A Vision for Wireless Access on the Road Network

David N. Cottingham, Jonathan J. Davies
Computer Laboratory, University of Cambridge
Cambridge CB3 0FD, United Kingdom
E-mail: {david.cottingham, jonathan.davies}@cl.cam.ac.uk

***Abstract*— Metropolitan area wireless networks are currently being deployed in major cities around the world, whilst in tandem there has been much research into vehicle-to-roadside communication. New applications for vehicular networking become possible as blanket, low-cost, wireless networks begin to exist across cities, resulting in Connected Traffic, rather than isolated Connected Cars. We classify these applications according to their distinguishing characteristics, and discuss their network architecture requirements. We outline our current work on the language and compiler support required in order to deploy applications on such networks, and how coverage mapping algorithms will enable better prediction of network conditions to optimise such deployments. We conclude that such analysis and tools are important for defining the future of wireless access for the road network.**

I. Introduction

Applications for vehicular networking are many and varied, with safety features such as collision avoidance, crash notification and information dissemination frequently cited. Others include electronic tolling, freight and asset tracking and traveller information systems. However, the majority of such applications do not focus on what would be possible were a large number of cars equipped with communication infrastructure that could be used in cities where (near)-blanket low cost wireless network coverage existed. Whilst at present this is not the case, we are already seeing limited deployments of WiFi networks in major cities. Hence, it is important to conjecture what applications might be possible. This paper seeks to outline that vision.

In the future, we are likely to see a combination of two initiatives for wireless access on the road network. Firstly, governments are already beginning to deploy converged networks along motorways, such as the UK's National Roads Telecommunications Services (NRTS) for carrying data from CCTV and roadside emergency phones. These networks will have sufficient capacity to allow both government and commercial enterprises to offer value-added services built on wireless access to the vehicles travelling on the road network. Secondly, both the private and public sectors will roll out wireless networks that cover large areas of cities.

Ubiquitous GSM/UMTS networks already exist. However, current pricing models make their usage infeasible for large amounts of data. Although prices may decrease, there still exists a very high cost barrier to new entrants competing in the market. Wireless networks such as WiFi, or other technologies, that can be deployed incrementally are an alternative. Competition between providers will drive costs down, since deployment costs are low compared to cellular networks. Bandwidth is also a key issue, with current cellular networks offering less than 1 Mb/s connectivity, though in the future this is set to increase. Future applications will require throughputs that exceed this provision.

Whilst we envisage that wireless access on the road network will come about, we note that such access will in the medium-term remain fragmented. Metropolitan networks will have "black spots", and we must therefore assume that handovers between networks and periods of disconnection will be the norm, rather than the exception. Hence, we do not rule out the use of cellular networks in providing network access for applications, but for the reasons described above we expect it to be used only where essential.

In addition to advances in wireless networking, modern vehicles contain many processors employed for diverse purposes. We believe that trends in decreasing power requirements, size and cost of manufacture mean that in future we can expect vehicles to provide embedded computing platforms for the execution of arbitrary user applications. This will allow applications to be deployed over the network that make use of such a distributed computing architecture.

II. Connected Traffic

The concept of the "Connected Car" has existed for some time, with many projects deploying a variety of communication technologies in vehicles. Small numbers of vehicles (such as individual companies' fleets) are connected to the Internet continuously, via whichever network is available at the time. We now propose the concept of *Connected Traffic* to mean wireless access on the road network that provides intermittent, high bandwidth, low cost connectivity to almost every vehicle. Connected Traffic therefore revolutionises our notions of what applications we can deploy over such a communications network. To aid our discussion, we first classify the types of applications into five categories from a network-oriented perspective.

A. User-Centric

Initial applications of wireless access for the road network will centre around information retrieval by drivers from existing Internet services, such as maps or coarse-grained traffic information. The network requirements for these services are thus quite low, with the utility of the network increasing with its coverage.

Such *User-Centric* services are already available, but their use will increase significantly if access to them becomes more convenient and at a lower price. Use of services such as podcasted traffic bulletins will increase, as mobile access becomes more viable. Information dissemination will move to a more on-demand approach as compared to the current fixed schedule regime.

B. Fixed-to-Mobile

Once network coverage transitions from isolated to near-ubiquitous, both private and commercial users will utilise applications that require two-way communication between a vehicle and a fixed host on the Internet. This will allow both information reporting, and remote operation of devices in either direction. We term these *Fixed-to-Mobile* applications.

In a similar vein to applications which are User-Centric, examples of Fixed-to-Mobile already exist. However, they currently utilise low throughput and high cost networks, and are hence only the domain of relatively niche commercial operations. With widely deployed wireless access, applications such as vehicle and asset tracking, automated accident notification, remote operation (e.g. of devices at home from a vehicle), and electronic payment on the move (e.g. for congestion charging) become possible.

C. Network-Pull

The use of vehicles, particularly private cars, as mobile sensor-platforms [7] will increase, developing from the current pay-as-you-drive insurance platforms where only GPS location data is gathered. Data from sensors conducting pollution or environmental monitoring could be used by central agencies to build up far more accurate maps of how these parameters vary on much finer scales than are currently achievable. These *Network-Pull* applications involve the network supporting the querying and collection of sensor data from potentially millions of vehicles. The supply of such data will be on a voluntary basis, and will most likely involve an incentives system (e.g. free access to the communications network if environmental data is uploaded) in order to encourage participation.

D. Network-Push

Applications will evolve from solely passively uploading information from sensors, to include obtaining aggregated dynamic information from the network. This is distinct from User-Centric applications, where static information not derived from other nodes in the network is downloaded. Applications will therefore move to a *Network-Push* model, where vehicles subscribe to streams of data from the network, which generates such data by processing large numbers of vehicle sensor readings. In this way, the incentive for a user to contribute is the utility of the information they obtain in return.

E. Peer-Centric

As networks for roads evolve, we will see users begin to share information between each other, rather than communicating solely with centralised services. Although such exchanges may happen using vehicle-to-vehicle communication, it is likely that a more structured and long-range service will be required, implying that a roadside communications network will be used for peer-to-peer communication. Applications in this category are therefore *Peer-Centric*.

III. Example Applications

In order to show how our classification scheme can be used, we now list various applications, briefly explain their purpose, and characterise their types. Table I summarises how each can be classified.

- **Entertainment.** In-vehicle web browsing and streaming media are instances of accessing existing Internet content in a mobile context and so are User-Centric.
- **Mobile Commerce.** Purchasing products whilst on the move [21] involves two-way communication to known Internet hosts and so are Fixed-to-Mobile.
- **Remote Operation.** The operation of devices in a remote location such as the home or the office, e.g. switching the heating on, are also Fixed-to-Mobile applications.
- **Asset Tracking.** A delivery or haulage vehicle reports its position and the items it contains in a Fixed-to-Mobile fashion.
- **Congestion Information.** In its most primitive form, ascertaining current road conditions can be performed in a User-Centric manner, perhaps where the in-vehicle navigation unit receives a broadcast of the relevant information. Instead, the information, known as floating car data [8], could be gathered in a Network-Pull manner by observing the movements of vehicles. Further, if the vehicles notify the network about their journey, pertinent information can be delivered in a Network-Push fashion.
- **Real-Time Weather.** Similarly, vehicles could download current weather observations and forecasts in a User-Centric manner. In addition, the network could gather data from vehicles containing meteorological sensors in a Network-Pull fashion, or further redistribute the aggregated data back to interested vehicles in a Network-Push manner.
- **Road Hazard Detection.** Potential hazards on the roads could be detected if data from vehicles' braking systems is gathered by the network in a Network-Pull manner. When a substantial number of vehicles brake sharply at a particular location, it could be marked as a potential hazard. Again, this could be a service to which vehicles subscribe and receive notifications of upcoming hazards from a Network-Push application.
- **Map Generation.** GPS traces are obtained from large numbers of vehicles, and combined to update digital maps in real-time [6] in a Network-Pull application.
- **Slot Booking.** Motorway slot-reservation and ramp metering [17] systems could be implemented by vehicles communicating with a known Internet host to negotiate timing and payment, so are Fixed-to-Mobile applications.

Application	U-C	F-M	N-Pull	N-Push	P-C
Entertainment	•				
Mobile Commerce		•			
Remote Operation		•			
Asset Tracking		•			
Congestion Information	•		•	•	
Real-Time Weather	•		•	•	
Road Hazard Detection			•	•	
Map Generation			•	•	
Slot Booking		•			
Fleet Management			•		
Gaming		•			•
Congestion Charging		•			•
Collision Avoidance					•
Accident Notification					•
Transport On Demand					•

TABLE I

CLASSIFICATION OF APPLICATIONS BY TYPE

- **Fleet Management.** Organisations owning a number of vehicles need to be able to manage them centrally, perhaps so that their routes can be optimised. This is possible by allowing a Network-Pull application to query the location of each vehicle.
- **Gaming.** Occupants of vehicles playing games—perhaps location-aware games—with occupants of other vehicles are participating in Peer-Centric applications. If non-mobile participants are also involved, it will be Fixed-to-Mobile.
- **Congestion Charging.** Many suggested implementations of electronic toll collection or congestion charging schemes involve the transmission of location data to a governmental organisation in a Fixed-to-Mobile fashion. Other researchers have proposed a Peer-Centric implementation of congestion charging that preserves the privacy of the users [2].
- **Intersection Collision Avoidance.** Self-organising co-ordination of the movement of traffic, such as the negotiation between vehicles approaching a road junction [9], are necessarily Peer-Centric in nature.
- **Accident Notification.** In-vehicle systems which detect collisions and automatically notify the emergency services of the location and the nature of the collision are instances of Fixed-to-Mobile applications.
- **Transport On Demand.** A user's network-connected device will proactively suggest not only routes to travel to a destination, but also modes of transport, using the aggregated information it receives. The user can then participate in Peer-Centric transport auctions, where different vehicles make bids for their respective charges to convey the user to their destination.

IV. NETWORK ARCHITECTURES

Each class of application introduced in Section II possesses different network architecture requirements. We now analyse those requirements in detail.

A. User-Centric

User-Centric applications are concerned with obtaining data from servers located in the fixed Internet. The quantities of data concerned can be small, such as traffic updates, or large, as in the case of podcasts or digital map updates. Quality of Service (QoS) requirements are nor particularly onerous; data is downloaded in the background to the vehicle for later use, or streaming media is heavily buffered.

Network conditions can therefore vary, handovers between different networks can take place, and patchy coverage is acceptable. Provided that the average throughput is sufficient to obtain the data in a reasonable time (e.g. traffic information within 5 minutes, or a connection to the streaming media source every 5 seconds), user requirements will be satisfied.

The communications paradigm in such applications is simple: data is either broadcast over a medium that the vehicle can listen to, or the system onboard the vehicle uses a ubiquitous connection such as a cellular link, or a localised one such as a WiFi hotspot to request the data. A lookup takes place for a server that is explicitly specified by name. Transfer of information takes place over standard unicast or broadcast protocols, and is largely in one direction.

B. Fixed-to-Mobile

Fixed-to-Mobile applications differ from User-Centric in that communication is no-longer predominantly one-way. The throughputs required vary from low in applications such as remote home-device operation, to high in the case of voice or video calling. However, whilst communication is two-way, it always involves the vehicle registering itself with the service concerned, whether this is initiating each connection, or providing its identifier to a lookup service. Depending on the application, synchronous or asynchronous communications may be required, though any application with which user interaction takes place is likely to require a fixed QoS. For this to be the case a (near-)ubiquitous network is needed.

Communication from a fixed source to the vehicle takes place over unicast connections. Data may be encrypted such that only that vehicle can decode it. Crucially, there must exist a naming service that keeps track of how the vehicle may be contacted. Technologies such as Mobile IP become important in allowing seamless mobility. In-network proxy servers that buffer packets whilst network handovers take place, and versions of transport protocols that are optimised for such varying network conditions [1] are also important.

C. Network-Pull

Network-Pull applications centre around retrieving data from vehicles that act as mobile sensor-platforms. They are therefore not concerned with contacting specific vehicles by their identifiers, but instead obtaining data concerning a particular geographical area or captured during a specified period of time. Applications will generate spatio-temporal queries which the network will then translate into vehicle-specific requests.

In order to achieve this, the network may be aware of the approximate locations of all vehicles, or it may broadcast queries to a particular geographic region. Privacy is therefore traded-off against the ability to obtain a richer dataset. Query languages, such as that used in the Cartel project [11], must be developed, along with scalable location directory services such as Grid [13] or the Landmark hierarchy [20].

An important consideration for a sensor data-rich network is security. Whilst data such as environmental conditions can be placed in the public domain, fleet management information will need to traverse the network in an encrypted form such that only the fleet owner can read it. Network designers must therefore ensure that query languages and location services take such data-hiding requirements into account. Privacy aspects of Network-Pull applications also influence how query languages are designed. Data suppliers will wish to decide whether information is artificially delayed or the position blurred before it is uploaded, whilst the network should also ensure that anonymity is aided by using, for example, mix zones [3].

The QoS required by Network-Pull applications will depend both on the update frequency and the degree of confidence requested. The network will need to give applications estimates of what levels can be expected for different geographical areas at different times. For example, it might be that pollution data can be obtained in large volumes from a city ring-road during rush hour (when there are many vehicles present), but overnight data is much more sparse. The lower the rate of sensor-platforms traversing the area, the more time will be needed for the necessary number of readings to be amassed to provide a particular level of confidence. Also, the QoS for a particular area will be affected by the number of users competing for access to radio channels, as well as the available network technologies. Hence, designers must consider demand not from a mobile client's perspective, but rather from the view of a fixed client requesting data from multiple mobile servers. In addition, the number of fixed clients who will aggregate data is likely to be few compared to the number of mobile servers.

D. Network-Push

An evolution of the Network-Pull paradigm, Network-Push applications will involve vehicles supplying sensor data, and subsequently receiving the processed values. The network will now also need to provide on-demand aggregated information, and provide mechanisms for automated pushing of data to mobile clients. Middleware using publish/subscribe message queues will need to be deployed, and QoS requirements will now also include the capacity of the message server.

The network architecture for the physical and data-link layers will be similar to that used for Network-Pull. However, multicast communication will be key in ensuring efficient bandwidth utilisation, and transport protocols will be more geared to ensuring that "enough" (rather than all) the data transmitted is received by the client, perhaps by using fountain codes [15].

Also, vehicles requiring aggregated information from the network will have a feedback effect on what data is required to be queried from the network. Network-Push applications will produce many mobile clients, each of whom will make a specific spatio-temporal request. This will in turn make the level of network provision required vary far more dynamically.

E. Peer-Centric

Peer-Centric applications will develop as processing resources on vehicles increase, and for some applications it becomes viable to bypass a central server altogether. This will involve the use of ad-hoc networking for gossip purposes, such as the MDDV protocol [22]. Vehicles will communicate information between each other in order to relieve the backbone of some of the overhead of transmission to every node, similar to the route optimisation technique used in the Mobile IP protocol [12]. Requests will be directed first to a geographically-local group of nodes, to ascertain whether the data required (or a subset of it) is available. If it is not, the request will then be made to the fixed network.

In addition, vehicles will act as data-mules, aggregating data from fixed sensors and transferring it up a hierarchy, such as private cars uploading their data to public buses, whose routes are fixed and have a well-defined schedule [14].

Peer-Centric applications will also need to cope with the entire network topology constantly changing, whereas with Network-Push only the final or initial hops were subject to such churn. Networks may require some fixed, but isolated, network nodes to provide predictable hops in the ad-hoc topology. In the case of safety-critical applications, such as intersection collision avoidance, it may be beneficial to involve a fixed node in order to provide QoS guarantees; certainly this type of application will be hardest to deploy on a network architecture which is purely ad hoc.

V. Application Construction

We now move from the network requirements of applications to consider the design, development and deployment of the applications themselves given the Connected Traffic paradigm. Our current work examines how applications can be split up to be executed in multiple processing nodes across the network.

In a road system with many vehicles, we have a rich network of processors with vehicles' on-board computers along with any fixed computing facilities provided by public or private bodies. These processors could be used to execute arbitrary parts of an application. Network-Pull and Network-Push applications typically work best when the results are derived from as large a number of inputs as possible. Thus, it is best to obtain data from a large number of vehicles. The decision about where the data is processed may affect the number of vehicles which contribute to the result.

For a given application, there is a broad spectrum of potential architectures which could be adopted. At one end of the spectrum is the *fully centralised* approach where all vehicles upload their data to a central server (or server farm) responsible for processing all of the data. However, if a large number of vehicles is to be supported, the large communication bandwidth required makes this approach impractical.

Alternatively, *regional servers* could be employed which are each responsible for processing the data gathered from within their geographical region. These servers may be interconnected by a backhaul network to allow sharing of the processed data. In many applications, the processed data will be more concise than the original data. This means that less bandwidth is required to transport the outputs from processing than the inputs, implying that it is most favourable to process data as close as possible to its source. However, this must be offset against the desire to incorporate as many vehicles' data into the processing as possible.

At the other end of the spectrum, there are no centralised processors and it is the responsibility of the vehicles themselves to perform the processing of the data. This could either be organised in a *peer-to-peer* fashion or with the assistance of *public data caches* which merely store data associated with a geographic region. These caches need not be connected with each other and act in a standalone manner, with vehicles acting as *data mules* [18] to share data between caches.

Unlike fixed networks, the network topology in Connected Traffic changes rapidly as vehicles move at high relative speeds. Network coverage is patchy and link bandwidths change. Thus, the application programmer can not make an optimal decision about which architecture is most appropriate to use. Rather than making the decision at design-time, it is preferable if the decision about how the application is split up and distributed to processing nodes is made automatically at run-time. Furthermore, the failure of processing nodes, or changes in the ability to communicate with them, also need to be dealt with. This demands language and compiler support to permit the application programmer to describe the tasks which constitute the application and the flow of data through them, to allow the application to self-organise its execution within the network in the optimal manner.

In order to establish the optimal mapping of application tasks onto processing nodes, the compiler needs to understand the notion of optimality. We can consider the suitability of the possible mappings with respect to various metrics. Three metrics which are relevant to many typical applications are:

- **Total execution time.** If the size of each task is known to the compiler along with the speed of the available processors, the duration of the execution of the application for a given mapping can be deduced.
- **Quality of the result.** The quality of the result of some processing is related to the number of vehicles whose data contributes to the result. A centralised configuration may process data from a larger number of vehicles than a peer-to-peer configuration.
- **Privacy.** The level to which the privacy of the originators of the data is respected is important in applications where personally-identifiable data is processed. Privacy is protected most strongly when individuals' data is mixed with others' data before the data becomes visible publicly.

For typical applications, the fully centralised approach has a higher execution time than the approach which uses regional servers because there is less parallelisation. On the other hand, the quality of result is better for the centralised approach. Privacy is poor in a centralised approach because an attacker observing the data arriving can see everything, but the attack must be distributed if a decentralised approach is adopted.

The problem of searching for the optimal mapping throughout the space of possibilities defined by any combination of such metrics is NP-complete, so is not practical to achieve in real-time. Instead, we must settle for an approximate approach to yield a near-optimal solution.

We have implemented a framework in which applications can be described in terms of constituent tasks; the network can be defined in terms of constituent processing nodes and the communication links; and the metrics against which to optimise are specified. A prototype implementation of the compiler uses heuristics to find a near-optimal mapping of tasks to processors.

VI. Coverage Prediction

In order for many of the applications we have described where vehicles are used as nodes for distributed processing, their levels of network connectivity must be predictable. Whilst in our vision of Connected Traffic network connectivity will be near-ubiquitous, there will still exist small areas where radio shadowing will cause connection drop-outs or reduced performance, as in the cellular network today. There is therefore a need for the coverage along roads to be accurately and efficiently mapped, and these data constantly updated. Our current work therefore examines methods for the aggregation, storage, and use of coverage maps.

Large amounts of data can be collected by using vehicles as sensor platforms, including data on the signal strength or throughput experienced on the various wireless interfaces that the vehicle possesses. This can be uploaded to a central server, or distributed between multiple nodes for processing. Signal strength data, particularly from moving vehicles, has a relatively large variability, due to obstructions by other vehicles, meteorological conditions, and time-varying multipath effects. Hence, this data will have significant amounts of noise.

The objective of coverage mapping is to aggregate this large corpus of data into a much simpler form, that can be efficiently stored in a database on the vehicle, and queried in minimal time. Simple smoothing of the data is not an option, since it is crucial that large drops in signal strength are captured, but that may cover only a small area of road. We are therefore developing algorithms that will preserve significant points of the graph, whilst removing random noise. Long term time variability of signals must also be included in any model.

Once a coverage map exists, vehicles can then predict what their likely network performance will be, and adjust parameters such as TCP timers, streaming media buffer sizes, and media bit rate. Handovers between different networks can be more easily anticipated, and the decision as to whether to perform route optimisation in Mobile IP-based connections will be more informed. Such coverage maps are also crucial for predicting how nodes' reachability will vary for task partitioning applications, where the throughput of the constituent communication links is a factor in determining the optimal mapping of tasks to nodes.

VII. Current Deployment

There has been a great deal of research into network access for vehicles, including the Drive-Thru Internet project [16], various projects examining 802.11b performance [19], [10], and work on 802.11a [5]. However, little has been done analysing the performance of larger scale network deployments in urban environments. We are now seeking to analyse how our work on distributed processing and coverage mapping performs in a real-life scenario.

We are currently in the process of deploying a small test network around the roads on our site in order to test our equipment and its backhaul system, which will be based on IEEE 802.16 WiMAX. In the medium-term, we have detailed plans for the deployment of a much larger (4 km^2) network, which will provide access along several roads in the city of Cambridge, UK. The area has been carefully selected in order that there are a variety of environments, including colleges, residential areas, university departments, a major access road, cycle paths, and minor roads. We also hope to deploy equipment along part of a nearby motorway to further increase the utility of our testbed. This network will enable us to continue to evaluate the underlying technology, and also to build some of the applications that we have described.

Since the spring of 2005 we have been developing a sensor platform for sentient transportation research, as part of the Sentient Vehicles project [4], [7]. This has allowed us to collect over 2 million data points of GPS-indexed data, including signal strength for cellular and WiFi networks in and around the city of Cambridge, as well as meteorological and pollution data, camera images, and engine performance data. Using this platform, we will be able to accurately log all pertinent information as we evaluate our network deployment.

VIII. Conclusion

In this paper, we have outlined our vision for how a near-ubiquitous wireless network for vehicular access will be used, where isolated Connected Cars then become Connected Traffic. The provision of wireless networks to support this wide range of uses is challenging, and requires both experimental characterisation of network deployments in urban environments and careful coverage planning. Techniques to easily allow the development of applications to run on the network are also key to generating value from it, necessitating research into compilers and query languages for such a large network of mobile processing agents. Another crucial consideration in application design is the admission that such a network will not be completely ubiquitous, thus the facility for applications to be aware of good estimates of future network conditions is also an important area of research. Overall, we conclude that wireless access for the road network is a challenging area, but one that holds promise to deliver a raft of useful applications.

Acknowledgments

The authors gratefully acknowledge the vision for, and support of, our work in this field from Andy Hopper. Thanks are also due to Alastair Beresford for his useful comments and suggestions.

References

[1] A. Baig, M. Hassan, and L. Libman. Prediction-based recovery from link outages in on-board mobile communication networks. In *Proc. IEEE GLOBECOM*, November-December 2004.

[2] Alastair R. Beresford, Jonathan J. Davies, and Robert K. Harle. Privacy-sensitive congestion charging. In *14th International Workshop on Security Protocols (to appear)*, LNCS, March 2006.

[3] Alastair R. Beresford and Frank Stajano. Location privacy in pervasive computing. *IEEE Pervasive Computing*, 2(1):46–55, March 2003.

[4] David N. Cottingham, Jonathan J. Davies, and Brian D. Jones. A research platform for sentient transport. *IEEE Pervasive Computing*, 5(4):63–64, Oct–Dec 2006.

[5] David N. Cottingham, Ian J. Wassell, and Robert K. Harle. Performance of IEEE 802.11a in vehicular contexts. In *Proc. IEEE VTC Spring*, April 2007. In press.

[6] Jonathan J. Davies, Alastair R. Beresford, and Andy Hopper. Scalable, distributed, real-time map generation. *IEEE Pervasive Computing*, 5(4):47–54, Oct–Dec 2006.

[7] Jonathan J. Davies, David N. Cottingham, and Brian D. Jones. A sensor platform for sentient transportation research. In *Proc. 1st European Conference on Smart Sensing and Context*, volume LNCS 4272, pages 226–229, October 2006.

[8] Peter Day, Jianping Wu, and Neil Poulton. Beyond real time. *ITS International*, 12(6):55–56, November/December 2006.

[9] Florian Doetzer, Florian Kohlmayer, Timo Kosch, and Markus Strassberger. Secure communication for intersection assistance. In *Proc. International Workshop on Intelligent Transportation*, March 2005.

[10] R. Gass, J. Scott, and C. Diot. Measurements of in-motion 802.11 networking. In *Proc. IEEE WMCSA*, 2006.

[11] Bret Hull, Vladimir Bychkovsky, Yang Zhang, Kevin Chen, Michel Goraczko, Allen K. Miu, Eugene Shih, Hari Balakrishnan, and Samuel Madden. CarTel: a distributed mobile sensor computing system. In *4th ACM SenSys*, November 2006.

[12] D. Johnson, C. Perkins, and J. Arkko. Mobility support in IPv6. Technical Report RFC 3775, IETF, June 2004.

[13] Jinyang Li, John Jannotti, Douglas S. J. De Couto, David R. Karger, and Robert Morris. A scalable location service for geographic ad hoc routing. In *Proc. ACM MobiCom*, pages 120–130, August 2000.

[14] G. Liu, B. Lee, B.Seet, C. Foh, K. Wong, and K. Lee. A routing strategy for metropolis vehicular communications. In *Proc. International Conference on Information Networking*, volume 2, pages 533–542, February 2004.

[15] Michael Mitzenmachert. Digital fountains: A survey and look forward. In *Proc. Information Theory Workshop*, 2004.

[16] Jörg Ott and Dirk Kutscher. Drive-thru Internet: IEEE 802.11b for "automobile" users. In *Proceedings of IEEE INFOCOM*, March 2004.

[17] Markos Papageorgiou and Apostolos Kotsialos. Freeway ramp metering: An overview. *IEEE Transactions on Intelligent Transportation Systems*, 3(4):271–281, December 2002.

[18] A. Seth, D. Kroeker, M. Zaharia, S. Guo, and S. Keshav. Low-cost communication for rural internet kiosks using mechanical backhaul. In *MobiCom '06: Proceedings of the 12th Annual International Conference on Mobile Computing and Networking*, pages 334–345, New York, NY, USA, September 2006. ACM Press.

[19] Jatinder Pal Singh, Nicholas Bambos, Bhaskar Srinivasan, and Detlef Clawin. Wireless LAN performance under varied stress conditions in vehicular traffic scenarios. In *Proc. IEEE VTC*, volume 2, Fall 2002.

[20] Paul F. Tsuchiya. The Landmark hierarchy: A new hierarchy for routing in very large networks. In *Proc. ACM SIGCOMM*, pages 35–42, August 1988.

[21] Upkar Varshney. Vehicular mobile commerce. *IEEE Computer*, 37(12):116–118, December 2004.

[22] H. Wu, R. Fujimoto, R. Guensler, and M. Hunter. MDDV: a mobility-centric data dissemination algorithm for vehicular networks. In *Proc. ACM VANET*, pages 47–56, 2004.