System Support for Smart Cars: Requirements and Research Directions

Raymond Cunningham and Vinny Cahill
Distributed Systems Group
Department of Computer Science
Trinity College Dublin
Ireland
{rcnnnghm, vjcahill}@cs.tcd.ie

Abstract

With the introduction of Wireless Local Area Networking technology in recent years, many exciting applications are now closer to becoming a reality. One such application is the smart car or automated vehicle - a car capable of transporting its passenger(s) to their intended destination without human intervention. For this application, automated vehicles need to be able to co-ordinate with each other in a reliable and timely fashion.

A logical approach to this problem would be to use existing real-time group communication techniques. However, existing techniques have been designed with a more static network in mind. In particular, group membership is more dynamic in the automated vehicle application. In this paper we outline new requirements for group communication suitable for such mobile applications. Such group communication must take into account the proximity of members of the group to each other.

The Anois project, at Trinity College Dublin, is currently developing such a group communication system addressing these new requirements that arise due to mobility.

1. Introduction

With the introduction of Wireless Local Area Networking technology such as the IEEE's 802.11 standard [6] in recent years, many exciting applications are now becoming more and more feasible. One such application is the smart car or automated vehicle. In this application the user literally takes a back seat only specifying his or her destination. The vehicle then takes control and brings the user to their specified destination.

Typically there is more than one vehicle on the street or highway at any one time and a first attempt to solve this problem has been to introduce platoons of automated vehicles travelling in the same direction. These platoons typically consist of two to twenty vehicles that initially travelled in their own lane, segregated from existing lanes. More recent work has allowed cars to travel at 85km/h on a freeway near San Diego [7].

For the automated vehicle to become a reality, communication between these automated vehicles is needed to support their co-ordination. This co-ordination will result in decisions being made by each individual automated vehicle. These decisions will need to be integrated with control systems already available on existing vehicles such as engine management and braking units. Recent work has integrated obstacle warning systems into the existing Cruise Control features of vehicles. Two examples of these obstacle warning systems are BMW's Electronic Assistant system that uses a digitised image, from a video camera, of the road ahead to keep a car within its lane and Mercedes Benz's Distronic system that automatically activates the brakes when there is traffic in front [9]. It uses radar control to scan the road ahead for 150 meters spanning across three lanes.

The automated vehicle system introduces new and interesting problems. The most obvious problem is collision avoidance by one automated vehicle of another vehicle. Co-ordination between automated vehicles within a platoon, and between platoons, when avoiding a collision and when an automated vehicle wishes to join or leave a platoon is another problem. These problems require real-time discovery of other automated vehicles and obstacles, timely dissemination of information between vehicles within a platoon and between platoons, and management of the membership of a platoon. The automated vehicles or platoons also need to be integrated with existing traffic management systems to provide co-ordination at traffic lights and allow optimal routing of traffic through the system. Solutions to these problems also need to address the characteristics associated with wireless networking such as variable bit error rates and variable end-to-end propagation delay to name but two.

From the above discussion, distributed systems researchers would be tempted to say that platoons logically correspond to the group abstraction. Existing real time group communication techniques could then be applied to the solution of the automated vehicle and similar problems. However, existing real time group communication techniques are not sufficient in trying to solve these problems.

In the next section we discuss in greater detail why existing real time group communication techniques are insufficient and some of the requirements of the automated vehicle system with particular reference to new requirements on group communication systems. Following this in section 3, we discuss the approach that will be taken in the Anois project to support the development of applications like the automated vehicle system that have these requirements. Finally in section 4 we discuss the current status of the Anois project and future directions.

2. Requirements

To allow automated vehicles and platoons of these vehicles to become a reality in the coming years, safety is a critical factor. Initial group communication systems have assumed a static network topology. In particular real time group communication protocols have relied on the characteristics of the underlying network or networks to allow predicable and reliable communication [2, 5]. Existing real time systems typically assume a fixed maximum membership model of hosts in the system [3]. This model would not be applicable in the automated vehicle application due to the very dynamic nature of the movement of vehicles. Therefore new group membership management techniques suitable for such a mobile and dynamic application are needed.

For stable control of an automated vehicle, the vehicle must update its acceleration and velocity every 20 milliseconds [1]. Therefore for a platoon of 20 vehicles each member of the platoon has to be able to send a message at least every millisecond to warn other members of the platoon of impending obstacles. The vehicle at the front of the platoon must then alert other members of the platoon within a well-defined time bound so that each member can take evasive action. Due to the mobility and possible failure of members of the platoon, neither one of them can ever be assured of the exact membership of the platoon. This implies that membership of the platoon needs to take into account the proximity and travelling direction of a member in relation to the other members of the platoon.

As well as co-ordination within platoons, automated vehicles would need to join and leave platoons. This joining and leaving needs to be time bounded with this time bound varying with the velocity at which the platoon is travelling. This would require the joining vehicle knowing their position and the proximity of other members of the platoon.

In addition to the above requirements, there is also a requirement that platoons should avoid each other. This requires communication between platoons. To achieve this a platoon needs to know its proximity to other platoons within a certain time bound and it then needs to be able to co-ordinate with that platoon or platoons to avoid the impending collision.

The co-ordination between platoons requires some form of wireless communication. A line of sight communication mechanism cannot be considered since two platoons may be arriving at a junction in a built up area. Early investigations of WaveLAN [1] have shown it to be unsuitable for this purpose due to the limited bandwidth of 2Mb/s. A more recent version has increased the bandwidth to 11Mb/s, with further increases likely, thus improving it's possible applicability for the automated vehicle application.

In summary, an automated vehicle system requires time-bounded communication, dynamic group membership management, real-time discovery of other automated vehicles and knowledge of the proximity of other members of the platoon.

3. Anois Architecture

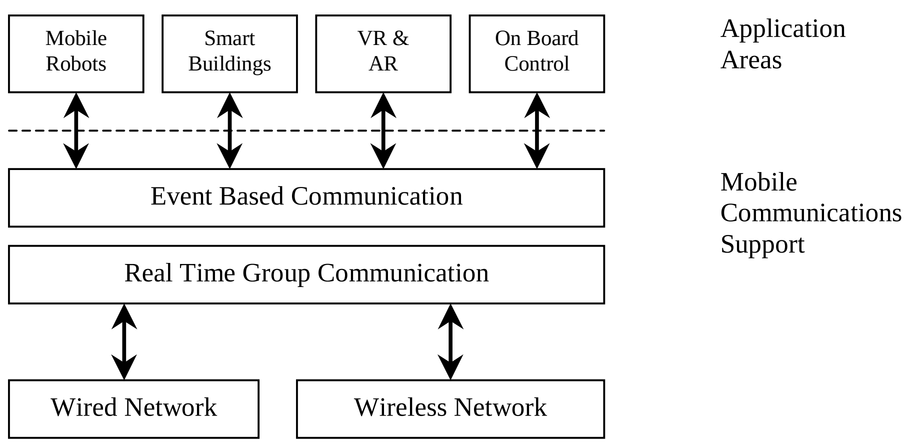

Figure 1. Anois Architecture

The Anois Architecture takes a layered approach to the development of mobile group applications. At the top layer are application areas that exhibit both mobility and group communication requirements.

These applications will be built using easy to use event-based programming paradigms to provide asynchronous one-to-many communication. Event based programming has the added advantage that it does not require the producers of events to have knowledge about the number and/or location of event consumers. As an example, in the automated vehicle system an “obstacle warning” event would need to propagate from the vehicle at the front of the platoon to other members of the platoon. However, the propagation of this event needs to be limited within a well-defined distance of the obstacle’s location.

The Event Based Communication layer relies on the Real Time Group Communication layer. This layer will build on top of both wired and wireless networking technology to provide a group membership abstraction based on proximity to members of the group to each other and direction of travel of members of the group. This notion of proximity to other group members will build on the physical position of the members of a group possibly using the Global Position System or a triangulation technique.

Continuing the “obstacle warning” event example, the producer of the event will include it’s current position as part of the message to other members of the group. The propagation of this message will take into account the position of the event producer. An “obstacle warning” event may result in the splitting of a group since some members may have passed the obstacle while others have not.

The Real Time Group Communication layer will also need to provide reliable and time-bounded communication between the members of a group in addition to providing primitives to join and leave the group.

4. Mobile Group Communication

Mobility introduces new problems into the real-time group communications arena. The way in which group membership is managed needs to change due to the dynamic nature of members of the group. Before discussing some of these problems, some of the characteristics of wireless networks, with particular reference to the IEEE 802.11 standard, are discussed.

4.1 Wireless Networks

Wireless networking technology breaks down into three distinct areas: Wireless Local Area Networking (WLAN), Wireless Metropolitan Area Networking (WMAN) and Wireless Wide Area Networking (WWAN). The last two types are impractical for the automated vehicle system due to their limited bandwidth and the latency involved, by requiring them to communicate through a fixed base station, in communicating between two mobile hosts in the same local area. There are a number of Wireless Local Area Networking technologies that are currently being used or have been proposed with the IEEE 802.11 Wireless LAN standard (WLAN) being the most widely implemented.

In the IEEE 802.11 standard, two different types of networks are possible. The first is an Infrastructure network where mobile hosts can communicate with mobile hosts outside their local area by first communicating with a fixed base station. The base stations are connected to each other by some backbone network, possibly IEEE

802.11 or some other networking technology. Base stations route packets, which are destined for other mobile hosts, to the correct base station that then forwards the message to the destination mobile host. The second type of network is an Ad Hoc network. What differentiates this type of network from Infrastructure networks is that there are no fixed base stations. Each node in the system has the potential to move and mobile hosts route packets on behalf of other mobile hosts to deliver their packets to their destination.

IEEE 802.11 uses two types of access to the wireless medium access protocol. The first, used in Infrastructure networks, requires a fixed base station to poll mobile hosts for data to send and the second employs a CSMA/CA mechanism that can be used in both Infrastructure and Ad Hoc networks. There is the possibility of the hidden terminal problem. This arises when two mobile hosts, outside each other's coverage area, wish to send a message to another host within the overlap of the two coverage areas. If both messages are sent at the same time then each message will corrupt the other at the host in the overlapping coverage area without either sending host observing that their message has been corrupted.

One way to overcome this problem is for the base station to poll the mobile hosts in its coverage area to see if they have data to send. This implies that the base station knows the number and identity of mobile hosts in its coverage area and that the polling can be carried out. This polling has been shown to perform badly [10]. More recent work has seen the development of a real-time medium access protocol on top of 802.11 [8].

4.2 Status of Anois

To date Cristian's multiple fault tolerant synchronous atomic broadcast protocol [2] has been implemented using Lucent Technologies WaveLAN [11] as the underlying network technology. This implementation is currently being evaluated.

The protocol specified by Cristian requires each host to be connected to every other host via independent communication networks and for the clocks at each host to be synchronised using a synchronisation algorithm [4]. The protocol can tolerate up to a well-known number of failures by delaying the delivery of each atomic broadcast by a known time Δ. The calculation of Δ relies on the time taken to send a message from one host to any other being bounded above by a known constant. In addition to this, the precision of the synchronisation algorithm relies on the timeout interval of the clock synchronisation algorithm.

The goal of the evaluation that is currently being carried out is to calculate values for Δ and the timeout interval of the clock synchronisation algorithm that then give a corresponding probability of correct message delivery at each host in the system and of consistent message delivery at two or more hosts. The precision of the clock synchronisation algorithm will also be evaluated.

From this evaluation we expect to gain a greater insight into the problems that existing real time group communication protocols face due to mobility and wireless networking. This insight will then allow us to develop extensions to cater for these problems.

Summary

In this paper we have reviewed some of the problems that need to be addressed if real-time group communication is to be applied to a mobile environment. We propose a new approach to group membership management based on proximity of group members to each other.

References

[1] Chen, "A Discussion of the WaveLAN Radio As Relevant To Automated Vehicle Control Systems", Department of Electrical Engineering and Computer Science, University of California, Berkeley CA, July 1995.

[2] Cristian, "Synchronous Atomic broadcast for redundant broadcast channels", The Journal of Real Time Systems 2, 1990.

[3] Cristian, "Reaching Agreement on Processor Group Membership in Synchronous Distributed Systems", Distributed Computing, Vol 4, pp 175-187, 1991

[4] Cristian, Aghili and Strong, "Clock Synchronisation in the Presence of Omission and Performance Failures and Processor Joins", 16th IEEE Int. Symposium on Fault Tolerant Computing Systems, Vienna, July, 1986.

[5] Gründsteidl and Kopetz, "A Reliable Multicast Protocol for Distributed Real-Time Systems", 8th IEEE Workshop on Real-Time Operating Systems and Software, Atlanta, GA, USA, May 1991.

[6] IEEE Computer Society LAN MAN Standards Committee, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications", IEEE Std 802.11-1997, The Institute of Electrical and Electronic Engineers, New York.

[7] NAHSC, The National Automated Highway System Consortium Technical Feasibility Demonstration. http://www.monolith-mis.com/ahs/Default2.htm

[8] Sobrinho and Krishnakumar, "Real-time traffic over the 802.11 medium access control layer, Bell Labs Technical Journal, 1996, pp 172-187.

[9] Tuckey, "The Driverless Car", http://businesscentre.bigpond.com/lifestyles/motoring/fact_sheets/driverless_car.htm

[10] Visser and Zarki, "Voice and Data Transmission over an 802.11 Wireless Network", Proceedings of PIMRC'95, Toronto, Canada, September 1995, pp 648-652.

[11] WaveLAN, http://www.wavelan.com/