Towards Proximity Group Communication

René Meier, Marc-Olivier Killijian, Raymond Cunningham, and Vinny Cahill

{rene.meier, marco.killijian, raymond.cunningham, vinny.cahill}@cs.tcd.ie

Distributed Systems Group
Department of Computer Science
Trinity College Dublin

Abstract: Group communication will undoubtedly be a useful paradigm for many applications of wireless networking in which reliability and timeliness are requirements. Moreover, location-awareness is clearly central to mobile applications such as traffic management and smart spaces. In this paper, we introduce our definition of proximity groups in which group membership depends on location and then discuss some requirements for a group membership management service suitable for proximity groups.

1 Introduction

The widespread deployment and use of wireless data communications is generally recognised as being the next major advance in the information technology industry. In the long term, wireless data networks will represent a key enabling technology underlying the vision of *ubiquitous computing* [1]. In this vision, interconnected computers will be embedded in a wide range of appliances ranging in size from door locks to vehicle controllers, and will co-operate to perform tasks on behalf of their human users ranging from automatically opening doors to routing vehicles to their intended destinations in co-operation with other vehicles' controllers. Mobility, and hence wireless networking, is clearly central to this vision. We believe that, as is the case for fixed networks, group communication [2], [3] will be a useful paradigm for many such applications of wireless networking in which reliability and timeliness are important requirements. A major feature of wireless communications is the fact that participants can be mobile and hence that their location can have an impact on the information in which they are interested or that they can provide. Hence, we believe that any wireless group communication system should support location awareness. Moreover, knowledge of the location of participants can be exploited in the implementation of the group communication system itself.

In this paper, we consider the problem of group communication in a wireless network. Much of the previous work in this area deals with routing protocols for group communication based on multicast or geocast [4], [5]. In this paper, we concentrate on the definition and semantics of group membership for location-aware mobile participants. While some research has already been done on groups in which membership is based on location information [6], [7], our definition of *proximity groups* takes into account both location and functional aspects. When group membership depends on location, it becomes important to understand what proportion of the area of interest is within wireless network coverage. For this purpose we defined a novel approach to coverage estimation [8]. Another important aspect of mobile computing is that partitions are very likely to occur. For this reason our

membership layer also includes a new failure and partition anticipation scheme that can take into account movement of nodes, battery life, etc. This algorithm can be set up to be either optimistic or pessimistic and tries to anticipate partitions/failures in order to maintain consistent group views.

2 Overview of Proximity Group Communication

At the heart of our approach to group membership management is the use of location for both functional and non-functional reasons.

- Firstly, in functional terms, it often makes sense to define a group in a mobile application in terms of a geographical area. We can easily imagine many cases where this would be interesting: in traffic management, for example, the area around a traffic-light could be used to define a group with cars in that vicinity becoming members of the group to receive notifications of changes to the state of the lights; in a similar way, we might want to define a group corresponding to the area around an ambulance in order to inform nearby cars to yield the right of way.
- Secondly, from the non-functional point of view, we can use location information to, for example, anticipate partitions and hence take preventative measures to ensure consistent group views when these partitions happen.

In classical group communications, a group is defined by its functional aspect, e.g. its name. Our notion of proximity group involves both location and functional aspects, i.e., to be able to apply for group membership, a node must firstly be located in the geographical area corresponding to the group and secondly be interested in the group. In the following paragraphs we discuss the various possibilities related to the location aspect of group membership.

To define a proximity group, we firstly have to define the area that it covers as a geometric shape with associated coordinates. Any kind of shape can be used, i.e., it need not necessarily be a circle or a square but can be arbitrarily complex. We can obviously imagine using 2 or 3 dimensions, but it is also possible to include time in the definition of the area, e.g., ``the area around the position at which the ambulance was located at 2.00pm". To define the coordinates of the area, we associate a reference point with the shape. We distinguish two cases: either the group is absolute, i.e. geographically fixed, or it is relative to a moving point, its so-called navel. In the absolute case, the reference point is attached to a fixed point in space. In the relative case, it is attached to the navel, i.e. an identified node.

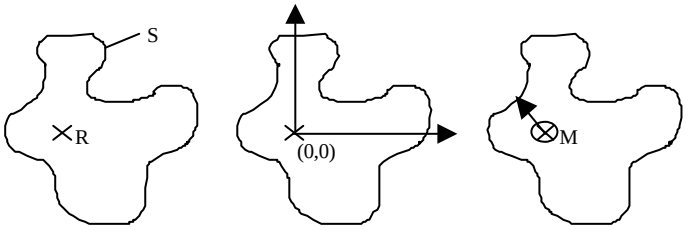

Figure 1: Area Definition

Figure 1 illustrates this notion of an area. The first shape S on the left is associated with a reference point R. This reference point is relative to the shape. The definition of the area is not complete since R has not been attached to a point (possibly moving) in space. For the second shape, R has been attached to the point (0,0), making the area absolute. The reference point of the third shape has been

attached to a node M that represents the navel of this relative proximity group. Table 1 gives some example group areas for two different application domains.

Domain	Absolute Proximity	Relative Proximity
Traffic Management	*Traffic light:* a traffic light informs nearby cars of its status. The shape is a circle; the reference point is the centre of the circle and is attached to the geographical coordinates of the traffic light.	*Ambulance:* an ambulance on call informs nearby cars to yield the right of way. The shape is a square, the reference point its centre and the navel is the ambulance itself.
Smart Spaces	*Resource access:* to use a printer, nearby people must reserve it using the printer proximity group. The system administrator defines the shape according to the available printers and offices; the reference point is at the printer and attached to the printer's coordinates.	*Centralised tour guide:* in a museum a group of tourists wear headsets and are remotely guided by an automatic guide. The area surrounds the group of tourists, while the navel is attached to one of them.

Table 1: Examples of Group Areas

In typical group communication, a group is roughly defined by a topic (or a name) and nodes can join this group if they are interested in its topic. We believe that this is also necessary for proximity groups because a node in the area of the group may or may not be interested in joining the group. We then add the functional aspect to the previous definition of a proximity group by associating a name with each group:

Definition 1. A Proximity group G is completely defined by the shape, the reference point, the navel and the name:

$$G = \{Shape, Reference\ Point, Navel, Name\}$$

3 Architecture Sketch

Our goal is to define a group membership management layer suitable for proximity groups. Since in this model, location is intrinsic to group membership, it is important to be able to provide applications with at least an estimate of the probability of there being one or more nodes, which while currently in the area of interest, are disconnected from the group, typically because of lack of network coverage. To address this issue we provide a coverage estimation tool [8] based on a novel algorithm that uses knowledge of the connectivity graph of the network. Coverage estimation can also be used to select the appropriate approach to our group communication. Our membership management protocols should also be failure aware and anticipate partitions, which are very likely to occur in the kind of networks that we are considering. This will lead us to the definition of a partition anticipation tool. Given these tools we will be in a position to define appropriate routing and geocasting protocols to be used by the group membership management layer. Figure 2 summarizes the resulting architecture.

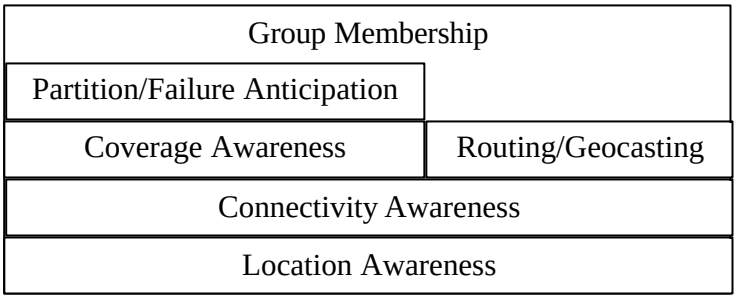

Figure 2: Summary of Requirements

4 Summary, Conclusions and Future Work

In this paper, we described a model of group membership for location-aware mobile participants that is at the heart of a new group communication toolkit for wireless networks that we are developing. In this model, eligibility for group membership depends on the location of the potential member and, in particular, each group is associated with a static or a mobile area of interest within which its members should be located. This model is aimed primarily at applications in the traffic management [9], smart space [10] and augmented reality domains. We also described some of the considerations underlying our approach to group membership management that exploits location information to achieve coverage estimation and partition anticipation. We are currently developing a suite of multicast protocols providing different ordering, reliability and timeliness guarantees based on this membership substrate.

5 References

[1] M. Weiser, "Ubiquitous Computing," *IEEE Hot Topics*, vol. 26, pp. 71-72, 1993.

[2] K. P. Birman, R. Cooper, T. A. Jospeh, K. P. Kane, F. Schmuck, and M. Wood, "ISIS-a distributed programming environment," Cornell University, Cornell 1990.

[3] F. Cristian, "Synchronous and Asynchronous Group Communication," *Communications of the ACM*, vol. 39, 1996.

[4] E. M. Royer and C. E. Perkins, "Multicast Operation of an Ad Hoc On-demand Distance Vector Routing Protocol," presented at MobiCom, Seattle, WA, 1999.

[5] Y.-B. Ko and N. H. Vaidya, "GeoTORA: A Protocol for Geocasting in Mobile Ad Hoc Networks," presented at 8th International Conference on Network Protocols (ICNP), Osaka, Japan, 2000.

[6] R. Prakash and R. Baldoni, "Architecture for Group Communication in Mobile Systems," presented at Symposium on Reliable Distributed Systems, West-Lafayette (IN), USA, 1998.

[7] G.-C. Roman, Q. Huang, and A. Hazemi, "On maintaining Group Membership Data in Ad Hoc Networks," Washington University, St Louis, Technical Report wucs-00-26, April 16, 2000 2000.

[8] M.-O. Killijian, R. Cunningham, R. Meier, L. Mazare, and V. Cahill, "Towards Group Communication for Mobile Participants," presented at Workshop on Middleware for Mobile Computing, Newport, Rhode Island, USA, 2001.

[9] R. Cunningham and V. Cahill, "System Support for Smart Cars: Requirements and Research Directions," presented at 9th ACM SIGOPS European Workshop, Kolding, Denmark, 2000.

[10] P. Nixon, S. Dobson, and G. Lacey, "Smart Environments: Some challenges for the computer community," presented at 1st International Workshop on Managing Interactions in Smart Environments, Dublin, Ireland, 1999.