Active Highways (Position Paper)

(Invited Paper)

Liviu Iftode, Stephen Smaldone
Dept. of Computer Science
Rutgers University, Piscataway, NJ
{iftode, smaldone}@cs.rutgers.edu

Mario Gerla
Dept. of Computer Science
UCLA, Los Angeles, CA
gerla@cs.ucla.edu

James Misener
California PATH
UC Berkeley, Richmond, CA
misener@path.berkeley.edu

***Abstract*—Highways are an essential component of our society because they are critical to quality of life and to local and national economies. Under good conditions, highways provide a safe and efficient route for people and goods to reach their destinations. However, as a direct consequence of their use, traffic congestion is ever-increasing, undermining the ability of highways to adequately provide an acceptable quality of service. It has become imperative for highway traffic to provide the same time guarantee quality as other transportation methods such as air and rail travel, while maintaining the convenience of flexible scheduling and destination for the individual traveler.**

In this position paper, we propose *Active Highways*, a fundamental departure from today's highway traffic management approaches that shifts the highway paradigm from a transportation infrastructure that monitors and controls traffic at the *aggregate level*, to a computer-based service that operates at the level of individual vehicles. In this sense, highways will become *active* managers of their own traffic similar to air traffic control. In our vision, future highways and future vehicles will communicate with one another, making the highway system aware of the drivers' travel plans and allowing it to cooperate with and actively instruct the driver on achieving them. In particular, *Active Highways* will allow drivers to reserve slots in special high-priority intelligent lanes. This fine-grained traffic management model will guarantee travel time bounds, handle exceptions and enforce global community and environmental policies using real-time information from vehicle- and infrastructure-based sensors.

I. Introduction

Highways are an essential component of our society. They are critical to quality of life and to local and national economies. At the same time, by absorbing the latest technological advances in computing and networking, highways are undergoing a transformation to a large system of systems, whose control and management are becoming orders of magnitude more complex. Clearly, there is both demand and opportunity for new solutions, which exploit available computing resources and whose design is inspired by a computational thinking. In this position paper, we argue for an inter-disciplinary approach to highway traffic management that combines transportation, computational, and social perspectives.

Highways as a Transportation Problem. Under good conditions, highways provide a safe and efficient route for people and goods to reach their destinations. However, as a direct consequence of their use, traffic congestion is ever-increasing, undermining the ability of highways to adequately provide an acceptable quality of service. Numerous highway transportation and highway management solutions have been proposed and developed, ranging from increased highway capacity, to deploying dedicated sensors to obtain congestion reports, to the recent trend towards demand management methods, such as Congestion Pricing. In all of these solutions, highway traffic is managed at the aggregate level with statistical guarantees, which cannot explicitly address individual traveler destinations and priorities. To this end, today's traffic management can be seen as solving *the common case* rather than *the worst case* scenarios, potentially leaving many travelers dissatisfied. Clearly, it has become imperative for highway traffic to provide the same time guarantee quality as other transportation methods, such as air, train, and subway travel, while maintaining the convenience of flexible scheduling and destination for the individual traveler.

Highways as a Computational Problem. From a computational perspective, highways are complex control systems that execute traffic models using data obtained from multiple sources starting from censuses (useful to capture commuter traffic), continuing with sensors deployed in the infrastructure, and ending with probing executed by individual cars. Shifting from coarse- to fine-grained traffic management requires new traffic models, whose development and verification are challenging data-to-knowledge computational tasks that rely on massive data sets. Finally, traffic management is a critical distributed application, which must be reliably and securely executed. To this end, the supporting computational framework must be designed highly available and scalable to resist faults, attacks [2], and temporary disconnections.

Highways as a Social Problem. Besides being a critical transportation problem, highways and their congestion affect the economy (e.g., additional delay in the shipment of goods), quality of life (e.g., personal time wasted due to extended commuting time), and local and global environments (e.g., wasted energy resources, global warming, etc.). In our project, we are exploring how stakeholders – users, designers, planners and social scientists – can benefit from our models, for example, by predicting the impact of pricing policies on driving patterns and eventually social behaviors. Finally, assigning social responsibility to highways is a major departure from the traditional highway concept and poses difficult challenges. Linking highway operation and a scheme of automated enforcement of global or local community policies such as greenhouse gas emission regulation can be seen as a compensatory measure on travel. However, most drivers are afraid that the system may abuse, or in some way misuse the information. Privacy is a “social concern” that is one of the multi-disciplinary transportation-computational-social dimensions of our project.

II.Basic Approach

Our approach, called *Active Highways*, represents a fundamental departure from today's highway traffic management approaches by transforming the highway paradigm from a transportation infrastructure that monitors and controls traffic at the *aggregate level*, to a computer-based service that operates at the level of individual vehicles. In this sense, highways will become *active* managers of their own traffic similar to air traffic control.

In our vision, future highways and future vehicles (still controlled by their drivers), will communicate with one another, making the highway system aware of the drivers’ travel plans (destination, waypoints, stops, speed, vehicle platoon, etc.) and allowing it to cooperate with and actively instruct the driver on achieving that plan. Powerful and highly-available highway servers will continuously execute the traffic management service based on adaptive, context-aware traffic models. These fine-grained traffic management models will guarantee travel time bounds, handle exceptions (incidents, special weather conditions, etc.), while also enforcing global community and environmental policies (level of noise, fuel emissions, etc.) using real-time information from vehicle- and infrastructure-based sensors.

The transformation we propose in highway traffic management is inspired and will be enabled by a computational approach to the problem. The entire highway system can be viewed as a virtual network whose member vehicles are managed using techniques inspired from computer networks and real-time communication, such as packet scheduling, congestion control, and quality-of-service mechanisms. *Active Highways* will allow drivers to reserve slots in special high-priority intelligent lanes (which coexist with existing regular, aggregate-level managed lanes). These reservations provide individual drivers with guaranteed bounds on the arrival times to their destinations. When crashes or exceptional events occur, *Active Highways* will reroute vehicles according to their priorities and destinations, within the available capacity of the alternate routes. Moreover, by continuously sensing the environmental conditions on each segment, active highways can instruct vehicles flowing on a given highway segment in a given interval to alter their speed or their route in order to maintain the level of sensed parameters under the prescribed thresholds.

III. The Lane Reservation System

Central to Active Highways is the Lane Reservation system we proposed in [1]. The figure below shows the model consisting of three of the five key components: reservation system, lane entrance assistance system, and enforcement system. The two remaining components, lane exit assurance system and exception handling system, are not shown in the figure but are discussed later in this section. At the core of

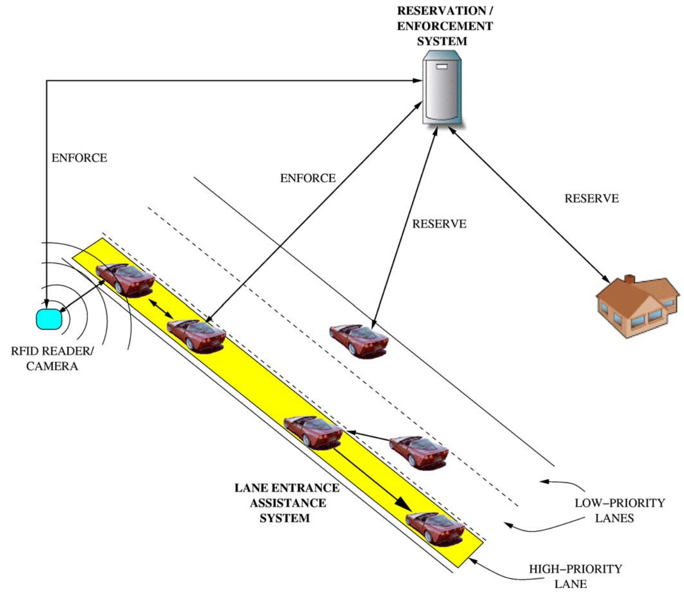

Figure 1. The Lane Reservation System.

lane reservation is the differentiation between high-priority and low-priority lanes. *High-priority lanes* are those that are under the direct management of the Active Highways system (highlighted in Figure 1). Any lane in a highway that is not under direct Active Highways management is called a *low-priority lane*. Low-priority lanes are utilized the same way highway lanes are utilized today, and are subject to all of the issues described earlier in this paper. The Active Highways system only makes guarantees to those drivers admitted to the high-priority lanes. We describe the key components of lane reservation in the remainder of this section. The *reservation* system is an online service that is responsible for issuing digitally signed tickets or tokens to drivers for traveling on a high-priority lane. Users may submit requests for a reservation ahead of time from some Internet connected location or may request an immediate reservation, on-the-go, while already in-transit. Additionally, the system allows for the identifying and priority treatment of frequent drivers. In this case, drivers would have to prove to the system, for example via location traces, that they frequently traverse a specific route (portion of or collection of highways). Given this proof, the system may favor the commuters and provide them preferential treatment in the reservation, for example, in the form of lower cost or automated scheduling.

Since we expect the relative speed of traffic in high-priority lanes to be greater than the speed in the low-priority lanes, the system must assist drivers as they transition between high- and low-priority lanes. High-priority lane entrance is managed by the *lane entrance assistance system*, which assists drivers to enter the high-priority lane according to their reservations with minimal disruption to the existing in-lane traffic.

We envision two lane entrance policies. The first policy requires drivers to possess a digital ticket, issued by the reservation system. These tickets contain a unique sequence number. Cars entering the high-priority lane issue a *request-to-yield* message (RTY) to oncoming vehicles. This message contains the entering car's sequence number and is passed back along the highway, in a vehicle-to-vehicle fashion. Once it has been passed back a sufficient number of hops, the local recipients hold an election to determine the vehicle that will yield to the entering car. Once completed, the winner of the election sends forward a *yield acknowledgment* (YA) message, which contains a unique identifier for the yielding vehicle, possibly accompanied by a physical description or image for the entering driver to recognize the yielding car. The yielding vehicle also sends a *slow down* (SD) message to cars behind it to alert them to the upcoming lane entrance event.

The second entrance policy is simpler than the first one. We consider this policy in the event that the first policy proves too complex to achieve in practice. In this policy, an entering vehicle sends an RTY message without a digital ticket. Then, the first car beyond a predetermined number of hops from the entering vehicle would be required to yield. It acknowledges the RTY with a YA and also notifies cars behind it of the upcoming event with SD messages.

A similar problem exists as vehicles transfer from high- to low-priority lanes. It is possible and likely that there will be unpredictable events in the low-priority lanes that affect the ability of vehicles to transit into the low-priority lanes. For example, consider a driver that wishes to exit at the next coming highway exit. If either the low-priority lanes or the exit lanes are congested with traffic, the system will not be able to meet its guarantee to the driver. Two possible solutions exist to handle the problem. First, it is possible to build additional infrastructure to accommodate high-priority traffic. Additional high-priority exit lanes may be built for popular highway exits to ensure capacity for high-priority traffic. Alternatively, it is possible to reserve a portion of existing exits for high-priority traffic, similar to the high-priority lanes.

To manage high-priority lane exiting, the *lane exit assurance system* maintains statistics on the exit rate for each exit and accordingly decides the upper bounds on the number of cars that should be allowed to exit at any given exit location (exit capacity). In general there will always be exceptional situations that must be handled by the Active Highways system, and this is described further at the end of this section.

The *enforcement system* serves to ensure that only drivers with a valid reservation travel on high-priority lanes. Enforcement can be handled through different mechanisms, including infrastructure-based cameras, RFID readers and other vehicle-to-vehicle communication. Roadside sensors (e.g., cameras, RFID readers, etc.) monitor the membership of high-priority lanes, and will report any violations that are encountered, to the central system for handling. Additionally, vehicles in the high-priority lanes can be equipped to challenge each other for enforcement purposes. This may be in the form of authorized vehicles (i.e., police cars) or other passenger vehicles on the high-priority lanes. The inter-vehicle validation occurs via an enforcement protocol over wireless vehicle-to-vehicle communication.

In order to maintain high quality of service with respect to travel time, the Lane Reservation system must incorporate mechanisms for handling exceptions. These exceptions include road incidents, crashes, cancellations, weather conditions and emergency evacuations. The *exception handling system* also allows reservations to be exchanged between drivers through the reservation system, either offline or on the highway in real-time. In the event of an accident on the high-priority lane, the exception handling system provides mechanisms for re-booking reservation slots.

Finally, there are both security and privacy issues in providing lane reservations. The first concern a user may have is in regards to their location privacy. By reserving slots on a highway, a user is essentially surrendering their location privacy to the system. For the system to be successful, a user must be given convincing guarantees that their private information will not be exposed. We plan to extend an existing solution [2] and apply it to address this problem. It is important to consider, though, that when comparing the system we propose to alternative forms of mass transit, such as airlines, railways, etc., it is common for users to be required to book their travel plans in advance, and to surrender their identification to the reservation/booking systems. Active Highways imposes no further privacy issues other than those already present in existing forms of commonly used travel. In fact, users may opt out of the Active Highways system and drive within the free low priority lanes, if they so choose.

Security issues, although not as critical as privacy issues, must be handled by the system. The system must handle reservation forgeries and "man-in-the-middle" attacks, especially in the case of vehicle-to-vehicle enforcement. To combat these issues we are drawing upon work in the field of secure network communication and authentication.

IV. Active Highways and Congestion

Active Highways bodes transformation in the way roads are traveled and managed. Policymakers have been grappling with this evolutionary, capital-intensive, and reactionary 'transformation' of our highway system. The consequences are excess highway capacity for much of the day and insufficient highway capacity during periods of high usage – where economic and environmental burdens are most pressing. Resources to 'build out of congestion' are precious and oftentimes nonexistent. Moreover, travel demand is out-pacing even the most ambitious highway network expansion or improvement programs.

Transportation policy has recently turned into a set of ideas that may make 'transformation' a real possibility: in the US Department of Transportation's recent *National Strategy to Reduce Congestion on America's Transportation Network*, otherwise known as the Congestion Initiative (see http://www.its.dot.gov/press/itscongestion.htm), the idea of demand management by applying spatially and temporally dynamic pricing – or tolling – to give travelers incentive to seek alternatives to traveling to congested areas during peak hours. Tolling may induce mode shift (e.g., from cars to transit), adoption of travel alternatives (e.g., telecommuting), or time shift (e.g., from the morning to midday).

Our Active Highways will multiply the potential benefits of the Congestion Initiative by channeling communication and network technology, not just to specific roadway lanes but rather to individual and connected vehicles. The transformation would be profound: individual vehicles would each comprise the edge of the Interactive Highway network, with real costs accounted and concomitant societal benefits accruing. With congestion information and highway availability instantaneously monetized, the traveler would make informed economic choices, both in pre-trip planning and during the trip. Hence, advice for route choice would essentially be in the form of price, allowing the traveler to make the price/time trade off – and under this paradigm, make a dynamic or en route selection of highways or even mode, for example, leaving a car and boarding a bus.

The system manager would have at his disposal a virtual lever, where with pricing strategies, highway sections, lanes or individual links could be filled or emptied, the mix of vehicle classifications (e.g., commercial vehicles versus light duty passenger vehicles) could be controlled, time of day for travel could be balanced or even the relative use of different travel modes could be regulated through tolling as a function of time, destination, highway, link, lane and/or vehicle class.

V. From Traffic Data to Traffic Models

The input data to be used for coarse-grained congestion pricing and fine-grained "lane reservation" pricing will be gathered from many sources. We can identify four sources that can provide suitable vehicular traffic statistics - census, historic, road sensors, and drivers.

The Census information is most useful in deriving background, long term "commuter type" and "goods delivery type" traffic patterns – the home/work place coordinates, work schedule, deliveries to markets, and highway map allow us to compute the average loads on the various sections of the urban grid [3]. The historic traffic information is obtained from past observations (coming from sensor and driver inputs) and from other statistics (e.g., average number of amusement park visitors, theater patrons, etc.). To these basic traffic levels we must superimpose the fluctuations that are caused by the dynamic interactions of vehicles at highway access ramps, traffic lights (for highways that cross urban areas), etc. Models exist to describe the statistical behavior of such fluctuations, taking into account average traffic patterns and urban grid characteristics. All of these data inputs amount to the "predicted" traffic model on a particular section of the highway. This "pre-computed" traffic model does not take into account unpredictable events such as accidents, weather, hostile activities etc. Thus, the background model must be validated and "augmented" with real time measurements from road side and road pavement sensors [4,5] as well as driver supplied measurements. In particular, we are exploring the use of smart phone with geo-location technologies (either via cell phone towers or via GPS) for highway traffic probing. In addition, the traffic model must take into account possible user reactions to the pricing schedule. High prices on one segment may cause substantial shifts to other segments - these secondary effects must be accounted for.

In this project, we plan to evaluate different techniques for the integration of background, prediction based data traffic models and real time measurements to yield an effective highway traffic model that can be interactively manipulated and used for coarse- and fine-grained pricing computation and strategy optimization.

VI. Enabling Factors

The technology enablers within Active Highways are emerging. Consider the concept Vehicle-Infrastructure Integration, or connected vehicle. FCC support in dedicated 75 MHz of licensed, free spectrum – and the ensuing IEEE 802.11p and 1609.2-4 standards are being coupled with the in-vehicle or applications standard, SAE J2739. Once these standards are in place, the potential for WiFi-like roadside equipment to essentially address, then 'talk to' and 'listen to' individual vehicles is enormous. Even less ubiquitous government investment and application of infrastructure may portend Active Highways: the prevalence of mobile telephony, many with geolocation services, and the expanded capabilities of these devices to encompass cellular, WiFi and someday, WiMax, wireless connectivity, combined with cheaper and cheaper connection costs with these communication links. Moreover, the trend for ever-more-accurate GPS devices is being coupled with advances in fusion with on-board vehicle kinematic knowledge and off-board investments in infrastructure-assisted correction services, e.g., the incipient High Accuracy-National Differential GPS system.

Finally, there are the questions of cost and incremental deployment. To solve these issues we can draw upon the success demonstrated in similar projects. High Occupancy Vehicle (HOV) lanes and EZPass are two good examples of successful highway systems that have been implemented. Both of these systems required substantial infrastructure changes and were incrementally deployed. Aside from the initial investment of stationary highway infrastructure (e.g., additional lanes/exit lanes, roadside sensors, etc.) Active Highways also requires on-board equipment. The premium price associated with reservations allows the system to be profitable, absorbing associated costs while attracting to potential investors.

VII. Related Work

This work shares a similar goal to work in providing virtual vehicle slots [8], work in providing managed motorways [9], Congestion Pricing [10], work in Personal Rapid Transit (PRT) or Guideways [11], and our previous work on Lane Reservations [1]. All these projects share the goal to address the large problem of traffic congestion with Active Highways. Our work also provides a general framework to administer and enforce other policies related to traffic, such as environmental issues (e.g., air quality, noise, etc.), energy conservation, etc. Active Highways utilizes both centralized and distributed computing techniques to solve the primary issue of traffic congestion while also enabling extended, related traffic policy enforcement.

There has been substantial work in the area of traffic sensing and monitoring. DASHExpress [12] is a system that utilizes cellular technology to collect and disseminate real-time traffic information. The VICS [13] and AHS [14] systems are infrastructure-based traffic information and driving assistance systems, while the FleetNet [15] and TrafficView [16] projects have explored the use of wireless ad-hoc inter-vehicular communication for the purposes of real-time traffic and safety information dissemination. The MIT CarTel project [17] utilizes in-car and road-side sensors as a sensor network platform to build applications such as traffic monitoring, pothole detection, and WiFi deployment mapping. Finally, the UCBerkeley PATH project [18] works to gather roadway and traffic related data to incorporate new technologies into vehicles and roadways to address various problems similar to those we describe in this paper.

Finally, utilizing vehicle-to-vehicle communication techniques, we have developed a prototype taxi booking system [19], a vehicular services [20,21], and adaptive traffic lights [22]. These projects highlight both the practicality and scalability of various applications utilizing the V2V approach. Additionally, our work in the "Experimental Vehicles, Intelligent Intersection & Instrumented Car" project [23] and the WHYNET project [24,25,26] have provided results from experimentation with real-world testbeds under various realistic vehicular conditions.

VIII. Conclusions

In this position paper, we have described our vision for the future of highway traffic management. Such large and important problems require new thinking and inter-disciplinary approaches to improve the serious conditions affecting modern highway infrastructures. We have presented Active Highways, a highway management approach that represents a fundamental departure from current approaches. In this project, we are working to provide a highway management system, analogous to those used in alternative modes of transportation (e.g., airlines, railways, etc.) to coordinate automotive highway traffic and offer guarantees to drivers in the form of travel time schedules and bounds. Additionally, the Active Highway approach enables the formulation of new traffic management policies through the gathering of important statistical data, and the enforcement of these policies within a framework that manages traffic in a fine-grained manner, down to the individual automobile.

References

[1] N. Ravi, S. Smaldone, L. Iftode, and M. Gerla. Lane Reservation for Highways (Position Paper). In the Proceedings of ITSC, 2007.

[2] F. Picconi, N. Ravi, M. Gruteser, and L. Iftode. Probabilistic Validation of Aggregated Data in Vehicular Ad-Hoc Networks. In the Proceedings of VANET, 2006.

[3] T. Nadeem, S. Dashtinezhad, C. Liao, and L. Iftode. TrafficView: Traffic Data Dissemination Using Car-to-Car Communication. ACM Mobile Comp. and Comm. Review, Vol 8, Issue 3 (July 2004), pp 6-19

[4] U. Lee, E. Magistretti, B. Zhou, M. Gerla, P. Bellavista, and A. Corradi. MobEyes: Smart Mobs for Urban Monitoring with a Vehicular Sensor Network, IEEE Comm Magazine, 2006.

[5] G. Marfia, G. Pau, E. Giordano, E. De Sena, and M. Gerla. Evaluating Vehicle Network Strategies for Downtown Portland: Opportunistic Infrastructure and Importance of Realistic Mobility Models. In the Proceedings of MoBiOpp, 2007.

[6] M. Roccetti, M. Gerla, C. E. Palazzi, S. Ferretti, and G. Pau. First Responders' Crystal Ball: How to Scry the Emergency from a Remote Vehicle. In the Proceedings of NetCri, 2007.

[7] Chen C., A. Skabardonis, and P. Varaiya. Systematic Identification of Freeway Bottlenecks. Transportation Research Record, No.1867, 2004.

[8] R. Morla. Vision of congestion-free road traffic and cooperating objects. Sentient Future Competition, 2005.

[9] V. Cahill, A. Senart, D. Schmidt, S. Weber, A. Harrington, B. Hughes, K. Singh. The Managed Motorway: Real-time Vehicle Scheduling - A Research Agenda. In the Proceedings of HotMobile, 2008.

[10] "Congestion pricing technical information summary," http://yosemite.epa.gov/aa/tcmsitei.nsf/Pages/ConPriInfoDoc.

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

[12] "DASHExpress," http://www.dash.net.

[13] "Vehicle Information and Communication System Center," http://www.vics.or.jp/english/.

[14] "Advanced cruise-assist highway system research association," http://www.ahsra.or.jp/index e.html.

[15] "FleetNet Project Homepage," http://www.informatik.uni-mannheim.de/pi4/projects/FleetNet/

[16] T. Nadeem, S. Dashtinezhad, C. Liao, and L. Iftode, "Trafficview: Traffic data dissemination using car-to-car communication," ACM Mobile Comp. and Comm. Review, vol. 8, no. 3, pp. 6–19, 2004.

[17] B.Hull, V. Bychkovsky, K. Chen, M. Goraczko, A. Miu, E. Shih, Y. Zhang, H. Balakrishnan, and S. Madden. CarTel: A Distributed Mobile Sensor Computing System. In the Proceedings of ACM SenSys, 2006.

[18] "California PATH Project Homepage," http://www-path.eecs.berkeley.edu/

[19] P. Zhou, T. Nadeem, P. Kang, C. Borcea, and L. Iftode, Ezcab: A cab booking application using short-range wireless communication. In the Proceedings of PERCOM, 2005.

[20] M. D. Dikaiakos, A. Florides, T. Nadeem, and L. Iftode. Locationaware services over vehicular ad-hoc networks using car-to-car communication. In the IEEE Journal on Selected Areas in Communications, 2007.

[21] O. Riva, T. Nadeem, C. Borcea, and L. Iftode. Context-aware migratory services in ad-hoc networks. In the IEEE Trans. on Mobile Computing, 2007.

[22] C. Gorgorin, V. Gradinescu, R. Diaconescu, V. Cristea, and L. Iftode. Adaptive traffic lights based on wireless communication between vehicles. In the Proceedings of VTC, 2007.

[23] C. E. Palazzi, S. Ferretti, M. Roccetti, G. Pau, and M. Gerla. How do you quickly choreograph inter-vehicular communications? In the Proceedings of the International Workshop on Networking Issues in Multimedia Entertainment, 2007.

[24] E. Magistretti, U. Lee, M. Gerla, B. Zhou, A. Corradi, and P. Bellavista. Mobeyes: Smart mobs for urban monitoring with a vehicular sensor network. IEEE Wireless Comm., vol. 13, no. 5, 2006.

[25] U. Lee, J.-S. Park, S. Y. Oh, and M. Gerla. Emergency related video streaming in VANETs using network coding. In the Proceedings of VANET, 2006.

[26] M. Roccetti, M. Gerla, C. E. Palazzi, S. Ferretti, and G. Pau. First responders' crystal ball: How to scry the emergency from a remote vehicle. In the Proceedings of the 1st International Workshop on Research Challenges in Next Generation Networks for First Responders and Critical Infrastructures, 2007.