

Privacy and Location-Aware Service Discovery for Mobile and Ubiquitous Systems

Leonardo Galicia Jiménez, J. Antonio García-Macías

Computer Science Department
CICESE Research Center
Ensenada, Baja California, México
{lgalicia, jagm}@cicese.mx

Abstract. Wireless networks play a major role in allowing the deployment of ubiquitous distributed systems. In these networks, service discovery should not only allow finding available networked services, but should also take into account the physical proximity of the entities requesting these services. However, physical proximity is not a sufficient criteria for service search and selection, as close attention should be paid to privacy issues. In this paper we present the design issues that should be considered in order to properly support service discovery based on the physical location of clients; these issues are taken into account for the proposal of an architecture for context-aware distributed systems that consider privacy concerns.

1. Introduction

As mobile technologies get cheaper and more widely available, demand for networked services has rapidly increased. Therefore, it is very important to provide users the possibility to find and make use of the services available in the network. Service discovery has been addressed by many commercial and academic entities and many proposals for solving related problems have been made, for instance SLP (Service Location Protocol) **Erreur ! Source du renvoi introuvable.**Jini [4] and UPnP (Universal Plug and Play) [5] However, some issues have been left out from these proposals, such as the selection of services based on their physical proximity; this particular issue is important because it is needed for supporting location-based and context-aware systems [3]

In order to support service discovery taking into account the physical location of the clients, it is necessary to have mechanisms to estimate this location and based on that search for the services. This type of search has the purpose of finding those entities –static or dynamic– that are within the space defined as proximal. Such entities can be devices or persons that the user wishes to interact with. However, when these mechanisms are used for finding persons, it is important to consider crucial factors related to the privacy of these persons. It is therefore of the utmost importance to include mechanisms that allow the protection of such privacy.

In this paper we present several aspects that need to be considered when designing systems that allow the discovery of services based on physical proximity; an

architecture that includes these factors is presented and a proof-of-concept prototype is implemented. In section 2 we present an overview of some of the related work, and section 3 deals with the design issues for developing this architecture, taking into account experiences gathered in public health institutions; privacy is an important issue here. Then in section 4, the architecture and its main components are presented. Some functional tests for the architecture are presented in section 5. Finally, section 6 offers some concluding remarks and outlines future developments.

2. Related Work

Previous work has been made regarding the provision of user support in location-based services. However, most lack features to address privacy concerns and to minimize the client's devices in order to automatically access an available network infrastructure and related services. Also, they do not provide the means to control and filter responses to service discovery requests, according to who issues the requests.

Some of these previous works offer a class framework [11] on top of which applications can be developed; others present architectures [10] [12] oriented toward location-based services. Most of these proposals rely on notification of events to the client that registers to a given service; they also use different technologies such as CORBA, RMI, Java Messaging Service (JMS), or Publish/Subscriber mechanisms. Our approach is event-based using the instant messaging and presence awareness (IMPA) paradigm.

Perhaps the proposal that is the most closely related to ours is Splendor [16]. However, there are some notable differences: first, Splendor has provisions for locating users, but we go farther than that and closely integrate local mobility issues by providing representation of physical spaces, proximity models, and other aspects that will be detailed later; second, with our approach, services are extended beyond the traditional device-centric view and even persons can be represented as service providers, according to their role or currently enforced policies; third, we not only rely on technical schemes (e.g., public and private keys, etc.) to provide security and privacy, but also take into account administrative policies defined by the authorities (persons) within certain scopes; also, we base our work on accepted standards such as XMPP, SLP, IETF's Geopriv, etc.

3. Design considerations

As previously stated, current proposals do not appropriately handle service selection based on their physical proximity, and so they do not support very well ubiquitous computing systems. Another important aspect left out by these proposals is privacy issues, which become particularly important when the discovered entities are persons. In this section we analyze these and other issues that should be considered for the design of systems that support service discovery based on physical proximity.

3.1 Local mobility

Some of the so called “knowledge workers” present a high degree of mobility in their daily activities. Local mobility refers to dynamic patterns of mobility that take place close to the worker’s office, or even within a building, when a worker is carrying out her duties, collaborating with colleagues, etc. [1] A clear example of this is the kind of work performed in a hospital [7] [6] [9] as it involves a high degree of mobility of patients, equipment, resources and personnel within the hospital facilities.

3.2 Localization of artifacts and persons

In order to provide location-based service discovery, it is first necessary to have some mechanisms to estimate the location of artifacts and persons through some device. Usually, location of devices is basically categorized as device-centered and network-centered. The first one allows the device, through some mechanism, to estimate its own position; that is, it is only the device and no other which can estimate and know its physical location. Meanwhile, network-based mechanisms require that a different entity within the network perform the estimation of a device’s physical location; this way, when a device want to know its location, it has to consult the entity in the network which is in charge of determining it. Some good examples of device-centered systems are RADAR [13] Cricket [14] and AeroScout¹. In this type of systems the estimation is usually performed by a PDA, mobile phone, or some other type of mobile device. So, in some way the estimation is person-centered, as the devices are normally carried by persons; but only the device is fully in control of the calculations for determining the current position and if the device is turned off for some reason, the location can not be determined. Some examples of network-centered systems are Active Badge [15] Ubisense², and Exavera³. In these types of systems, objects and persons can carry small devices to aid the network in determining their location.

We think that the network-centered model is more appropriate for the type of scenarios that take place in hospitals, which are our focus for technological development. In a hospital environment some artifacts, such as wheelchairs, stretchers, portable EKG equipment, and others, are good candidates to be located. Moreover, the network-centered model allows the possibility of continuous tracking.

3.3 Representation of physical spaces

It is necessary to have a computational model to represent all those artifacts and persons that are moving within a physical space. This model should represent, at least, the physical space, the entities that move within it, as well as those that are static.

¹ <http://www.aeroscout.com>

² <http://www.ubisense.com>

³ <http://www.exavera.com>

There are currently different models that allow the representation of physical spaces [12] including the geometric model, set theory, graphs, and the semantic model. The geometric model includes definitions based on Euclidian geometry, through coordinates in a Cartesian plane; this plane is a direct consequence of cartographic representation, where information and participating entities are superposed on maps, planes or images. The semantic model offers descriptive information about the geometric areas that represent physical spaces. Under these considerations, both the geometric and semantic models are appropriate for the requirements and technological needs identified in the type of environments that we are interested in (*i.e.*, hospitals). These models have been previously used in the projects and commercial systems mentioned above, namely Cricket, RADAR, Exavera, Ubisense and Radianse.

3.4 Proximity model

A fundamental concept for proximity-based service discovery is, not surprisingly, proximity. The key for a correct association between services and physical spaces is the geographical criterion to be used when services are searched based on their proximity. Two models are widely used for the selection of services: the distance-based model and the scope-based model. In the distance-based model, clients select the services that are within a certain distance from the current position. Given that proximity is a relative value and what is perceived as proximal can vary drastically according to the activities being performed by the client, some mechanism should be present to dynamically change the proximity range.

In a scope-based model, each service is associated with a scope that explicitly represents the context of use of the service within a physical space. The client selects those services whose scopes include the location of the requesting client; that is, a client can discover services if it is inside a certain scope, as well as the services. The main characteristic of this model is that the correlation between context and proximity is assured. When services are discovered, no matter what their distance, they have a high probability of being relevant for the requesting client.

We consider that the scope-based model better suits our needs, mainly because it allows the representation of physical sub-spaces as geometric shapes; this is very adequate for indoors environments such as hospitals where the definition of rooms, working areas, etc. is very useful.

3.5 Definition of services

A service is an entity that can be used by a person, a computer program, or any other entity [4] examples of services are files, a storage device, a printer, a server, etc. When service discovery is performed within a physical space or scope, it is not enough to know the available services, but also the persons within it, in order to be able to interact with them. Thus, it is convenient to consider artifacts and persons as services associated to a given scope; it should be noted that some of these entities will not strictly provide a service, but it is necessary to know their attributes (type, role,

etc.), in order to know what kind of entities there are and where they are in a certain space.

3.6 User identification

In a ubiquitous computing environment, the search for services responds to contextual variables such as the location of the user and its identification. While the location allows obtaining those services near to the user, the identification allows the user to access them. This way, a user could make an anonymous request for available services within a scope and she may obtain a list of them, however, it could be possible that more services would exist; this is because some of those services may require additional information about the user (such as her role, an ID, etc.) in order to be possible to discover them. This way, services are discovered not only based on their location or proximity, but also based on who is requesting them.

3.7 Access control to location information

In spite of the benefits that ubiquitous computing environments offer, one of the main barriers for the adoption of related technologies is the privacy concerns of users. A clear example is shown in the results obtained by Intel Research Berkeley through a series of interviews and scenarios [7] where users manifested concerns regarding being located. Recently the IETF⁴, through the Geopriv⁵ working group, has been studying privacy issues that arise when the geographical information of people and resources is used. The focus of the group, as stated in their charter, is “to assess the authorization, integrity and privacy requirements that must be met in order to transfer such information, or authorize the release or representation of such information through an agent”. Then, attention is paid to presence and geospatial information, commonly used in instant messaging systems, location-based services and others. Up to date, the Geopriv working group has generated a couple of RFCs (request for comments) and recommendations aimed at proposing a standard that guarantees the privacy of users. Much of this work is currently under review, but rules and mechanisms have already been defined (represented in XML formats); these control when, to whom, in what place, and under what circumstances geolocation information can be released. The Geopriv group uses formats and architectures previously approved by the IETF, for instance, XMPP (Extensible Messaging and Presence Protocol) or SIMPLE (SIP for Instant Messaging and Presence Leveraging Extensions).

The user should be able to define when (day and time), to whom, where (place), and under what circumstances (status of the user) location information can be released. The mechanisms for this are defined by the Geopriv group, and can be easily applied to artifacts, not for protecting their privacy as this would not make much

⁴ <http://www.ietf.org>

⁵ <http://www.ietf.org/Geopriv>

sense, but for extending service discovery and indicate who and under what circumstances can be discovered. This way, artifacts and persons (viewed more generally as services) can be discovered based not only on their location, but also taking into account who is requesting them.

4. An architecture for privacy and location-aware service discovery

Considering the issues presented in the preceding section, we have designed an architecture that allows the discovery and selection of services based on their physical location and considering privacy concerns. A key component for implementing a functional prototype incorporating the elements of this architecture is SALSAs (Simple Agent Library for Seamless Applications). SALSAs is a framework that provides a set of abstract classes and mechanisms for the development of autonomous agents that represent users and services [7]. SALSAs agents can be executed on PDAs, cell phones, servers, or any other computing device. XMPP (Extensible Messaging and Presence Protocol) is used for communication of XML messages between clients and servers (implemented with Jabber). Basically, a SALSAs agent contains a Jabber server and a subsystem that implements the autonomous behavior of the agent. Many components are present in the agent: a protocol for registration in an agent directory; a Jabber client through which users, and agents representing users and services, interact via XML messages; and finally, a subsystem including components for perception, reasoning and actions.

We outline the main components of the proposed architecture:

- **Location system.** We assume the presence of a system that estimates the current location of entities on a physical space; we have previously given examples of these systems. Then, each entity (person or artifact) is represented as a SALSAs object and gathers its (x,y) coordinates from the location estimation system; of course, this information may vary with time. Each SALSAs agent representing an entity is autonomous and independent from others. Also, each agent resides in a centralized server.
- **Physical space server.** Physical spaces or areas of interest are geometrically represented through polygons. These have a semantic value for the user, such as a floor, a room, a department, etc. These polygons are centrally stored in an entity (the physical space server) which is responsible for translating from geometric coordinates to semantic values more understandable by the final user; so, in a way it acts similarly to a DNS server. The physical space server is represented by a SALSAs agent that is loaded on a different computer than the location system. It is the main entity to be discovered by the rest of the system's entities, as without it the scope of the services is lost.
- **Policies generator.** We have previously discussed the importance of controlling the release of location information, and we have also mentioned the efforts of the IETF Geopriv group in this regard. The policy generator is the entity charged with presenting an interface for editing the policies that control when, where, and

who has the right to discover particular services. This entity can be on a mobile device for allowing editing policies for that device, or it can be centralized for editing the policies of a set of services.

- **Location server.** This server, contrary to the physical space server, is optional. The location server is an entity represented by a SALSA agent, and it has the purpose of containing information about the location of entities that subscribe to this server. When an entity requests a subscription it first gets validated, and after that it will periodically send location information to the server. During the validation process the entities indicate the policies that will restrict access from non-authorized entities. The location server has two auxiliary components: a policy container and a location container. The first one serves as a temporary storage for policies that are sent during the subscription process. The second one stores, information about the location of an entity (which is periodically updated).
- **General purpose SALSA agents.** In the SALSA framework, agents can represent services, devices, and users. Thus, any entity represented by an agent can potentially search for services. A SALSA agent can reside in a user's computer and perform a service search, consume location information from other entities, or have an associated scope and receive requests for service discoveries.

4.1 Control policies for service discovery

These policies are inspired mainly by RFC 3669 of the Geopriv working group, as well as by the policy and common policy working drafts of that same group. RFC 3669 deals with the authorization, integrity and privacy involved in releasing information about the location of users. The Geopriv group is currently exploring how to represent information about location and presence, as well as how to protect this information.

In its Common Policy draft, the Geopriv group defines the base mechanisms for delivering location information in presence messages; these mechanisms, which allow access control form information regarding presence and location of users, can be extended and translated easily to other application domains. These mechanisms define an XML document (figure 1) that represents policies associated to an entity. Requests from entities contain policy rules, and these are checked against the policies defined in the entity that receives the request; if one or more of the rules match, then the location and presence information are released, else it is denied. Each rule is composed by three sections: conditions, actions, and transformations. The conditions section defines all the restrictions that should be satisfied by an entity in order for it to obtain the requested information. The actions section is a set of processes that the user requests to an entity to perform; these actions have not been defined by Geopriv and are meant to be defined at the application (and not at the user) level. The transformation section indicates those modifications that should be made to the location information before being released. For instance, even if a user complies with all the rules imposed by an entity, a transformation could be imposed to reduce the precision of the geographic location, or to indicate the floor where the entity is located

but without saying in what room it is. These mechanisms are used on our proposed architecture.

```
<rule id="f3g44r1">
  <conditions>
    <identity>
      <id>bob@example.com</id>
    </identity>
    <validity>
      <from>2003-12-24</from>
      <to>2003-12-24</to>
    </validity>
  </conditions>
  <actions/>
</rule>
```

Figure 1. XML document showing a rule with the conditions section defined by the Geopriv working group.

4.2 Discovery of the Jabber server

The proposed architecture utilizes a IMPA (Instant Messaging and Presence Awareness) paradigm, so for implementation purposes we used a Jabber system. In fact, the SALSA framework relies on Jabber, so each SALSA agent should have a Jabber identifier and normally has to be manually configured to indicate to which server it will connect. This last point goes against the idea of automatic service discovery. With this in mind, our architecture incorporates the SLP API so SALSA agents can automatically be configured to find the Jabber server to which they should connect. Although other service discovery protocols could have been used, we think that SLP offers an important advantage: it allows the integration with DHCP; this means that a device can automatically send a request to a DHCP server and obtain not only the network parameters for autoconfiguration, but also the information regarding the SLP server to which it should connect.

4.3 Interaction of components

When an architectural component is initialized it sends a multicast request for a SLP server. Once this server is located, the component locates the Jabber server to which it should connect. Then, the SALSA agent representing the component is initialized. The first interaction of a SALSA agent is made with the physical space server in order to obtain all the spaces that it contains. This way, the agent can determine to which space it belongs and can be associated to a certain scope. As its (x,y) coordinates change with time, the agent can detect these changes and translate them to semantic values, meaning that it can determine at all times to which scope it is associated. The scopes have an associated identifier which is also used for joining conference rooms in the Jabber server. This way, all agents within the same scope are also present in the same conference room. A consequence of this design feature is that

service discovery is greatly simplified, as all involved messages are sent to only one conference room (and to all agents within it), and not to all services system-wide.

There are also other characteristics of conference rooms that motivated their use in this architecture. One is that every time that a Jabber entity enters a conference room, a presence message is sent to all entities participating in that room. Likewise, when an entity leaves a room, all entities are notified of it. These way, all SALSA agents within a conference room are aware at all times of which other entities are within a conference room (which in turn represents a scope), guaranteeing the integrity of services within a scope. Another characteristic is that every time a Jabber entity enters a conference room, it enters with a pseudonym, and not with its Jabber-ID. This means that the Jabber-ID for an agent will not be disclosed, thus allowing for complete anonymity. Of course, when an agent leaves one room and enters another, its pseudonym also changes.

4.4 Service discovery

Once a SALSA agent is aware of its scope, and once it is associated to a conference room, it is capable of discovering other agents (representing entities) within that scope. The discovery process can be performed in a proactive or reactive way. When performed proactively, each time a SALSA agent detects that a new entity has entered the room it requests to know the service(s) associated to that entity and this information is stored in its cache. When the entity leaves the room, the SALSA agent eliminates the information from its cache. On the other hand, with a reactive approach, an explicit request for a list of the entities in a room is made. This list reflects the entities present at the time when the request is made, and for each entity the SALSA agent can request its associated services. No matter what method (proactive or reactive) is used, the mechanism for service discovery works in the same way (figure 2). Each SALSA agent has a discovery policy, and when a service discovery request is made, the SALSA agent issuing the request can either provide its Jabber-ID or make an anonymous request (providing “anonymous” as its identifier).

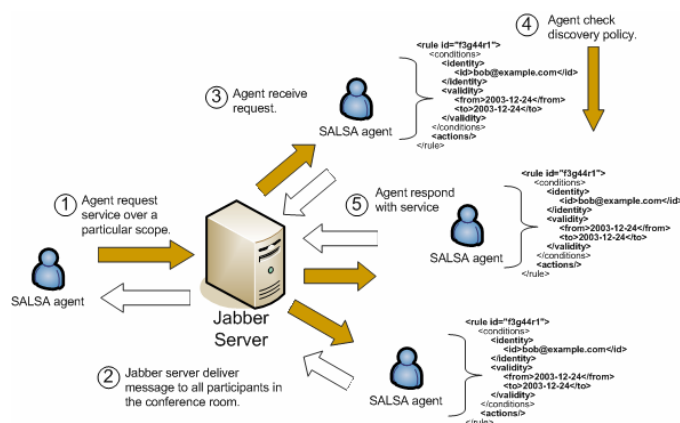


Figure 2. Service discovery mechanism within a certain scope, represented by a Jabber conference room.

5 Functional tests

The proposed architecture is supported by two key elements: the physical space server and the location server. The physical space server, the location server, the mechanisms for service discovery, as well as the policies for this discovery, have all been implemented by extending the SALSA framework. However, we have left out the implementation of the location estimation subsystem to third parties, in order to provide more flexibility; the only restriction here is to obtain standard (x,y) geographic coordinates.

For the purpose of testing the functionality of our architecture, we implemented a prototype. The SALSA agents of this prototype were given (x,y) coordinates from a simulator, which simulated the random movement of entities within a predefined physical space. We then conducted functional tests in order to determine the differences between the functional requirements and the actual system. These functional requirements are derived from the design considerations presented in section 3. For the evaluation of these functional requirements we used the following configuration:

5.1 Location system

The location system implemented is able to run a number of independent SALSA agents. These agents were assigned a random lifetime between 10 and 60 minutes. Likewise, random services are defined between five predefined ones: doctor, patient, nurse, assistant, and artifact. The location system used a uniform distribution to randomly assign (x,y) coordinates each second, taking care of making continuous movements (entities can not jump from one point to another distant one). For the tests, 40 simultaneous and independent SALSA agents were run (figure 3).

We observed that each SALSA agent determined its space from the physical space server, and was able to deduce its corresponding scope. Also, each agent subscribed to the location server. It should be noted that the location server provided a graphical interface to visualize the movement of entities in a simulated 2-D physical space during the time of the simulation, showing also the scopes where the entities are and their interactions. Service discovery was performed proactively for each agent in the simulation, entering and leaving different scopes. Service discovery in the scenario shown in figure 2 was performed as expected, getting all services in the scope during intervals of 0 to 2 seconds.

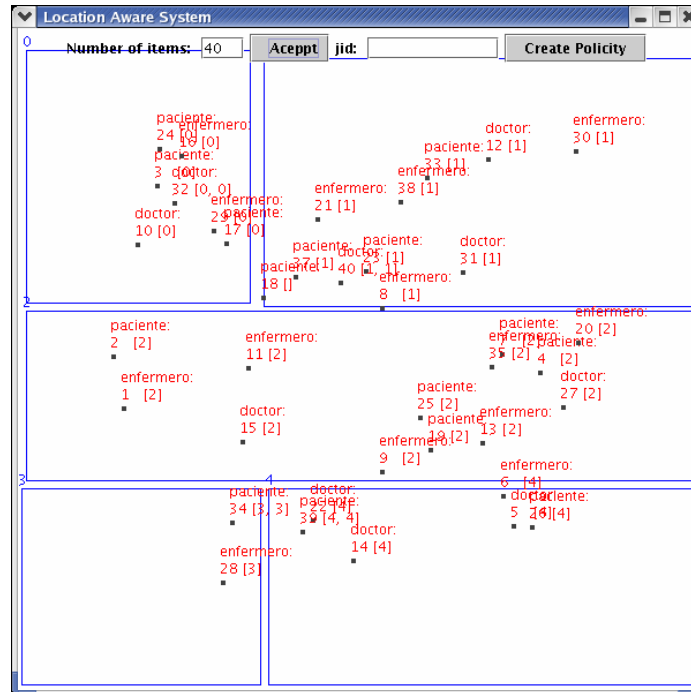


Figure 3. Forty SALSA agents moving simultaneously within a given physical space. Each point represents a mobile SALSA agent identified by a number from 1 to 40. Numbers between braces represent the current scope of the agent. The physical space is divided in fives scopes named from 0 to 4. They represent rooms, divisions, etc.

5.2 Physical space server

Five scopes (figure 2) were defined arbitrarily as regular geometric shapes through the physical space server. For the determination of the scope to which a SALSA agent belongs, and taking into account (total or partial) intersections between scopes, the prototype implemented R-trees [2]. The evaluation platform was developed using Java2 SE, a Jabber v2.2 Jive Jive Messenger⁶ on a Pentium II computer running Linux RedHat v2.4.20-8, and of course using the SALSA framework.

⁶ <http://www.jivesoftware.com>

6. Conclusions

We have presented the design considerations that should be taken into account for the design of ubiquitous distributed applications that allow the discovery of service based on their physical proximity and considering important privacy concerns. These considerations were modulated by the requirements of ubiquitous computing applications to be used in public health facilities such as hospitals. With these design considerations, an architecture was proposed and later evaluated through the development of a functional prototype. This prototype uses an extension of the SALSA framework, which defines autonomous agents that represent entities in a physical space.

The use of privacy policies, as proposed by the IETF Geopriv working group, were applied in order to restrict the way in which services are discovered and the potentially sensitive information is disclosed. Even if the proposed architecture supposes a network-centered model for the location of persons and artifacts, it can easily be extended to adopt a device-centered model.

Acknowledgements

Financial support for this project was provided by the Mexican Council for Science and Technology (CONACyT).

References

- [1] Bellotti, V., and S. Bly. Walking Away from the Desktop Computer: Distributed Collaboration and Mobility in a Product Design Team. In Proceedings of CSCW, ACM Press. 209-218 p. (1996).
- [2] Guttman, A. R-trees: A dynamic index structure for spatial searching. ACM SIGMOD Conference on Management of Data, 47-57 p. (1984).
- [3] Hodes, T., Katz, R., Servan-Schreiber, E. and Rowe, L. Composable ad-hoc Mobile Services for Universal Interaction. In Third ACM/IEEE International Conference on Mobile Computing. 1-12 p. (1997).
- [4] Johansen, T. Jini Architectural Overview. White Paper Sun Microsystems. (1999).
- [5] Microsoft Corporation, Universal Plug and Play Device Architecture Reference Specification, Version 1.0. Technical report. Microsoft Corporation.
- [6] Muñoz, M., Rodríguez, M., Favela, J., Gonzalez, V.M., and Martínez-García, A.I. Context-aware mobile communication in hospitals. IEEE Computer. 36(8):60-67 p. (2003).
- [7] Barkhuus, L., and Anind D. Location-Based Services for Mobile Telephony: a study of users' privacy concerns. INTERACT 2003, 9th IFIP TC13 International Conference on Human-Computer Interaction. (2003).

- [8] Rodríguez, M., and Favela, J. Autonomous Agents to Support Interoperability and Physical Integration in Pervasive Environments. Atlantic Web Intelligence Conference, AWIC 2003, Springer-Verlag. 278-287 p. (2003).
- [9] Santana, P., Castro, L.A., Preciado, A., Gonzalez, V.M., Rodríguez, M. D. and Favela, J. Preliminary Evaluation of UbiComp in Real Working Scenarios. 2nd Workshop on Multi-User and Ubiquitous User Interfaces (MU3I). (2005).
- [10] Chen, X., Chen, Y. and Rao, F. An efficient spatial publish/subscribe system for intelligent location-based services, Proceedings of the 2nd international workshop on Distributed event-based systems. (2003).
- [11] Coulouris, G., Naguib, H. and Sanmugalingam, K. FLAME: An Open Application Framework for Location-Aware Systems, UbiComp Adjunct Proceedings. (2002).
- [12] José, R., Moreira, A., Rodrigues, H., and Davies, N. The AROUND architecture for dynamic location-based services. Mobile Networks and Applications. 4(8): 377-387 p. (2003).
- [13] Bahl P. and V.N. Padmanabhan, RADAR: An In-Building RF-Based User Location and Tracking System, IEEE INFOCOM, Vol. 2, Tel-Aviv, Israel (March 2000), pages 775-784. (2000).
- [14] Nissanka B. Priyantha, Anit Chakraborty, and Hari Balakrishnan. The cricket location-support system. In Proceedings of MOBICOM, pages32-43, Boston, MA. (2000).
- [15] R. Want et al., The Active Badge Location System, ACM Trans. Information Systems, pp. 91-102. (1992).
- [16] Zhu, F. et al., Splendor: A Secure, Private, and Location-Aware Service Discovery Protocol Supporting Mobile Services. In Proceedings of the First IEEE international Conference on Pervasive Computing and Communications, p 235. (2003).