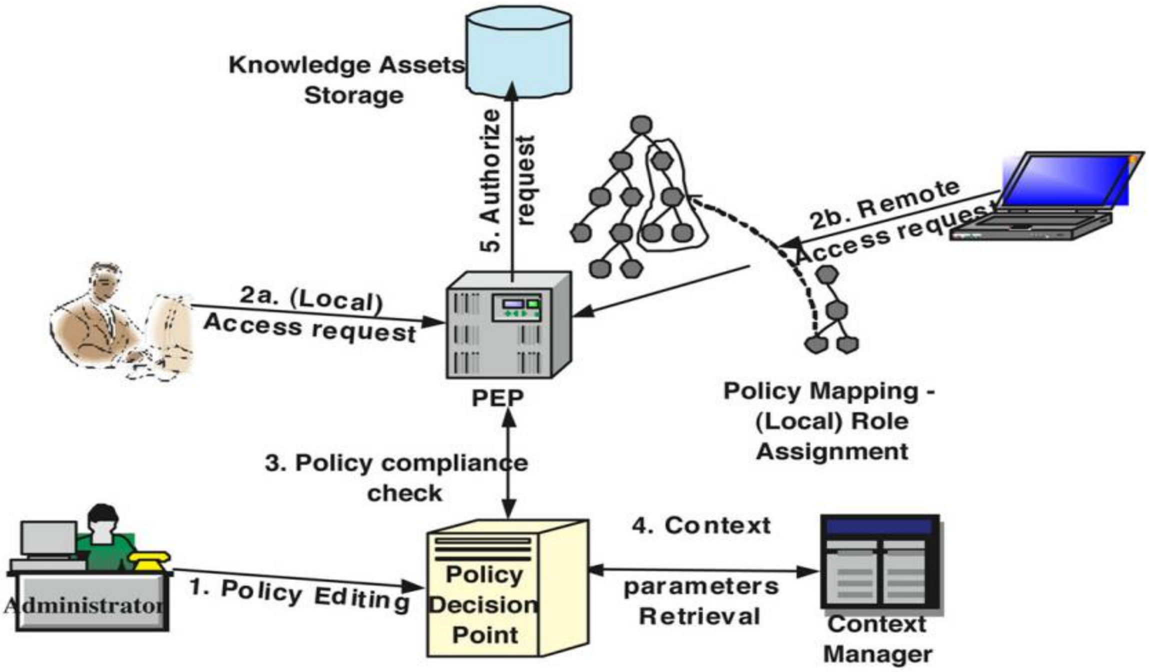
Knowledge Assets Storage
5. Authorize request
2b. Remote Access request
2a. (Local) Access request
PEP
Policy Mapping - (Local) Role Assignment
3. Policy compliance check
4. Context parameters Retrieval
Administrator
1. Policy Editing
Policy Decision Point
Context Manager

Fig. 7. Authorization process in distinct steps.

2a), it has to be validated for its consistency with the local security policy prior to its execution. Accordingly, each request (from the same or from remote domain) is directed to the Policy Enforcement Point (PEP). The request is constructed in an appropriate XML message and directed to the Policy Decision Point (PDP). Prior to the validation of the request, the context handler is sending additional subject, resource, action and environment attributes to the PDP. Accordingly, the request is validated from the PDP and a response message is sent to the policy enforcement point (PEP), which handles the details for providing authorization to the requester.

Taking into account the researched environment, we have extended the XACML authorization scheme by deploying redundant PDP (Policy Decision Point) and PEP (Policy Enforcement Points) entities in the network – instead of only one – to cope with instability issues under the pervasive infrastructure. So instead of a single, centralized PDP, we share this responsibility among the members of the CDS, in the same way that the members of the DLS act collectively as a conceptual centralized server. Several redundant PEPs are also deployed, ensuring that prior to sharing a knowledge asset, the requester's privileges are compatible with the predefined domain's policy.

C. Remote domain's access control enforcement

In our implementation we defined a set of four roles with varying authorization privileges, for two medical domains: For the first we define WardDirector, Surgeon, TraineeSurgeon and Nurse. For the second we define WardManager, GP (General Practitioner), TraineeDoctor and Nurse. Our aim is twofold: first to enable the autonomous security management within each VON separately and second for certain selected roles of one domain within the second domain's coverage area. For the first research aim, the concept of the distributed PEP/PDP has been defined. The second target can be achieved through the concept of policy mappings [22] [25]. Policy mappings define correspondence of roles of other domains to roles of a target domain. Consider the following scenario: Doctors who are assigned to work twice a week as surgeon's in the emergency department of the Hospital, are also working two more days as general practitioners. When an emergency appears it is crucial to be able to retrieve a patient's medical history record collected from his previous possible visits to the hospital's general department. If the two hospital departments can be considered as independent VONs, then it is necessary to define a mechanism to define a role correspondence. The necessary information can be stored -together with the policy- at the (distributed) PDP.

In the following we discuss two important issues, namely preserving privacy and storing policies. Special consideration is necessary due to the inherent limited resources of the mobile devices that usually participate in pervasive environments.

1) *Privacy preservation.* A lot of portable devices nowadays enable user authentication using PKI (Public Key Infrastructure) keys. User authentication at device level is enabled by allowing the user to a PIN identifier, while the private key can be carried in a secure removable media, which enables the matching of the holder's identity with that of the owner. Even though computational resources are a burden for handheld devices, many manufacturers support 128-bit symmetric encryption at a hardware level. It is possible to encrypt all necessary data, while the identification of both communicating parties can be verified through digital certificates.
2) *Policy storage.* There are various approaches relative to policy storage and enforcement on handheld wire-

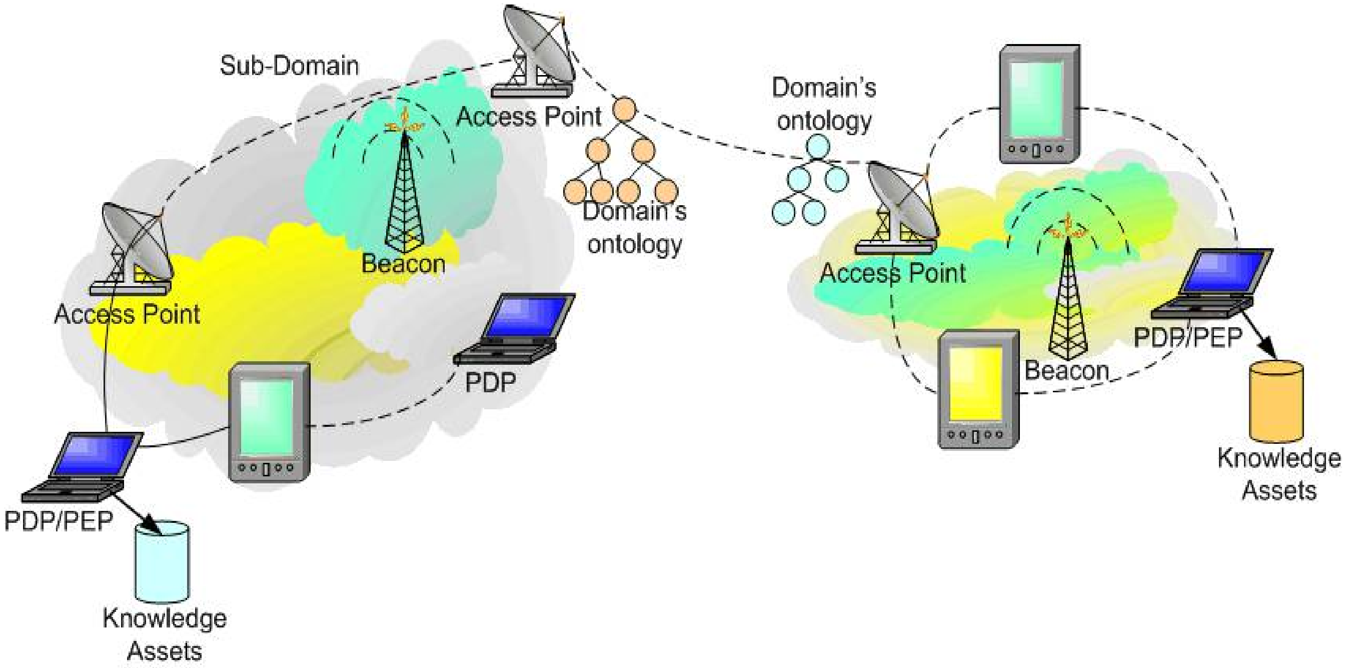

□

Fig. 8. Distributed Access Control Architecture. The presence of a beacon on each domain can notify the device about the local policy enforced and direct easier user-authentication.

less devices. According to [29], policies can be stored on smart cards while the device can monitor all the time whether the smartcard has been removed from the system or not. In our approach not all the devices need to be aware of the organizational policy, since not all the devices would enable provision of services for users or devices. The policy is maintained at the distributed PDP which reasons about a user's access privileges over shared assets – in our previous scenario medical records that are kept only in distinct locations and where security controls are applied prior to any authorization. Instead of making the policy available to any device, we just make it available on critical points, which store sensitive knowledge information, like medical data and where any misuse can lead to breaches of confidentiality and every access request is accompanied by provision of requester's credentials that are therefore evaluated according to the organizational policy. Users through their handheld (or portable in general) devices can roam from one domain to the other and request from the domain to which they belong or a remote domain access to knowledge assets. The presence within a specific domain can be identified by providing the devices with a number of certificates that enables them to identify the domain in which they are at any moment by receiving digitally signed messages transmitted from a beacon [20] [23] (Fig. 8). Each device is also provided with a user-indicant digital certificate, which accordingly is associated with a domain role. Therefore, the devices are aware constantly according to the received messages whether they are in the default domain, whether they are present in a collaborative domain where a local role can be assigned to them through the pre-defined policy mappings, or whether they are present in a non-trusted zone.

VII. Implementation Details

In order to test the proposed system's validity we proceeded in building a prototype implementation. We are also planning on testing the system's efficiency by simulation experiments; this though remains on our future agenda. We chose to first evaluate the functionality of the system in a realistic scenario and then test its scalability and dependence on various parameters by simulation, as indicated in [16].

The prototype has been implemented using the Java programming language and in particular J2ME since we are targeting mobile and pervasive environments. Being lightweight is a necessity so the CLDC (Connected Limited Device Configuration) was selected. Another reason for using Java was to grasp interoperability since in pervasive environments a variety of devices in terms of hardware/software combinations is present. The communication protocol amongst devices regarding the protocols described and the knowledge querying are implemented using the SOAP web services protocol. We understand that a more lightweight solution like XML-RPC or even CORBA-based approaches would yield faster results, yet we chose SOAP since we want to exploit its advanced features. As we have already mentioned access and authorization rules have been built based on XACML and XML is used to keep records of neighbors and their properties (Fig. 6). For the semantic ontology we used part of the Open Directory Project (dmoz.org) to serve our needs.

Each device has both the functionalities of being a simple knowledge source in the pervasive environment and of being member of the DLS. The results of the DLS construction algorithm as described earlier will indicate which of the two will be activated for each device. To be precise, the DLS members also have the simple knowledge source activated since they too host knowledge information. The memory footprint for the simple and the advanced services are minimal, requiring 21Kb and 52Kb of memory respectively. These numbers do not include the memory footprint of the J2ME platform of course.

The platform has been tested on our experimental hardware platform that consists of 10 personal computers. In order to be able to emulate mobile ad hoc scenarios we used the MobiEmu emulation environment [17]. We tested the platform with various mobility scenarios derived from the ns2 simulator and the initial findings prove its robustness and the ability to accurately locate and retrieve knowledge information in pervasive environments even in cases of high mobility. The links between the computers used for the pervasive scenario emulation are wireless, in particular based on 802.11b. We understand that the scenarios lack validity as far as the used devices are concerned since personal computers are more powerful from the devices that normally exist in pervasive domains; we plan to ameliorate that in the future with the use of devices with limited resources.

VIII. Conclusions

In this paper we presented our approach on providing secure and robust knowledge management solutions in pervasive environments based on the ad hoc networking paradigm. We introduced the concept of Virtual Overlay Networks in order to adjust the demands of our architecture to the specific low resource specifications and the instability characterizing the ad-hoc infrastructure. We also introduced a policy-based solution that enables the cooperation between different VONs, which act as autonomous sub-domains. A prototype implementation has been built and tested on our experimental setting and the initial findings are promising, achieving adequate degrees of robustness and access to requested information. Our future work includes extensive testing of the platform in real scenarios and simulation of the DLS performance and efficiency in terms of network parameters.

Acknowledgment

This work was co-funded from the E.U. by 75% and from the Greek Government by 25% under the framework of the Education and Initial Vocational Training Program - Archimedes.

References

[1] C. E. Perkins. *Ad Hoc Networking*. Addison Wesley Longman, 2001.

[2] S. Corson and J. Macker. *Mobile Ad hoc Networking (MANET): Routing Protocol Performance Issues and Evaluation Considerations*. IETF RFC 2501.

[3] S. Chakrabarti and A. Mishra. QoS Issues in Ad Hoc Wireless Networks. *IEEE Communications Magazine*, **39**(2): 142–148, 2001.

[4] J.-P. Hubaux, T. Gross, J.-Y. Le Boudec and M. Vetterli. Toward Self-Organized Mobile Ad Hoc Networks: The Terminodes Project. *IEEE Communications Magazine*, **39**(1): 118–124, 2001.

[5] Z. H. Haas *et al.* Guest Editorial. *IEEE Journal on Selected Areas of Communication, Special Issue on Wireless Networks*, **17**(8): 1329–1332, 1999.

[6] I. Chlamtac, M. Conti and J. J.-N. Liu. Mobile ad hoc networking: imperatives and challenges. *Ad Hoc Networks*, **1**(1): 13–64, 2003.

[7] S. Giordano. *Handbook of Wireless Networks and Mobile Computing*. John Wiley & Sons., New York, 2002.

[8] M. S. Corson, J. P. Maker and J. H. Cernicione. Internet-based mobile ad hoc networking. *IEEE Internet Computing*, **3**(4): 63–70, 1999.

[9] M. Weiser. The Computer for the 21st Century. *Scientific Computer*, September 1991.

[10] R. Friedman, M. Gradinariu and G. Simon. Locating cache proxies in MANETs. *ACM MobiHoc 2004.*

[11] U. Kozat and L. Tassiulas. Network layer support for service discovery in mobile ad hoc networks. *IEEE Infocom 2003.*

[12] P.-J. Wan, K. M. Alzoubi and O. Frieder. Distributed construction of connected dominating set in wireless ad hoc networks. *IEEE Infocom 2002.*

[13] A. Agrawal, D. K. Anvekar and B. Narendran. Optimal Prioritization of Handovers in Mobile Cellular Networks. *Proceedings of the Annual IEEE International Symposium on Personal, Indoor and Mobile Radio Communication (PIMRC)*, The Hague, Netherlands, September 1994, pp. 1393–1398.

[14] B. Narendran, P. Agrawal and D. K. Anvekar. Minimizing Cellular Handover Failures Without Channel Utilization Loss. *Proceedings of the IEEE Global Telecommunications Conference (GLOBECOM)*, San Francisco, CA, December 1994, pp. 1679–1685.

[15] W. Su, S.-J., Lee and M. Gerla. Mobility prediction and routing in ad hoc wireless networks. *International Journal of Network Management*, **11**(1): 3–30, 2001.

[16] C. Tschudin, P. Gunningber, H. Lundgren and E. Nordstrom. Lessons from experimental MANET research. *Ad Hoc Networks, Special Issue on Ad Hoc Networking for Pervasive Systems*, **3**(2): 221–233, 2005.

[17] Y. Zhang and W. Li. An Integrated Environment for Testing Mobile Ad-Hoc Networks. *ACM MobiHoc 2002.*

[18] K. Kida and H. Shimazu. Ubiquitous knowledge management - enabling an office-work scheduling tool for corporate knowledge sharing. *Proceedings of IEEE Workshop on Media Networking*, 2002.

[19] I. Nonaka. A Dynamic Theory of Organizational Knowledge Creation. *Organization Science*, **5**(1): 14–37, 1994.

[20] R. Choudhri, L. Kagal, A. Joshi, T. Finin and Y. Yesha. PatientService: Electronic Patient Record Redaction and Delivery in Pervasive Environments. *Fifth International Workshop on Enterprise Networking and Computing in Healthcare Industry (Healthcom 2003)*, Santa Monica, June 2003.

[21] A. Seleznyov, A. Mohamed and S. Hailes. ADAM: An agent-based Middleware Architecture for Distributed Access Control. *Twenty-Second International Multi-Conference on Applied Informatics: Artificial Intelligence and Applications*, 2004.

[22] P. Belsis, S. Gritzalis and S. Katsikas. A Scalable Security Architecture enabling Coalition Formation between Autonomous Domains. *Proceedings of the 5th IEEE International Symposium on Signal Processing and Information Technology (ISSPIT'05)*, December 2005, Athens, Greece, IEEE Press.

[23] P. Belsis and S. Gritzalis. Security Control Schemes for Pervasive Medical Environments. In: P. Georgiadis, S. Gritzalis and Y. Marias (Eds.), *Proceedings of the 1st IEEE International Conference on Pervasive Services ICPS 2005 - Workshop on Security, Privacy, and Trust in Pervasive and Ubiquitous Computing (SecPerU'05)*, July 2005, Santorini, Greece, Diaylos Press.

[24] Organization for the Advancement of Structured Information Standards (OASIS), XACML Extensible access control markup language specification 2.0, OASIS Standard (available at `http://www.oasis-open.org`, accessed February 2005).

[25] J. B. D. Joshi, R. Bhatti, E. Bertino and A. Ghafoor. Access Control Language for Multi-Domain Environments. *IEEE Internet Computing*, **8**(6): 40–50, 2004.

[26] Z. J. Haas, J. Deng, B. Liang, P. Papadimitratos and S. Sajama. Wireless Ad Hoc Networks. In: J. Proakis (Ed.), *Encyclopedia of Communications*, John Wiley, December 2002.

[27] M. Kumar, B. Shirazi, S. Das, B. Sung, D. Levine and M. Singhal. PICO: A Middleware Framework for Pervasive Computing. *IEEE Pervasive Computing*, **2**(3): 72–79, 2003.

[28] A. Malatras, G. Pavlou, P. Belsis, S. Gritzalis, C. Skourlas and I. Chalaris. Secure and Distributed Knowledge Management for Pervasive Environments. *IEEE International Conference on Pervasive Services*, Santorini, Greece, 2005.

[29] W. A. Jansen, T. Karygiannis, S. Gavrila and V. Korolev. Assigning and Enforcing Security Policies on Handheld Devices. *Proceedings of the Canadian Information Technology Security Symposium*, May 2002.

Apostolos Malatras (A.Malatras@surrey.ac.uk) is a researcher and PhD candidate at the Center for Communication Systems Research, Department of Electronic Engineering, University of Surrey, United Kingdom, where he is an active member of the Networks Research Group. He holds a Diploma in computer science from the University of Piraeus, Greece and a M.Sc. degree in Information Systems from the Athens University of Economics and Business, Greece. His research interests focus on software engineering, context awareness, knowledge management, network management, networking, and service engineering, policy-based systems, multimedia service control, programmable networks, and communications middleware

George Pavlou (G.Pavlou@surrey.ac.uk) is a professor of communication and information systems at the Center for Communication Systems Research, Department of Electronic Engineering, University of Surrey, United Kingdom, where he leads the activities of the Networks Research Group. He holds a Diploma in engineering from the National Technical University of Athens, Greece, and M.Sc. and Ph.D. degrees in computer science from University College London, United Kingdom. His research interests focus on network management, networking, and service engineering, covering aspects such protocol performance evaluation, traffic engineering, quality of service management, policy-based systems, multimedia service control, programmable networks, and communications middleware.

Petros Belsis (pbelsis@aegean.gr) is a researcher and PhD candidate at the Information and Communication Systems Engineering Department, University of the Aegean, Greece. He holds a Diploma in Physics, from the University of Athens, Greece, a Diploma in Computer Science, from the Computer Science Department of Technological Education Institute, Athens, Greece and M.Sc. degree in Information Systems from the Athens University of Economics and Business, Greece. His research interests focus on Distributed Knowledge Management Systems, policy based systems and security in Distributed Environments.

Stefanos Gritzalis (sgritz@aegean.gr) is an Associate Professor, Head of the Department of Information and Communication Systems Engineering, University of the Aegean, Greece and Director of the Info-Sec-Lab. He holds a BSc in Physics, MSc in Electronic Automation, and PhD in Informatics all from the University of Athens, Greece. His published scientific work includes several books on Information and Communication Technologies topics, and more than 100 journal and national and international conference papers. The focus of these publications is on Information and Communication Systems Security.

Christos Skourlas (cskourlas@teiath.gr) is a professor of Databases at the Computer Science Department of Technological Education Institute, Athens Greece. He holds BSc in Mathematics and PhD in Informatics from the University of Athens, Greece. His research interests focus on Information Retrieval, Knowledge Management, Multilingual Systems, Disambiguation and Natural Language Processing., Medical Informatics.

Ioannis Chalaris (ixalaris@teiath.gr) is a professor of Software Engineering at the Computer Science Department of Technological Education Institute, Athens Greece. He holds a Diploma in Physics, from the University of Athens, and a Phd in Informatics from the University of Berlin, Germany. His research interests focus on Software Engineering, Software Quality and Assurance and Business Modeling, e-Government.

This article has been cited by:

1. Petros Belsis, Christos Skourlas, Stefanos Gritzalis. 2013. A Wireless System for Secure Electronic Healthcare Records Management. *International Journal of Advanced Pervasive and Ubiquitous Computing* **5**:4, 16-32. [Crossref]
2. Dimitris Vassis, Petros Belsis, Christos Skourlas. Secure Management of Medical Data in Wireless Environments 427-432. [Crossref]
3. Petros Belsis, Christos Skourlas, Stefanos Gritzalis. 2011. Secure Electronic Healthcare Records Management in Wireless Environments. *Journal of Information Technology Research* **4**:4, 1-17. [Crossref]
4. Petros Belsis, Stefanos Gritzalis, Sokratis K. Katsikas. 2007. Partial and Fuzzy Constraint Satisfaction to Support Coalition Formation. *Electronic Notes in Theoretical Computer Science* **179**, 75-86. [Crossref]
5. Petros Belsis, Christos Skourlas, Stefanos Gritzalis. Secure Electronic Healthcare Records Distribution in Wireless Environments Using Low Resource Devices 697-712. [Crossref]
6. Petros Belsis, Christos Skourlas, Stefanos Gritzalis. A Wireless System for Secure Electronic Healthcare Records Management 1509-1525. [Crossref]
7. Petros Belsis, Christos Skourlas, Stefanos Gritzalis. Secure Electronic Healthcare Records Management in Wireless Environments 202-219. [Crossref]
8. Petros Belsis, Christos Skourlas, Stefanos Gritzalis. Secure Electronic Healthcare Records Distribution in Wireless Environments Using Low Resource Devices 247-262. [Crossref]