

Towards Real-Time Middleware for Vehicular Ad Hoc Networks

Barbara Hughes, René Meier, Raymond Cunningham and Vinny Cahill

Distributed Systems Group,

Department of Computer Science, Trinity College, Dublin, Ireland

{barbara.hughes, rene.meier, raymond.cunningham, vinny.cahill}@cs.tcd.ie

ABSTRACT

Applications of inter-vehicle and vehicle-to-roadside communication that make use of vehicular ad hoc networks (VANETs) will often require reliable communication that provides guaranteed real-time message propagation. This paper describes an event-based middleware, called RT-STEAM. Unlike other event systems, RT-STEAM does not rely on a centralized event broker or look-up service while still supporting event channels providing hard real-time event delivery. RT-STEAM event filtering can be based on subject, content and/or proximity. To guarantee real-time communication, we exploit proximity-based event propagation to guarantee real-time constraints within the defined proximities only. The proximity within which real-time guarantees are available is adapted to maintain time bounds while allowing changes to membership and topology, typical of VANETs. This Space-Elastic Model of real-time communication is the first to directly address adaptation in the space domain to guarantee real-time constraints.

Categories and Subject Descriptors

C.2.4 [Computer Communication Networks]: Distributed Systems

General Terms

Design, Reliability

Keywords

Event-based middleware, Real-time, Proximity adaptation

1. INTRODUCTION

Ad hoc wireless networks comprise sets of mobile nodes connected by wireless links that form arbitrary wireless network topologies without the use of any centralized access point or infrastructure. Ad hoc wireless networks are inherently self-creating, self-organizing and self-administering.

While most research in ad hoc networks has assumed a random waypoint mobility model in a network of a particular shape, e.g., rectangular, the specific patterns of vehicle movement make inter-vehicle and vehicle-to-roadside networks distinctive. In particular, the potential for high speeds and the limited dimensionality afforded by a confined roadway differentiates vehicular ad hoc

networks (VANETs) from other ad hoc networks.

By enabling inter-vehicle and vehicle-to-roadside communication a broad range of applications in the areas of cooperative driver assistance and mobile information systems are being developed. One of the more sophisticated applications for inter-vehicle communication is the platooning of vehicles. For example, the lead vehicle in such an application may broadcast sensor information to coordinate movement and potentially reduce the consumption of fuel.

A possible application of vehicle-to-roadside communication is in next generation urban traffic control (UTC) systems. A traffic light controller at a junction could be informed of a vehicle's pending arrival and change the traffic light sequence to allow the approaching vehicle to pass through the junction without stopping. If the traffic light controller is informed of a number of vehicles approaching the junction, then the controller could optimize the traffic flow across the junction. This time constrained communication between vehicles and traffic light controllers should continue reliably during peak times when a large number of vehicles are approaching the junction from a number of different directions. The potential for contention increases as the number of vehicles communicating with the controller increases.

Both the vehicle platooning and UTC applications require a communication paradigm that supports dynamic changes to the topology of the underlying wireless network, e.g., to accommodate vehicles joining and leaving a platoon, as well as delivery guarantees for time-critical messages, e.g., sufficient time for the traffic light controller to change the flow of traffic across the junction.

This paper describes an event-based middleware, called RT-STEAM¹ [1], that is based on an implicit event model [2] and has been designed for mobile applications and ad hoc networks. RT-STEAM differs from other event-based middleware in that it does not rely on the presence of any separate infrastructure and supports distributed techniques for identifying and delivering events of interest based on location. RT-STEAM supports decentralized approaches for discovering peers, for routing event notifications using a distributed addressing scheme, and for event filtering based on combining multiple filters. RT-STEAM provides a programming model based on the concept of event channels. A number of event channel classes with different temporal and reliability attributes are available to integrate real-time support into the event channel model. Depending on the guarantees available from the underlying network, the proximity

Copyright is held by the author/owner(s)
VANET'04, October 1, 2004, Philadelphia, Pennsylvania, USA.
ACM 1-58113-922-5/04/0010

¹ Real-Time-Scalable Timed Events And Mobility

characteristically associated with an event channel may require adaptation to maintain the required real-time guarantees while allowing changes to membership and topology as is typical of VANETs. This Space-Elastic model is the first to directly address adaptation in the space domain to maintain real-time guarantees.

2. THE SPACE-ELASTIC MODEL

Hard real-time event communication in a highly dynamic VANET is challenging. We exploit the proximity filters defined by RT-STEAM to reduce the scope of a VANET where real-time communication is required to the area within the defined proximity bounds only. The dynamics of a VANET impact the real-time guarantees available within the proximity bound. In this case, we dynamically adapt the proximity bound to maintain the required real-time guarantees. This dynamic proximity, or space-elastic, adaptation is at the core of our Space-Elastic model.

2.1 Space-elastic Adaptation

The timeliness properties of hard real-time applications are invariant and require guarantees that $P(\text{within } T \text{ from } t(e)) = 1$, i.e., the probability that a time-dependent execution will complete **within** (a time interval) T **from** (start time) $t(e)$, is guaranteed for all timeliness properties. Assuming $\text{space}(E, \text{size})$ is a function that bounds a geographical space of a defined **size** in relation to entity E , (a mobile or stationary node in a VANET), the timeliness property within a specified proximity becomes $P(\text{within } T \text{ from } t(e) \text{ over space } (E, \text{size}) \text{ only}) = 1$. In Figure 1, the circle represents the proximity-bound for real-time guarantees, relative to an entity, E , for example representing a traffic light.




Figure 1: Timeliness properties within a proximity-bound only

The Space-Elastic model assumes that real-time applications are space-aware and defined proximity bounds are adaptable. A space-elastic application knows the minimum proximity where real-time guarantees are critical and adaptation between the maximum (desired) and minimum (critical) proximity is tolerated. Failure to guarantee real-time communication in the critical proximity is a real-time failure, the consequences of which and the actions to arise are determined by the characteristics of the real-time class. For example, failure to guarantee hard real-time properties in the *critical* region is a critical failure, and transition to a fail-safe or fail-operational state may be possible.

Figure 2 illustrates proximity adaptation, in this case motivated by route failures in the desired proximity rendering real-time communication no longer possible. In this example, real-time guarantees are available in the area highlighted as the adapted proximity-bound only.




Figure 2: Adaptive space-elasticity

For a space-elastic application to benefit from proximity adaptation the behavior of the space-elastic application must also be adaptable. In the vehicle platooning scenario, a desired proximity for inter-vehicle real-time communication spans the platoon. If communication within the platoon starts to degrade, for example, due to obstacles in a city environment, the platoon may adapt behavior accordingly, e.g., reducing speed, until such a time when real-time communication within the original desired proximity is achievable.

2.2 Real-time Communication Architecture

To achieve real-time guarantees, the communication architecture used by RT-STEAM combines a proactive, mobility-aware routing and resource reservation protocol (PMRR), at the network layer, with a predictable time-bounded medium-access control (TBMAC) protocol [3]. The PMRR, attempts to discover and maintain real-time constrained routes within a proximity-bound, e.g., as specified by the Space-Elastic Model, utilizing the cellular structure of TBMAC coupled with admission control to establish real-time routes spanning potentially multiple TBMAC cells.

3. ACKNOWLEDGMENTS

The work described in this paper was partly supported by the Irish Higher Education Authority's Programme for Research in Third Level Institutions cycle 0 (1998-2001) and by the FET programme of the Commission of the European Union under research contract IST-2000-26031 (CORTEX).

4. REFERENCES

- [1]. R. Meier and V. Cahill, "Exploiting Proximity in Event-Based Middleware for Collaborative Mobile Applications". in *Proceedings of the 4th International Conference on Distributed Applications and Interoperable Systems (DAIS'03)*, Paris, France, 2003, Springer-Verlag Heidelberg.
- [2]. René Meier and Vinny Cahill. Taxonomy of Distributed Event-Based Programming Systems. in *Proceedings of the International Workshop on Distributed Event-Based Systems (ICDCS/DEBS'02)*, Vienna, Austria, 2002, 585-588.
- [3]. R. Cunningham and V. Cahill, "Time bounded Medium Access Control for Ad Hoc Networks". in *Principles of Mobile Computing (POMC'2002)*, Toulouse, France, October 30-31, 2002.