# The Managed Motorway: Real-time Vehicle Scheduling - A Research Agenda -

Vinny Cahill
vjcahill@cs.tcd.ie
Distributed Systems Group
Trinity College Dublin, Ireland

Aline Senart
senarta@cs.tcd.ie
Distributed Systems Group
Trinity College Dublin, Ireland

Douglas C. Schmidt
d.schmidt@vanderbilt.edu
Vanderbilt University
Nashville, Tenessee, USA

Stefan Weber
webers@cs.tcd.ie

Anthony Harrington
harrinac@cs.tcd.ie

Barbara Hughes
hughesb@cs.tcd.ie

Distributed Systems Group
Trinity College Dublin, Ireland

## ABSTRACT

Air pollution, traffic congestion, stress and accidents are common features of today's road transportation experience. New approaches to improving the efficiency and safety of transportation systems are therefore required. Existing work on safe high-speed motorway driving, however, either assumes that vehicles are driverless and/or is limited to local decision making or to one-lane-only reservation systems. This paper describes a novel approach to vehicle scheduling based on real-time hierarchical scheduling, local real-time coordination and real-time inter-vehicle communication. We describe a new model in which road users reserve variable-size slots on motorway lanes, which enables enforcement of timeliness guarantees and adaptive scheduling based on a combination of local and global decisions. We present the research challenges that must be tackled to ensure that this vision of managed motorways becomes a reality.

## Keywords

Managed motorway, real-time, adaptive scheduling.

## 1. INTRODUCTION

Road transportation is the life blood of industrialised economies. Unfortunately, the existing road network, including the motorway system, is becoming increasingly congested due to an increase in the number of vehicles to accommodate and a decrease in the ability to build new and larger motorways. According to [1], congestion costs amount to 50 billion € per year or 0.5% of the European Union's (EU) Gross Domestic Product (GDP), and by 2010 this figure could rise to 1% of EU GDP.

In the EU, the number of cars per thousand persons increased from 232 in 1975 to 460 in 2002. The overall distance travelled by road vehicles has tripled in the last 30 years and, in the last decade, the volume of road freight has grown by 35% contributing to 7,500 km or 10% of the network being affected by traffic jams daily. In the future, therefore, it will be necessary to manage road capacity more efficiently to reduce the effects of road congestion including delay, unpredictable journey times, excessive pollution, and driver frustration. One possible approach is to exploit vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communication and related technologies to coordinate road use.

The approach in this paper allocates available road space (or equivalently travel time) on designated roads to vehicles for the duration of their intended journey based on potentially prioritised requests. This approach is similar to the ways in which road space is now dedicated to public transport vehicles (via bus lanes) or multi-occupied private vehicles (via car pool lanes). In our approach, however, each vehicle is allocated a moving time slot in which to travel for the duration of its journey similar to how time-division multiple access (TDMA) data communication systems allocate slots to messages in transit through a network.

Adherence to travelling in its allocated slot ensures congestion-free travel from source to destination for each vehicle (in the absence of unexpected events considered later). A motorway control system is in charge of dynamic real-time global vehicle (actually slot) scheduling along the motorway, including dynamic adaptation of slot sizes to allow greater speed in the presence of low traffic demand (when greater inter-vehicle safety distances can be accommodated) and slower speeds when nearing maximum road capacity. Vehicle priority is accommodated by allowing different lanes to travel at different speeds subject to service-level agreements in place between the motorway management and drivers.

Either drivers or a virtual on-board cruise control system are responsible for ensuring vehicles remain within their assigned slots. In the former case, drivers are provided with feedback on the appropriate speed at which to travel and their adherence to slots via an on-board driver-information system possibly providing a head-ups display. In this case, enforcement mechanisms are needed to penalise drivers that do not comply with their slot allocation.

An alternative to driver control is a semi-autonomous approach in which drivers retain control of vehicle steering but

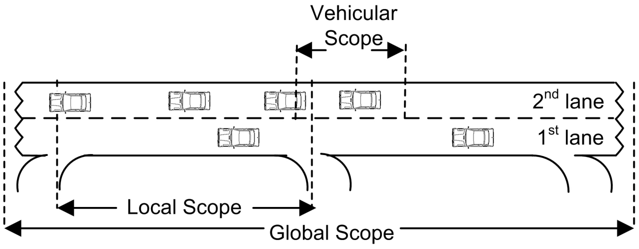


**Figure 1: The Managed Motorway**

their speed is controlled by a virtual cruise control system that ensures the vehicle remains within its assigned slot. We do not foresee vehicles being fully autonomous. However, inter-vehicle coordination may be used to optimise braking times and thus minimise the required slot sizes for collaborating vehicles.

The basic approach of TDMA-based scheduling is well established for data networks. When applied to physical entities, however, it must be adapted to accommodate unexpected local events that might impact on schedules, ranging from drivers who drive outside of their allocated slots, to vehicles that break down, to unexpected obstacles (*e.g.*, pedestrians, animals, or other vehicles) appearing on the road. Global scheduling must thus be supplemented by local real-time coordination of vehicles to accommodate rapid adaptations to local events. The goal of this hybrid global/local scheduling is to allow vehicles within an area to adaptively manage road usage in a way that is compatible with the global schedule.

While similar models have been proposed in the past, we believe that our approach is unique in that it neither requires that vehicles be autonomous nor that any additional road-based infrastructure be provided. Our approach is based on real-time communication and coordination between vehicles and can in principle be applied on any existing stretch of roadway to support existing vehicles that install a suitable driver-information system. Apart from this model, the principle contribution of this paper is the identification of a research agenda that will help realise the model by overcoming technical challenges.

The remainder of the paper is structured as follows. Section 2 presents our model for managed motorways; Section 3 identifies the research challenges that must be tackled to realise this model; Section 4 compares our approach with related work; and Section 5 presents concluding remarks.

# 2. THE MANAGED MOTORWAY

This section describes the architecture of a managed motorway and its characteristics. We explain its elements and the tasks and activities associated with its individual elements. This description will be used in the subsequent sections to explain the research challenges that must be resolved when creating a managed motorway.

## 2.1 Overall Architecture

Our architecture for a managed motorway is shown in Figure 1. The motorway consists of a number of entry and exit ramps and a set of lanes in each direction. In the following discussion, we limit ourselves to just one direction of the motorway and ignore spaces such as hard shoulders, the median, bridges or tunnels. The length of the motorway can be divided into stretches of lanes called *spans* that are demarcated by entry and exit ramps. Cars may enter the motorway through entry ramps into the first lane and may subsequently migrate to other lanes. Vehicles leave the motorway by merging from the first lane to one of the exit ramps.

Entry ramps are fixed elements that include tollgates and queueing lanes. Tollgates provide a means to ensure that only qualified vehicles enter the motorway. Qualifications could include sufficient clearance and proper equipment or software required to use the motorway. Vehicles wait in a queueing lane for their slot to become available before entering the motorway.

## 2.2 Hierarchy of Scopes

The motorway is managed through a hierarchy of scopes ranging from global to local to vehicular scope. Each scope encapsulates certain elements of the motorway and controls access to these elements. Each type of scope is associated with a number of tasks, as described below.

**Global Scope.** The global scope encompasses the motorway in its entirety including a motorway control system and infrastructure to link local scopes to the control system. The responsibilities of the global scope include the management of the capacity of the motorway and the coordination between local scopes that may be needed to control the use of motorway capacity.

**Local Scope.** A local scope covers a span of a motorway, *i.e.*, between entry/exit ramps, and includes a local control system and communication elements, such as 802.11 access points or masts. The responsibilities of a local scope are limited to management of vehicles travelling through that region in compliance with decisions made globally. Local scopes collaborate with neighbouring local scopes to coordinate their activity including handing off vehicles travelling between scopes. Local scopes typically communicate with vehicles within the local scope in real-time.

**Vehicular Scope.** The vehicular scope encompasses a vehicle and its immediate vicinity. It includes physical elements, such as the sensors and actuators mounted on a vehicle, and communication elements, such as local and global transceivers. The area that a vehicular scope covers is limited by the range of a vehicle's sensors. The vehicular scope is responsible for communicating the assigned speed and other information to the driver, and for the distribution of information that has been acquired in the vehicular scope as well as for implementing inter-vehicle coordination protocols as discussed later.

Elements in each scope described above need to collaborate in real-time to manage the motorway, optimise its throughput, and ensure the safety of its vehicles. The novelty of our approach is a TDMA-style management strategy that coordinates the movement of vehicles through a slotted solution.

## 2.3 Slot Management

In our approach, each vehicle on a motorway will be assigned one slot that defines the physical space it may occupy. The management of the slots is used to maximise the throughput of the motorway at the global scope. These slots are managed in a collaborative manner by the local scopes in response to requests from the global scope motorway control system.

For example, tollgates at an entry ramp might decide when a vehicle at the entry ramp may join the first lane and which slot will be assigned to the entering vehicle. The slot size is determined at the local scope as it depends on vehicle demand and the characteristics of the environment. The local scope, however, may also be influenced by the vehicular scope. For example, an accident can be detected at the vehicular scope by sensors mounted on a vehicle. This information is forwarded to the local scope that initiates immediately vehicle re-scheduling.

# 3. RESEARCH CHALLENGES

This section describes key research challenges that must be addressed to achieve real-time vehicle scheduling on managed motorways. We explain why existing TDMA-based approaches are not appropriate and show that hierarchical and adaptive scheduling of vehicles is a challenging problem. We then present local real-time coordination as a mandatory requirement to support real-time adaptive scheduling of vehicles. Finally, we describe the challenge of real-time communication in dynamic vehicular networks.

## 3.1 Real-time Vehicle Scheduling

The high predictability and safety requirements needed by managed motorways motivates the need for real-time dynamic scheduling of vehicles. Ensuring efficient road usage, however, that accounts for deadlines and drivers' constraints while offering dynamic replanning in response to last minute reservations, cancellations or emergencies in real-time is a particularly challenging problem. The control system in charge of the dynamic real-time vehicle scheduling imposes requirements on service-level agreements, admission control and scheduling that make it different from other reservation-based approaches.

### *3.1.1 Service-Level Agreement*

A Service-Level Agreement (SLA) must be negotiated between the drivers and the motorway management system prior to the use of the motorway, *i.e.*, either on-the-fly as vehicles approach the motorway or *a priori* for advance reservations. An SLA typically contains quality-of-service (QoS) commitments, pricing policies and timeliness guarantees. After an SLA has been mutually agreed upon, the motorway control system enforces all the SLA commitments, such as maintaining safe speeds between vehicles. In turn, drivers must behave in accordance with their commitments, *e.g.*, vehicles must travel in their allocated moving time slots (via cruise control or manual driving) to ensure they arrive at their destination on time.

To enforce SLAs, penalties are incurred when the commitments are not respected. For example, vehicle behaviour is typically monitored and repeated failure to adhere to the SLA results in sanctions ranging from monetary penalties, a reduced QoS threshold for future journeys, or even a vehicle being banned from the system.

Conversely, if the motorway cannot meet its SLAs it provides refunds. For example, if an accident prevented vehicles from reaching their destination on time, a refund could be offered based on the number of minutes vehicles were delayed. Rewards can also be disbursed when services useful to achieving SLAs are provided. For example, an SLA may include an extra bonus for vehicles that act as leaders in a platoon, as they consume more fuel than following vehicles.

### *3.1.2 Admission Control*

To ensure safety and that SLAs are met, the number of available slots on a managed motorway cannot exceed a peak. It is therefore necessary to employ admission control that rejects excessive incoming reservation requests so that the remaining reservations can be serviced within their SLAs.

Tollgates located on each access ramp check slot availability via the global scope controller to make sure that the capacity of the motorway is not exceeded. Such admission control can be designed either to respect existing schedules or in the presence of suitably prioritised requests to cause adaptation to existing schedules either locally or globally. For example, the arrival of an emergency vehicle might precipitate local/global rescheduling.

An entrance assistance system can also judge whether vehicles have proper software versions and necessary hardware to accept them on the motorway. Example software include warning applications, secure payment, as well as infotainment applications and hardware, such air pollution sensors and adequate communication capabilities. Vehicles that desire access to the motorway wait in a queueing lane for a slot to become available and/or for software updates. Tollgates use short-range communication and relays between queued vehicles to regulate the admission of cars and perform time-bounded software updates when necessary.

### *3.1.3 Hierarchical and Adaptive Scheduling*

To accommodate unplanned events that might impact vehicle schedules, real-time scheduling should be adaptive rather than static. For instance, when vehicles with advance reservations miss their slot or do not show up, their slot should be reassigned dynamically to queued vehicles. Similarly, if vehicles drive outside their slots, the schedule of nearby vehicles will be affected and therefore must be updated.

Recomputing the schedule for all the cars in a centralised way might not always be feasible when rapid adaptations deviate from the global/local schedule, *e.g.*, if there is an accident or drivers do not fulfill their commitments. In such situations, global/local scheduling is supplemented by local vehicle coordination. In this hybrid approach, scheduling only invoked when the time window available to vehicles to adjust their behaviour is not sufficient to allow safe reaction without approaching vehicles also adapting their behaviour.

Scheduling is multi-level, where global decisions impact local policies and vice-versa. The global level generally focuses on broad scheduling decisions, *e.g.*, ensuring fairness for all access ramps on a motorway. In contrast, the local level executes the global plan in the local context, *e.g.*, the speed of a vehicle and the position of its slot within a local span are typically local decisions. Decisions at a vehicular scope are exceptional and focus on safety-critical adaptations, *e.g.*, braking when animals cross the motorway.

When local adaptations span several vehicles, a feedback loop to the global motorway management system is required to re-evalute new schedules globally.

While a traditional real-time system might attempt to account for the worst-case occurrence of such events in determining journey schedules [2], both the scale and dynamism of a motorway make such an approach infeasible. Even if it were feasible, the impact on vehicle scheduling would likely be such that achievable speeds would be significantly degraded to account for worst-case behaviour. An alternative approach might involve dynamically recalculating journey schedules in response to the occurrence of local events. This approach also appears infeasible, however, due to the need to communicate the occurrence of such an event to the local scope scheduler and then apply the resulting updated schedule in sufficient time, *e.g.*, to avoid hitting an unexpected obstacle.

Our approach provides a more effective multi-level scheduling strategy, where vehicles within a local scope can self-organise locally to adapt to such situations. Such adaptations may affect only a single vehicle (effectively requiring it to depart from its planned schedule). More likely, however, our approach will involve a set of collocated vehicles adapting their behaviour, *e.g.*, stopping, slowing down, speeding up or changing lane. Depending on the system load at a given time, such adaptations could be (1) limited to a set of vehicles in a particular area or (2) propagate the effects to the local scope scheduler as input into schedule recalculation.

## 3.2 Local Real-time Coordination

To ensure safe operation, the self-organisation described above requires that vehicles continuously coordinate their behaviour in such a way that their aggregated behaviour respects system-wide safety constraints, such as maintaining safe distance to other vehicles. Ensuring these system-wide safety constraints when vehicles are moving at high speed implies stringent real-time requirements on coordination.

### 3.2.1 Impediments to Real-time Coordination

A traditional approach to coordination relies on V2V communication to synchronise their behaviour [3]. Since the vehicles are mobile, moreover, they would typically communicate over wireless (possibly *ad hoc*) networks. In general, wireless network communication—and in particular real-time wireless communication—is unreliable and the achievable timeliness varies greatly over time and location [4]. Environment-mediated communication [5]—where vehicles communicate by changing their environment and detecting changes made by other vehicles using sensors—can be used to supplement direct communication. The range and accuracy of sensors, however, are inherently limited and sensor information might be affected by environmental conditions, such as temperature.

These limitations imply that the information vehicles have about other vehicles and their environment varies significantly over time and space. Vehicles may not be able to reach a consensus on their collective behaviour, may not even be able to consult each other, and may not have a complete view of their environment. Vehicles will therefore need to make decisions independently of each other—using only limited information—while still ensuring system-wide safety constraints.

### 3.2.2 Real-time Feedback

To address this challenge, our approach exploits real-time feedback on currently available information provided by novel real-time sensing and communication models (see Section 3.3). In our approach, vehicles adapt their behaviour depending on currently available information, hence making progress when it is safe to do so, while ensuring that safety will not be compromised. Building on the Comhordú coordination model [6], the notion of consensus is replaced by a set of contracts that bind vehicles *a priori* and allow them to make predictions about other vehicles behaviours even when they do not have current information about them. These contracts capture the necessary information for a vehicle to safely begin or continue performing any given action.

At any time during execution, each vehicle can derive the set of actions it can safely undertake, given the information that it currently has about its environment. An intermittent interruption in communication might therefore lead to a following vehicle slowing down conservatively. This adaptation, however, does not necessarily require all vehicles to adapt their schedules.

## 3.3 Real-time Inter-vehicle Communication

Underlying our approach to global journey scheduling and adaptation to unexpected events is a wireless communication service that supports real-time V2I and V2V communication. Due to the potential absence of fixed infrastructure, vehicular networks are usually modelled as mobile *ad hoc* networks (MANETs), *i.e.*, vehicular *ad hoc* networks (VANETs). There is a body of on-going research addressing wireless communication in VANETs [7] as well as ongoing standardisation efforts based on 802.11p [8]. In general, however, this work (especially 802.11p) only addresses best-effort—or at most soft real-time—communication suitable for non-safety-critical applications, such as electronic tolling, media sharing, and emergency-vehicle arrival warning.

### 3.3.1 Impediments to Real-time Communication

The challenge of real-time V2I and V2V communication stems from the specific characteristics of MANETs, particularly their dynamically changing membership and topology. Real-time communication can be characterised as the achievement of bounded and known message delivery delays, in the presence of disruptive factors, such as other real-time traffic, variable load or faults. Achieving hard real-time communication in fixed networks relies on assumptions of known network load, guaranteed connectivity, deterministic communication latency and guaranteed resource availability [9]. These assumptions are not applicable in wireless networks due to their intrinsic dynamics. For example, the movement of wireless nodes yields time-varying connectivity between nodes.

In a typical MANET, the lack of fixed infrastructure limits the transmission range of mobile nodes and requires wireless hosts to act as routers if needed. The result is a distributed multi-hop network with a time-varying topology where routes are typically short-lived [10]. Neither the assumption on known communication latency nor on guaranteed connectivity is therefore applicable in this domain. In addition, multi-hop *ad hoc* communication has the potential to incur non-deterministic medium-access latency at each hop leading to non-deterministic latency for end-to-end communication. In summary, achieving hard real-time communication in MANETs is impossible in the general case since the above assumptions cannot be achieved without restrictions on network dynamics.

#### 3.3.2 Space-elastic Communication

To address this issue, in prior work, we defined the space-elastic model for real-time communication in wireless networks, including *ad hoc* networks [11]. In (*ad hoc*) wireless networks, varying link quality and network topology mean that it is impossible to guarantee communication with hard real-time requirements in a fixed geographic area. The space-elastic model exploits the rationale observed in [12], *i.e.*, that context information is relevant to particular geographical areas, to ensure periodic communication with real-time requirements only within a dynamic (varying over time) proximity. In addition, real-time message senders are notified of changes in the proximity within a bounded time, hence allowing them to adapt their behaviour to current communication conditions.

In the space-elastic model, message senders specify the required latency for message delivery and the area (called the "desired coverage") in which messages should be delivered. The size of the area in which timely delivery of messages is provided (called the "actual coverage") can change over time. This model ensures that message senders can communicate within the specified latency in the actual coverage and will be notified within a bounded time if this coverage changes. Moreover, receivers will receive all messages of types in which they have expressed interest if they are present within their senders' actual coverages at the delivery time of the messages.

The space-elastic model assumes that applications are space-aware, *i.e.*, that they can specify and interpret bounds in space, reliably. The feasibility of the space-elastic model with low-jitter real-time communication and time-bounded adaptation notification has been demonstrated in real (as opposed to simulated) settings [11].

#### 3.3.3 Space-elastic Communication in VANETs

The space-elastic model shows the feasibility of providing real-time wireless communication in arbitrary MANETs. The protocols underlying the current implementation of the space-elastic model, however, are unsuitable for safety-critical cooperative automotive applications. This limitation stems from their being designed to handle the general case of MANETs with slow-moving nodes and, in particular, to trade off protocol performance (in terms of achievable latency and bandwidth usage) for predictability.

A suitable communication substrate for the managed motorway thus requires the design of a suite of novel protocols to support the space-elastic model that are tailored to particular features of automotive applications, including (1) their constrained and predictable movements patterns (*i.e.*, along roads), (2) the fact that groups of vehicles tend to travel together for long (minutes) periods of time while traversing the same routes, and (3) the possible presence of roadside communication infrastructure as foreseen by initiatives such as the Vehicle Infrastructure Integration (VII) initiative [13].

## 4. RELATED WORK

This section compares our work with related research on safe and congestion-free managed motorways.

### 4.1 Motorway Traffic Monitoring

Work on motorway traffic monitoring includes the Vehicle Information Communication System (VICS) [14], the Advanced cruise-assist Highway System (AHS) [15] and the FleetNet project at NEC [16]. VICS has been in service in Japan since 1996 and delivers traffic information to car drivers nationwide. This information system gathers real-time traffic condition from roadside sensors deployed on motorways and main roads throughout the country, and provides road traffic information to drivers via telematics units. Most Intelligent Transportation System (ITS) experts consider VICS as one of the most successful ITS deployments in the world [17] with over 26.1 million units sold in 2007 [18].

AHS is a more advanced system that provides a practical driving-assistance system in Japan by using vehicle sensing and communication capabilities in addition to those offered by the motorway infrastructure. The cooperative vehicle-motorway system advises drivers of the best route to take depending on traffic condition and warns them about out-of-sight hazards. Similarly, NEC Europe Network Laboratories are developing ITS technologies to enhance road safety. In their FleetNet project, a wireless multi-hop *ad hoc* network for inter-vehicle communication has been developed that transmits traffic information to improve driver comfort and safety.

The three systems described above allow drivers to obtain real-time congestion information from other vehicles or/and the road infrastructure and can therefore avoid the congested areas. Road congestions are not prevented before they occur, however, which is a novel element of our managed highway approach.

### 4.2 Safe Automated Transport

Personal Rapid Transit (or PRT) is an alternative mode of transportation that offers a potential solution to traffic congestion and improved travel accessibility [19]. PRT consists of small vehicles under automatic control running on a network of specially-built guideways. Even though PRT addresses a number of public mass transit weaknesses, such as fixed timetabling and limited routes, a number of issues remain unresolved [20]. The limitations of PRT are similar to rail-based automated guiding vehicles [21], which both suffer from limited capacity and requires the construction of additional expensive infrastructure.

Rather than building specialised infrastructure, new projects have moved towards the design of Automated Highway Systems (AHS) that provide transportation for driverless vehicles [22]. These vehicles organise themselves in platoons to reduce traffic congestion on motorways [23, 24]. While these projects are exciting, they assume the existence of fully autonomous vehicles, which is well beyond the state-of-the-art today and thus do not provide viable solutions for today's transportation challenges.

### 4.3 Virtual Vehicle Slots

The last category of related work that investigates managed motorways considers the use of virtual vehicle slots to guarantee time-bounded journeys. [25] presents a vision of human-controlled cooperating vehicles that explores the concept of time-space corridor to guarantee motorway lanes free of traffic congestion. This approach, however, only uses local coordination between vehicles to regulate the flow of vehicles. In contrast, our approach combines real-time local and global decisions, which makes it more appropriate to reschedule of road usage in ways compatible with the global schedule.

The lane reservation system for highways [26] follows a similar slot-based approach where drivers must reserve a slot when entering a motorway. The system however does not provide congestion free travel for all the lanes as only one high-priority lane is supported. Moreover, this approach does not focus on real-time scheduling and coordination, whereas ours does.

# 5. CONCLUDING REMARKS

This paper described a new approach to vehicle scheduling for managed motorways. In this approach drivers travel in their assigned moving slots from departure to arrival to help reduce traffic congestion, accidents and car emissions. These slots have guaranteed speed and safety distances to other slots, allowing the assurance of vehicle SLAs. Such enforcement of timeliness guarantees relies on adaptive vehicle scheduling.

We have identified the key research challenges that must be tackled to realise these managed motorways: real-time hierarchical scheduling, local real-time coordination and real-time inter-vehicle communication. Our current work is tackling these challenges both analytically and experimentally by prototyping this model in a representative testbed environment.

# 6. ADDITIONAL AUTHORS

Additional authors: Kulpreet Singh (Distributed Systems Group, Trinity College Dublin, Ireland.
email: `singhk@cs.tcd.ie`)

# 7. REFERENCES

[1] Commission of the European Communities. Raising awareness of ICT for smarter, safer and cleaner vehicles. http://ec.europa.eu/information, 2006.

[2] H. Kopetz. *Real-Time Systems: Design Principles for Distributed Embedded Applications*. Springer, 1997.

[3] A. L. Murphy, G. P. Picco, and G.-C. Roman. Lime: A coordination middleware supporting mobility of hosts and agents. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 2006.

[4] B. Hughes and V. Cahill. Achieving real-time guarantees in mobile wireless ad hoc networks. In *24th IEEE Real-Time Systems Symposium Work-in-progress Session (RTSS'03)*, pages 37–40, December 2003. Work In Progress.

[5] E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press US, 1999.

[6] M. Bouroche, B. Hughes, and V. Cahill. Real-time coordination of autonomous vehicles. In *IEEE Conference on Intelligent Transportation Systems (ITSC 06)*, pages 1232–1239, September 2006.

[7] ACM Press, editor. *Proceedings of the 4th International Workshop on Vehicular Ad Hoc Networks*, Montral, QC, Canada, September 2007.

[8] IEEE P802.11 Task Group p. Status of project IEEE 802.11 task group p. http://grouper.ieee.org/groups/802/11/Reports/tgp_update.htm. Accessed October 2007.

[9] P. Verssimo and L. Rodrigues. *Distributed Systems for System Architects*. Kluwer Academic, 2001.

[10] K. Wang and B. Li. Efficient and guaranteed service coverage in partitionable mobile ad hoc networks. In *21st Annual Joint Conference of the IEEE Computer and Communications Societies*, New York, USA, 2002.

[11] B. Hughes. *Hard Real-Time Communication for Mobile Ad Hoc Networks*. PhD thesis, Department of Computer Science, Trinity College Dublin, 2006.

[12] H. Hartenstei, B. Bochow, A. Ebner, M. Lott, M. Radimirsch, and D. Vollmer. Position-aware ad hoc wireless networks for inter-vehicle communications: the FleetNet project. In *2nd ACM International Symposium on Mobile Ad Hoc Networking and Computing*, Long Beach, California, USA, 2001.

[13] United States Department of Transportation. Vehicle infrastructure integration. http://www.its.dot.gov/vii/. Accessed October 2007.

[14] Vehicle Information and Communication System Center. Accessed October 2007. http://www.vics.or.jp/english/.

[15] Advanced cruise-assist highway system research association. Accessed October 2007. http://www.ahsra.or.jp/index_e.html.

[16] M. Maekawa. ITS (Intelligent Transportation Systems) Solutions. *NEC Journal of Advanced Technology*, 1(3), Summer 2004.

[17] S. S. Saad. ITS in Japan, a different approach to transportation policy. *World Review of Intermodal Transportation Research*, 1(1):45–54, 2006.

[18] Infrastructure Japan Ministry of Land and Transport. VICS is in full steam ahead, 2007. http://www.its.go.jp/ITS/ITSinJapan/navi.html.

[19] N. X. H. Viet, H.-S. Song, and K. Nam. Locating power supplies on a personal rapid transit system to minimize system losses. *IEEE Transactions on Industry Applications*, 40(6), 2004.

[20] W. D. Cottrell. Moving to the mainstream. In *10th International Conference on Automated People Movers*, Orlando, Florida, USA, May 2005.

[21] L. Yang and F.-Y. Wang. Driving into intelligent spaces with pervasive communications. *IEEE Intelligent Systems*, 22(1):12–15, January 2007.

[22] S. M. Easa. Automated highways. *Wiley Encyclopedia of Electrical and Electronics Engineering*, 1999.

[23] R. Horowitz and P. Varaiya. Control design of an automated highway system. *Proceedings of the IEEE*, pages 913–925, July 2000.

[24] B. Li, X. Wang, H. Cai, and C. Wang. The R&D strategy for intelligent highway system of China. In *IEEE Intelligent Transportation Systems*, volume 1, pages 569– 573, October 2003.

[25] R. Morla. Vision of congestion-free road traffic and cooperating objects. Sentient Future Competition, November 2005. http://www.embedded-wisents.org/competition.

[26] N. Ravi, S. Smaldone, L. Iftode, and M. Gerla. Lane reservation for highways (position paper). In *10th International IEEE Conference on Intelligent Transportation Systems*, Seattle, Washington, September 2007.