# AGAPE: a Location-aware Group Membership Middleware for Pervasive Computing Environments

Dario Bottazzi, Antonio Corradi, Rebecca Montanari
Dipartimento di Elettronica, Informatica e Sistemistica
University of Bologna
Viale Risorgimento 2, 40136 Bologna, Italy
{dbottazzi, acorradi, rmontanari}@deis.unibo.it

## Abstract

*The widespread diffusion of mobile computing along with the integration of telecommunication systems and the Internet enables a scenario where the promise of ubiquitous computing is starting to be realised. This scenario calls for novel services the can be deployed on-demand close to the user and customised to not only client needs, but also to the client current location. In particular, group membership management services should exploit the visibility of client location to organize effective solutions for group communication and interoperation and to promote the design and development of advanced collaborative applications, such as traffic management and e-care ones. The paper describes a middleware for group membership management (AGAPE) that exploits both the visibility of users position and the heterogeneous characteristics of the access terminals to facilitate interoperation in a pervasive computing scenario.*

## 1. Introduction

The increasing proliferation of low-cost devices equipped with computing capabilities, the widespread availability of network connectivity in the environments where users live and work, and the growing market of portable devices with wireless connectivity enable a pervasive computing scenario where the vision of being connected anytime anywhere is becoming a reality [1]. Users expect to access their needed data and services from ubiquitous points of attachments and from heterogeneous devices, ranging from fixed to portable ones, and even when changing physical locations.

This scenario promotes the development of innovative classes of applications that, on the one hand, can enable users to exploit traditional Internet services, such as Web and e-mail services, independently from their location and used access terminal, and, on the other hand, can also take into account relative positions of users, service components and available resources to tailor service results. For instance, a city guide assistant should dynamically retrieve maps, audio descriptions, and detailed textual information about the buildings and restaurants close to the current physical position of the user.

New advanced collaborative applications, such as e-care, virtual cafes, troop car and traffic management, can also stem from increased possibility of mobile device interaction with each other by exploiting both wired and wireless connectivity and from the availability of location systems that detect mobile user/device position.

However, the provisioning of pervasive collaborative applications raises new challenges due to the characteristics of the new computing scenario. The topology of the network is no longer fixed and predefined and a-priori assumptions on the status and availability of resources are not possible. Possibly mobile nodes appear and disappear in an unpredictable manner; disconnection and network partitioning are common events; the scarce bandwidth provided by wireless network technologies makes the congestion to be a normal condition; network links are implemented employing different wireless technologies such as Bluetooth [2] or IEEE 802.11 [3] and used access terminals exhibit a high degree of heterogeneity in terms of screen size, battery constraints, computing power and available memory.

Traditional group membership systems can hardly support group management operations in pervasive computing environments [4], [5]. They typically assume a network environment composed of nodes connected by reliable, usually fixed, wired networks in which link failures and network partitions are rare. In addition, group membership management operations typically depend only on the state of the interacting processes and on their connectivity status and do not take into account the current location of collaborating users/devices.

Several efforts are underway to construct the kind of

This investigation is supported by the University of Bologna (Funds for Selected Research Topics: "An integrated Infrastructure to Support Secure Services"), by CNR (Funds for Young Researchers: "An Integrated Security Infrastructure for Multimedia Services") and by MIUR within the framework of the FIRB Project "WEB-MINDS".

group management infrastructure required to support collaborative applications in environments characterized by wireless connectivity and user/device mobility [6], [7], [8].

The paper claims that novel middleware solutions for pervasive group membership management should consider the visibility of the physical position of users/devices connection to the network in addition to the traditional aspects related to the status of the network operating conditions, in order to enable effective group collaboration and to maintain and organize group members views.

The paper presents the architecture and implementation of the AGAPE (**A**llocation and **G**roup **A**ware **P**ervasive **E**nvironment) framework which provides a set of support facilities for collaborative applications in pervasive scenarios. AGAPE adopts the metaphor of group, i.e., a logical set of interacting entities associated together accordingly to application-specific requirements, as the basic mean for enabling service provision and resources sharing. The main distinguishing characteristic of AGAPE is that it is a location-aware middleware for group membership management that bases group management decisions on the location of users/devices. The underlying assumption is that users located in the same network locality are likely to require reciprocal collaboration more frequently than with members located in different physical network areas. Following this consideration, AGAPE is designed to relax the demand for a global view of all users in all network localities and to exploit location-dependent views rapidly and directly available even in the case of network partitions.

The remainder of the paper is organized as follows. Section 2 identifies current ongoing research proposals and solutions for group membership management in mobile computing and ad-hoc network environments. The AGAPE framework is presented in Section 3, while Section 4 describes the AGAPE implementation in terms of middleware components. Section 5 presents an AGAPE-based collaborative e-care application for supporting emergency rescue. Concluding remarks follow.

## 2. Emerging Solution for Group Management in Pervasive Computing

Many research efforts have been spent in the past around the development of group membership management systems. The group membership concept has been identified and clarified in [4], [9] and first introduced for synchronous systems. A group management system is generally defined as the set of services in charge of maintaining a list of the currently active and connected processes in the group, of reporting variations to the group members when the list changes and of delivering messages to all members.

Several solutions exist for traditional distributed applications that are generally classified into two main types: "Primary Partition" and "Partitionable" group membership services. Systems that fall in the former class such as [4], [10], [11] maintain a single agreed view of the actual composition of the group. Those systems assume to operate in network environments that never partition. Systems that fall in the latter class such as [5], [12], [13] allow multiple and concurrent views of the same group to coexist and evolve autonomously in different network partition. Views are reconciled whenever the network merges.

Traditional systems assume high bandwidth availability and reliable multicast and total ordered communication channels, thus providing unsuited solutions for the pervasive scenario. In addition, the main effort of traditional group membership systems is directed toward the maintenance of a global view of all interacting users, i.e., global lists of currently active and connected group members. However, it could be inefficient and unviable to ensure a global view of all users in pervasive computing environments with frequent network partitions and device disconnection. Furthermore, due to the need to guarantee global agreement on group membership, the proposed group management protocols tend to be characterized by relevant overhead and by a high degree of needed coordination among group members and require devices with relevant computational resources.

The new scenario and, in particular, the need to enable interoperation among wireless connected and resource-limited devices with scarce bandwidth availability, require the investigation of innovative solutions. Solutions should take into account issues related to user mobility and group service provisioning should depend upon the position where the user is actually located. To our knowledge, only few solutions are starting to emerge that address group management in pervasive environments.

Mobile System is a multi-level architecture for group communication in mobile systems [14]. The location of mobile nodes plays an important role and constitutes the criteria for grouping members. A Proximity Layer is used to determine the set of mobile nodes in the proximity of a given node in the network. This information is used by a three-round group membership protocol derived from [9] in order to enable processes to join/leave the group. The construction of groups takes into account only the physical proximity of mobile nodes, thus, ignoring application specific needs.

Another solution for group management in mobile ad-hoc networks is described in [6]. In particular, an algorithm is presented to provide application developers with the ability to maintain a consistent membership view in mobile ad-hoc networks. The algorithm assumes the possibility to monitor the migration and location of group members, exploits location information to decide when a host within communication range has to be admitted to or eliminated from a group and is able to accommodate both the merging of groups and the partitions of groups.

In [7] an initial proposal of a group membership middleware for mobile ad-hoc networks is introduced. The membership system follows a peer to peer model and is organised depending on device proximity. The main characteristic is related to the modelling of group members based on the notion of "Fuzzy Membership". Membership is not modelled as a binary status, i.e., a given entity may only be either a member or not a member, but it is associated to a fuzziness level that is determined at run-time and represents the level of reliability of the member. For instance, if a member is frequently unreachable, it may be associated to a low level of fuzziness.

Another group communication service targeted at pervasive computing and ad-hoc networks is described in [8]. The main contribution is that eligibility for membership depends of not only on the location where the device is actually positioned, but also on the level of its interest in the group within wireless network coverage.

## 3. The AGAPE Framework

AGAPE is a location-aware group membership framework that support the rapid design, development and deployment of pervasive collaboration applications. Possibly mobile users can exploit AGAPE to create, to join and to leave from a specific group and to receive group membership views.

A prominent feature of AGAPE is that it provides in each network locality group membership views that maintain the list only of neighbours users. As a consequence, users have direct visibility of only the locally available set of group members. When group members enter, depart or disconnect from a network locality, AGAPE reports the view changes to all co-located group members. When a group member moves from one locality to another, AGAPE is in charge of providing the mobile group member with the group view of the new locality. A global view of group members can also be provided, when needed, via the merging of all local views obtained through the coordination of AGAPE middleware components.

In addition, the AGAPE framework has been designed to address the heterogeneity of used client devices and to take into account their limited computing capacity. AGAPE overcomes device heterogeneity by providing views in different formats suitable to fit differentiated characteristics of access terminals. AGAPE also enables resource-limited portable device accessibility to group management facilities via the deployment of middleware components that act over the fixed network on behalf of users/devices in order to carry on group management-related operations.

### 3.1. The AGAPE Model

The network support for the AGAPE framework is constituted by wireless networks modelled accordingly to cell and virtual cell models [14]. In both cells and virtual cells a base station (BS) provides network coverage to the devices that are located within the scope of the associated cell (see Figure 1). The radius of the cell and the maximum number of devices that can interoperate within a determined cell depend upon the specific wireless network technology adopted for network implementation [2], [3]. Any network cell can be also augmented with a wired network that provides connectivity for fixed devices. The resulting networking environment provides connectivity for both the fixed and the mobile devices located within the coverage area provided by the base station and constitutes what we refer in the paper as a network locality.

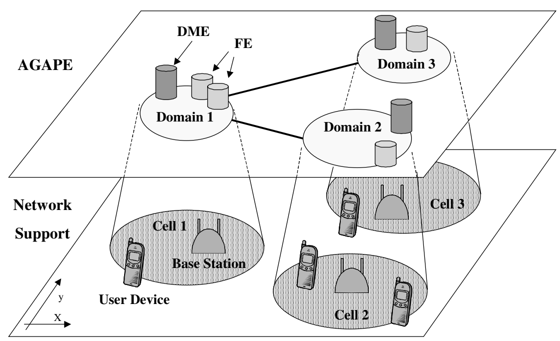


**Figure 1.** The AGAPE model in a wireless network scenario.

The AGAPE model is centered around the domain abstraction that represents a set of entities grouped together depending on relationships raised by common properties. The domain abstraction hides to application developers the details of the network support and provides them with an easy way to organize and to map logical groups over the physical network infrastructure. Each AGAPE domain is mapped over one single network locality. Let us note that whereas there is a one-to-one relationship between a domain and a locality, a single locality may contain different domains. In addition, AGAPE domains are organized hierarchically implementing an n-nary tree in order to achieve scalability. Each node of the tree is a domain. Application specific needs determine the tree structure, the hierarchical relationships and the choice of the mapping of domains over network localities. The concept of group at the application level is realized by composing one or more domains.

As it stems from Figure 1, it is possible to identify two classes of entities in the AGAPE model: the Domain

Management Entity (DME) and the Final Entity (FE).

In particular, the DME is responsible for performing group management operations, e.g., allowing users to join/leave a group domain and maintaining an updated list of the group members currently active within a specific domain. The AGAPE model assumes that there is one single DME instance for each domain. It is worth noting that a DME performs intra-domain management operations autonomously without requiring any coordination with external DMEs.

Coordination among DMEs is only required for enabling interoperation among users of the same logical group that are located in other domains in different localities. The ability of DMEs to autonomously manage group domains provides robustness to the model and allows to maintain the availability of group management support even in the case of network partitions. In fact, even if the network infrastructure connecting the various group domains partitions, causing the group itself to partition, the resulting semi-groups can evolve autonomously and can be still managed. In particular, both the autonomy of domain management and the loose coupling between DMEs of different domains ensure that network partitioning does not influence the interoperation of users within a locality.

FEs represent users/devices exploiting the group management support of AGAPE in their current network locality. When entering a locality, one FE has visibility of which AGAPE group domains are locally available and can decide to which group domain, i.e., to which logical group, to connect. If the FE is not already a member of a group, the FE has to join the group. The joining phase consists in first contacting the DME of the group domain of FE interest. The FE is, then, authenticated and authorized and if admitted into the group, the FE becomes a group member meaning that it may provide services to other group members or share specific resources or simply behave as client. The status of group member remains valid until one FE explicitly requests to leave the group. On the contrary, if the FE is already a recognized group member when it arrives at a new locality, it can decide either to immediately connect to the group domain of its interest or to defer the connection to the group domain until expressly required. In the former case, the FE is considered an active group member, in the latter case a non active group member. One FE can become a non active group member also in the case of temporary disconnection from the network. The distinction between active and non active group members is the criteria adopted in AGAPE to guide interoperation among FEs. FE interoperation is enabled only for active group members.

## 3.2. The Architecture

The AGAPE middleware is designed according to a layered architecture, depicted in Figure 2, built on top of the Java Virtual Machine (JVM). The AGAPE basic layer provides various services. The ***Location*** service allows to determine the entrance/exit of group members into a locality and to propagate location information up to the application level. The ***Discovery*** service enables FEs to discover available group domains in a locality before starting to interoperate within a specific group. The ***Event*** service permits to detect and notify to interested entities modifications in the system state, such as changes related to the location of group members or failures of neighbouring servers and clients

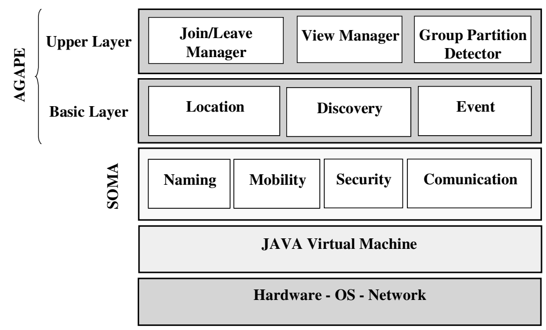


**Figure 2**. The AGAPE architecture.

The AGAPE upper layer includes the basic mechanisms and protocols to enable the interoperation of group members both when residing in the same group domains and when migrating across domains of the same group in different localities and to support the dynamic addition/removal of group domains. In particular, the ***Join/Leave Manager*** service allows AGAPE entities, both FEs and DMEs, to join or leave from a group. With regard to FEs, the service enables FEs to join/leave a specific group and to possibly register as active group members. When one FE is enabled to join a group depending on the defined group management policies, it receives an unique member group identifier and, on its turn, makes available to the group domain application-dependent information that may be of interest for the other group members, such as user/device profiles describing user preferences and device characteristics or the list of resources that the new FE is willing to share with the group. Also DMEs can exploit the Join/Leave Manager service in order to activate/de-activate at run-time a new group domain. In particular, the join protocol for domains requires DMEs to specify the father group domain to which the new domain should connect and returns a unique group domain identifier for the new domain.

The ***View Manager*** is the service in charge of creating, disseminating valid group views to AGAPE entities and updating them when changes in both group membership and group domain structure occur. AGAPE distinguishes between intra-domain and inter-domain views. Intra-domain views provide the perception of the composition of the group that is limited to the scope of a given domain. Intra-domain views are maintained in a centralized fashion by the View Manager service, locally to the domain itself. In particular, intra-domain views contain the list of currently active group members within a specific domain and for each member the associated application-dependent data that group members have provided during the join phase, e.g., user/device profiles. Intra-domain views updates occur when active group members join/leave the group and at arrival/departure of group members to/from a domain. Inter-domain views provide the perception of the composition of the group in terms of the set of domains composing the group. Inter-domain views are managed and updated through the cooperation of all the View Manager services of the domains constituting a given group. The View Manager service modifies inter-domain views when domains join/leave the group hierarchical structure and after group partitions deriving from network partitions.

The ***Group Partition Detector*** services provides the mechanisms and protocols for detecting and handling group partitions. The Group Partition Detector in each domain exploits beacons to sense group partitions. On the one hand, the Group Partition Detector sends an "I am alive" beacon to the father domain and to the possibly children domains at regular interval time. On the other hand, it waits for "I am alive" beacons from its father and children domains. If these beacons do not arrive after a determined time, the Group Partition Detector assumes that a network partitioning may have occurred. In this case, the Group Partition Detector coordinates with the View Manager service in order to update inter-domain views accordingly.

The AGAPE middleware layer is implemented on top of SOMA, a Java-based Mobile Agent (MA) platform with a rich set of facilities for the design and development of MA-based applications. Other details about the SOMA platform are available in [15]. In particular, AGAPE exploits the following subset of SOMA services: the ***Naming*** service for associating FEs and DMEs, i.e., group domains, with globally unique identifiers and to organize these identifiers in name systems to make possible the tracing of entities even if they move, the ***Communication*** service for enabling the communication between FEs and DMEs according to an asynchronous model, the ***Mobility*** service for supporting group member mobility, and the ***Security*** service for supporting authentication and authorisation controls when FEs/DMEs are willing to join a group.

## 4. Middleware Components

The AGAPE middleware in each domain is implemented in terms of three main interacting components, as shown in Figure 3: Cluster Head, Proxy and Group Access Module.

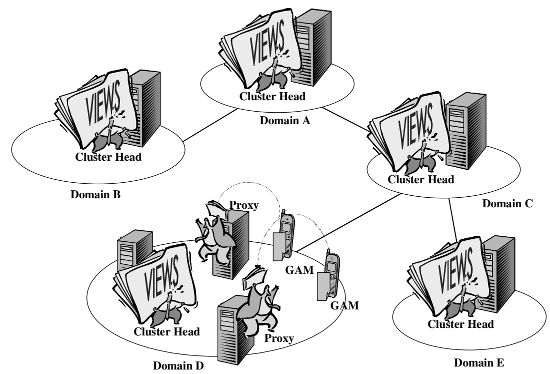


**Figure 3.** AGAPE Middleware Components.

The Cluster Head (CH) component represents the Domain Management Entity within the scope of a given domain and implements the server-side join/leave protocols along with the functionality of the View Manager and of the Group Partition Detector. In particular, CH is implemented as a SOMA agent in charge of admitting FEs into a domain and DMES, i.e., new group domains, within a specific group structure, of performing view installation and update and of cooperating with other CHs to reorganize the group structure even in the case of network partitioning.

In order to enable the collaboration between users operating from resource-limited devices we have implemented AGAPE so as to free the user device from the computational load needed for supporting client-side group management operations by moving part of the computation to the fixed network. Few solutions are starting to explore active middleware components that can be deployed at run time to act over the fixed network on behalf of users/devices. Recent research activities are recognizing the relevance of implementing middleware components as mobile entities that follow the client movements to offer the needed support only where and when needed [16]. Following these considerations, we provide any user device with a SOMA-based proxy that operates on the behalf of the user. Within the scope of a given group, a one-to-one relation exists between the user and the associated proxy. In the case a user is member of several groups, she is associated with several proxies, one for each group she belongs to. A proxy is created when its responsible user joins a group and destroyed when the user voluntarily communicates the will to leave the group.

The use of proxies has the main benefit of allowing a final user to be member of more than one logical group and at each time in each network locality to be also an active member of more than one group.

Proxies are in charge of managing, transparently to the users, their reconnection to group domains when users change their physical locality, of handling the communication between group members, of retrieving on the behalf of their responsible users the user/device profiles needed during the joining phase. Proxies can also filter or transcode the intra-domain view data received from the CH in the suited format to fit the display and memory of the current user access device and can carry on group communication and management operations even in case of temporary device disconnection. A key feature of AGAPE proxies is that they can also follow the possible movements of portable devices by maintaining the group session state. As soon as the user via its portable device reconnects to the network, her associated proxy moves to the domain where the user is actually located.

Group Access Modules (GAM) are the only AGAPE middleware components that run on user (possibly) portable devices. GAMs offer minimal functionalities and are tailored to suit also resource-limited devices, being implemented using J2ME-CLDC-MIDP1.0. When a user enters a locality, the GAM component running on its device coordinates with the discovery service in order to retrieve the locally available groups and associated domains, i.e., to discover the CHs in the locality. Then, GAM enables the user to select a specific group and to join it. In addition, if a user joins several groups, GAM permits her to dynamically choose the group in which to operate at each instant of time. The default behaviour in a scenario characterized by the possibility to work with multiple groups is to transparently maintain the connection to the same group upon user migration toward a new locality. At run-time, when a user is already an active group member, GAM acts as simple interface between the user and her associated proxy in order to exploit the group management services. Finally, when a user is willing to leave the group, the GAM component forwards the request to the CH.

## 5. AGAPE at Work in an E-Care Scenario

To illustrate how AGAPE supports and facilitates the development of location-aware collaborative services, let us consider an e-care scenario. We have developed an emergency healthcare application prototype on top of the AGAPE framework that allows subscribed cardiopath patients moving in outdoor environments to promptly communicate a status of emergency and to be instantly assisted. In particular, the application exploits the visibility of the position of patients and physicians to forward emergency calls to the nearest physician in addition to the nearest hospital. Contacted physicians can then provide initial assistance and can start to insert initial patient-related information, such as currently observable symptoms.

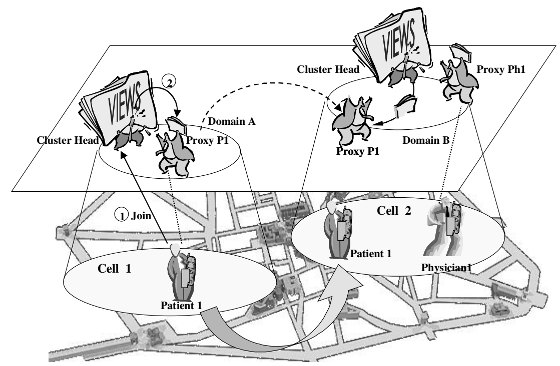


**Figure 4.** AGAPE at work in an emergency healthcare scenario.

Figure 4 shows our test-bed setting for the application prototype in one wireless metropolitan network. The application aggregates together in one single logical group cardiopath patients and physicians and decomposes the group into disjoint domains mapped over cells with 802.11 access points. Each domain contains a AGAPE Cluster Head and proxies components operating on the behalf of both the patients and physicians currently in the cell. Cardiopath patients and physicians access and interact with the AGAPE infrastructure via Group Access Modules running on wireless connected Compaq iPAQ PDAs with IEEE 802.11-cards.

The GAMs installed on PDAs provide similar group management operations operations to both patients and physicians. The only differences are in the provided application-dependent services: patient GAMs allow patients to forward emergency call, whereas physician GAMs are augmented with the possibility to retrieve, elaborate, insert patient information and to collaborate with colleagues.

Let us suppose that one patient is moving around the urban area. When the patient enters a locality and her wireless device connects to the network via the local access point, the location service collects information related to the new connected client device by exploiting the Simple Network Management Protocol (SNMP). In particular, the location service receives from an SNMP agent installed on the local IEEE 802.11 access point roaming-specific SNMP traps that are triggered by the client terminal entrance in the locality [17]. The first time the patient accesses the emergency healthcare service, the GAM on her PDA allows her to authenticate and to insert profiling data that contain personal medical information, such as potential allergies, thus producing the discovery of the CH in the current cell and the activation of the group join protocol. The CH coordinates with the SOMA underlying Security service to authenticate the patient and if authentication succeeds, the patient is assigned with a group unique identifier and a Proxy component. The CH, then, updates the intra-domain view and distributes the view to all the Proxies of the group members currently active within the local domain. Views contain the patient identifier and related medical information. Similar actions take place at the joining phase of physicians.

Once registered as active group members, both patients and physicians may freely move and connect to different domains at different times. When one user disconnects from a domain, this event is detected and notified to the CH that updates the intra-domain view accordingly. When the same user reconnects to a new domain as an active group member, the event is communicated to her Proxy that moves to the new domain by exploiting the Mobility SOMA service and contacts the local CH that updates the intra-domain view to reflect the arrival of the new group member and distributes the view to all other members.

In a heart attack emergency, the patient requests medical assistance via her GAM. The client request is notified to the associated Proxy that checks if there is one physician within the scope of the associated domain. This operation is handled simply by consulting the intra-domain view. If one physician is present in the same domain, an help message reaches the associated physician Proxy. The message brings data related to user ID and actual position to allow the physician to approach the patient.

Let us note that for the sake of description simplicity of the AGAPE components at work we have here simplified the presentation of the healthcare application prototype. It is worth stating that AGAPE also enables the interoperation of neighbour physicians. In the depicted example, for instance, if the patient is also pregnant, AGAPE allows the first contacted physician to request help to a gynaecologist colleague in the nearest locality and to exchange messages while waiting for her arrival. The AGAPE middleware will provide the infrastructure support to exchange messages between the physicians along with providing all the patient-related information to the contacted gynaecologist.

## 6. Conclusions and Future Work

The increasing mobility of users, the resource-limited and heterogeneous characteristics of user access terminals, the unreliability and scarce bandwidth capacity of wireless links raise challenging problems for the design, development and deployment of group membership management systems in pervasive computing environments and require to re-think and re-design traditional group membership management solutions. AGAPE intends to give a contribution to the research area of middlewares to support pervasive group-aware applications.

AGAPE supports a concept of group that allows application developers to decide the eligibility of group membership depending on application-specific needs. But, as a key feature, AGAPE considers the location of group members as a first-class concept to take into account for group management decisions. In particular, AGAPE exploits the visibility of the position of group members to structure the group views. This offers an increased level of robustness, efficiency and scalability to group member interoperation and promotes the development of novel collaborative applications where it makes sense to organize service provisioning depending on user/device co-locality. In addition, AGAPE provides group management support also to ultra-constrained devices by adopting a proxy-based approach that frees mobile clients from group management computations.

For the sake of simplicity, we have only presented a simple e-care application prototype built on top of AGAPE, but we are experimenting the AGAPE middleware in a wide variety of scenarios. First experiences in the use of AGAPE have shown that our middleware can simplify the design and implementation of location-aware collaborative services. These results are stimulating further research along different guidelines to improve the current prototype and to develop more complex services on top of it. We are currently working on the integration of AGAPE with a policy-based management framework to improve the dynamic reconfigurability of the structure and the topology of the group and to simplify the long-term manageability of the system.

## References


[1] M.Weiser, "The Computer for the Twenty-First Century," Scientific American, Vol. 265, No.3, September, 1991.

[2] Bluetooth SIG, http://www.bluetooth.org/.

[3] IEEE 802.11 Wireless LAN Working Group, http://grouper.ieee.org/groups/802/11/index.html.

[4] B. Glade, et al., "Light-weight Process Groups in the ISIS System", Distributed System Engineering, Vol. 1, No. 1, IOP Press,1993.

[5] Ö.Babaoğlu, et al. "On Group Communication in Large Scale Distributed Systems", in Proceedings of the ACM SIGOPS European Workshop, ACM Press, Dagstuhl, Germany, September, 1994.

[6] Q. Huang, et al., "Relying on Safe Distance to Ensure Consistent Group Membership in Ad Hoc Networks", Technical Report WUCS-01-31, Washington University, Department of Computer Science, St. Louis, Mis-

souri.
[7] R. Friedman, "Fuzzy Group Membership", in Proceedings of the International Workshop on Future Directions in Distributed Computing (FuDiCo 2002), Springer-Verlag LNCS 2584, June 2002.
[8] R. Meier, et al., "Towards Proximity Group Communication", in Proceedings of the Workshop on Middleware for Mobile Computing, ACM Press, Heidelberg, Germany, November 2001.
[9] F. Cristian, "Reaching Agreement on Processor Group Membership in Synchronous Distributed Systems", Distributed Computing, Vol. 4, No. 4, Springer-Verlag, April 1991.
[10] R. Friedman, et al., "Strong and Weak Virtual Synchrony in Horus", in Proceedings of the IEEE 15th Symposium on Reliable Distributed Systems, IEEE Press, Niagara-on-the-Lake, Ontario, Canada, October, 1996.
[11] M.F. Kaashoek, et al., "Group Communication in the Amoeba Distributed Operating System", in Proceedings of the 11th International Conference on Distributed Computer Systems, Arlington, TX, May 1991.
[12] L. Rodrigues, et al., "Partitionable Light-Weight Groups", in 20th IEEE International Conference on Distributed Computing Systems (ICDCS'20), IEEE Press, Taipei, Taiwan, April 2000.
[13] A. Fekete, et al., "Specifying and Using a Partitionable Group Communication Service", in 16th Annual ACM Symposium on Principles of Distributed Computing, ACM Press, Santa Barbara, CA, August, 1997.
[14] R. Prakash, et al., "Architecture for group communication in mobile systems". In 17th IEEE Symposium on Reliable Distributed Systems, IEEE Press, West Lafayette, Indiana , October, 1998.
[15] P. Bellavista, et al., "Protection and Interoperability for Mobile Agents: A Secure and Open Programming Environment", IEICE Transactions on Communications, Vol. E83-B, No. 5, May, 2000.
[16] P. Bellavista, et al., "Mobile Agent Middleware for Mobile Computing", Computer, Vol. 34, No. 3, March 2001.
[17] P. Bellavista, et al., "Application Domain Accounting for Roaming Services", in Proceedings of the 8th IEEE Workshop on Future Trends of Distributed Computing Systems (FTDCS'03), IEEE Press, San Juan, Puerto Rico, May, 2003.