Bounded-Latency Alerts in Vehicular Networks

Rahul Mangharam* Raj Rajkumar* Mark Hamilton* Priyantha Mudalige† Fan Bai†

*Dept. of Electrical & Computer Engineering
Carnegie Mellon University, U.S.A.
{rahulm, raj, mthamilt}ece.cmu.edu

†General Motors Research & Development
General Motors Corportation, U.S.A
{priyantha.mudalige, fan.bai}gm.com

***Abstract*— Vehicle-to-vehicle communication protocols may be broadly classified into in three categories; bounded-delay safety alerts, persistent traffic warnings and streaming media for telematics applications. We focus on the first category of time-critical messaging as is it of greatest value to the driver and passengers. Safety alerts are transmitted from a vehicle during events such as loss of traction, sudden braking and airbag deployment. The objective for a safety protocol is to relay messages across multiple vehicles within a 1.5-2km distance to alert approaching vehicles within a bounded end-to-end delay (e.g. 1.5 sec). Due to high mobility and ephemeral connectivity we must employ broadcast protocols, as well as mitigation strategies to curtail inherent issues associated with broadcast protocols, such as broadcast storm problem. In this paper, we present a Location Division Multiple Access (LDMA) scheme to suppress the broadcast storm problem and ensure bounded end-to-end delay across multiple hops. This scheme requires participating vehicles to time synchronize with the GPS time and receive the regional map definitions consisting of spatial cell resolutions and temporal slot schedules via an out-of-band FM/RDBS control channel. We use the GrooveNet vehicular network virtualization platform with realistic mobility, car-following and congestion models to evaluate the performance of LDMA in simulation and on the road.**

I. Introduction

Vehicles equipped with short-range IEEE 1609 Wireless Access in Vehicular Environments WAVE/DSRC-enabled wireless interfaces are poised to make driving safer, more efficient and more enjoyable [1]. As most traffic events are locally relevant, multi-hop messaging between vehicles may achieve lower latencies and utilize network resources more efficiently than cellular-based centralized schemes. Vehicle-to-vehicle (V2V) applications may be divided into three primary categories: (a) time-critical on-road safety alerts, (b) non-critical local traffic informational updates and (c) multimedia exchange as listed in Table I.

The protocol designed for these applications are quite unlikely to be end-to-end Internet protocols such as TCP/IP as they are broadcast-based and not primarily destined to a particular set of logical addresses. V2V messages are broadcast to nodes based on the vehicles' physical properties such as position, direction of travel, speed and communication capability as shown in Fig. 1. V2V messages may be routed based on fixed waypoints rather than across a set of mobile nodes. Furthermore, unlike TCP/IP, upon reaching the destination the traffic alert and warning messages are not terminated but may persist and be rebroadcast in the relevant geographic region until the traffic incident is resolved.

TABLE I

Protocol Requirements for VANET Applications

VANET Application	Primary Goal	Primary Problem	Solution Approach
Safety Alerts	Bounded Latency	Broadcast Storm	Scheduled Flooding
Traffic Warnings	Message Persistence	Disconnected Network	Adaptive Broadcast
Telematics (Streaming)	End-to-End Connectivity	Rapid Topology Changes	Heterogeneous Network

The primary goal for a V2V safety protocol is to deliver the alert to all approaching vehicles in the 1.5km to 2km range from the incident so that drivers may be alerted prior to their natural visual reaction. This imposes a tight end-to-end delay requirement on all messages and is on the order of 1.5 to 2 seconds. Due to the high mobility and ephemeral nature of link connections, V2V protocols are broadcast-based and hence suffer from the broadcast storm problem [3] with a large variance in end-to-end delay. In this paper we study a deterministic solution to this problem in the context of vehicular ad hoc networks (VANET). We achieve this by synchronizing all vehicles and scheduling grid-based spatial regions in a pipelined manner. As each vehicle is equipped with a Global Positioning System (GPS) receiver, we are able to synchronize the vehicles using the Pulse Per Second (PPS) signal from the GPS receiver. A vehicle may transmit in the current time slot if its current region is scheduled to be active during that time slot. Spatial slot resolution and temporal schedule assignments are made via an out-of-band signaling scheme. As all modern vehicles are equipped with an FM-based Radio Data Broadcast System (RDBS) receiver [2], we are able to centrally and economically schedule large areas through the regional radio station. By using this Location Division Multiple Access (LDMA) scheme we are able to reduce the number of conflicting transmitters and pipeline message broadcasts with bounded end-to-end delay.

A. Driver Reaction Times

A driver's reaction to an on-road event may be categorized into three components: mental processing, body movement and device response. On average, reaction times are 1.25sec including 0.3sec in body movement time [4]. Drivers tend to

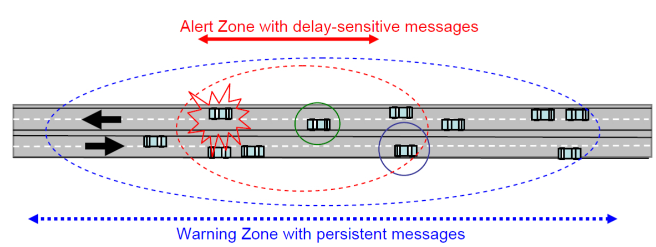

Fig. 1. Message Alert and Warning zones for vehicular networks

respond more slowly when there is high cognitive load, either from driving (complex roadway) or non-driving (use of in-car displays and cellular phones) factors. According to the California Department of Motor Vehicles [5], the reaction time of such drivers is approximately 2.5 seconds.

In order to see the importance of the mental processing component of the driver's reaction, consider for example a person is driving a car at 55 mph (80.67 feet/sec) during the day on a dry, level road. He sees a pedestrian and applies the brakes. Total stopping distance consists of three components:

1) Reaction Distance - For a reaction time of 1.5 seconds, the car will travel 1.5 x 80.67 or 120.9 feet before the brakes are even applied.
2) Brake Engagement Distance - The time for body movement till the moment the foot touches the brake pedal is 0.3 seconds on average. This accounts for another 24.2 feet.
3) Physical Force Distance. Once the brakes engage, the stopping distance (determined by a deceleration of $1/2g$ and an initial speed of 55mph) is 134.4 feet.

Total Stopping Distance = 120.9 ft + 24.2 ft + 134.4 ft = 279.5 ft. We notice that almost half the distance is created by driver reaction time. A late warning causes drivers to overreact and brake extra-hard which is a primary cause of automobile pile-ups. Human safety studies show that auditory signals generally produce faster response times and computer-based early-warning systems can significantly reduce the reaction time. Moreover, auditory transduction is mechanical, whereas visual transduction requires a relatively slow, biochemical process. For drivers with cognitive load, the reaction time distance and deceleration distance (based on CA DMV data) are shown in Table II. The goal of V2V Safety Alerts is to deliver a message with a bounded end-to-end latency of 1 and 2 seconds to approaching drivers within 1km and 2km respectively from the source. We assume a well-connected network and our goal is to ensure rebroadcasts by forwarders are mainly collision-free without the need of prior message handshaking.

B. Overview of LDMA

Probabilistic schemes for re-broadcast based flooding result in redundancy, contention and collision, ultimately resulting in long and non-deterministic end-to-end delay. We approach the desired level of determinism by the use of global hardware-based time synchronization with position information. With the use of GPS PPS signal, we are able to achieve a pairwise synchronization accuracy of sub-2ms. Once the vehicles are time synchronized with a common reference, we don't schedule individual nodes but spatial cells across the street map. This allows us to bypass the NP-hard graph coloring problem of independent set slot assignment of nodes which is not a feasible solution due to the high mobility and rapidly changing topology [6].

TABLE II
TYPICAL HUMAN REACTION DISTANCES

Initial Speed	Reaction Distance	Deceleration Distance	Total Stopping Distance
35 mph	128.3 feet	81.7 feet	210 feet
45 mph	165.0 feet	135.2 feet	300 feet
55 mph	201.6 feet	201.9 feet	403 feet
65 mph	238.3 feet	282.1 feet	520 feet

By scheduling spatial cells in a pipelined manner in time, a vehicle in an active cell is allowed to broadcast a message. Due to the broadcast nature of the shared wireless channel, cells adjacent to an active cell are inactive and the nodes are in receive mode. In the next time slot, the active cell is at a different spatial location. By separating active cells by a distance greater than the WAVE/DSRC interference range, we are able to minimize the number of collisions due to concurrent transmissions. The use of tightly-coupled global time synchronization to schedule time slots across spatial cells forms the basis of Location Division Multiple Access (LDMA). Consider a simple one-dimension road example as in Fig. 2. We observe the left-most vehicle is involved in an accident and has triggered a Safety Alert message for all approaching vehicles. In order to pipeline the message (via local rebroadcasts) we choose a cell size based on the communication and interference range of the transceiver and assign slot schedules to minimize the delay of eastward-bound messages. Assuming the DSRC radio's interference range is 300m, we assign a three-slot schedule 0, 1, 2 to all 100m-long cells along the highway. This allows for interference-free concurrent communication as all active cells are separated by the maximum interference range.

LDMA includes the facility to re-program slot schedules and spatial cell resolutions via an out-of-band control channel. This capability allows us to adapt the scheduling scheme for different traffic densities, street topologies and traffic incidents where messages are needed to proceed fast in certain directions or be persistently re-broadcast for the duration of the event. Due to the use of a low-cost, low-rate control channel such as FM/RDBS, we expect slot and cell reprogramming to occur on longer time scales of the order of a few seconds to a few hours.

Finally, LDMA integrates location-based routing with a list of waypoints, specified in terms of GPS coordinates or cell identifiers. LDMA has been implemented in the GrooveNet vehicle network virtualization platform [7] for both simulation and on-road studies. The three components including fine-grained scheduling, low-rate re-programmability and integrated location-based routing enable LDMA to exploit cross-layer optimizations across a large range of vehicle densities and network topologies.

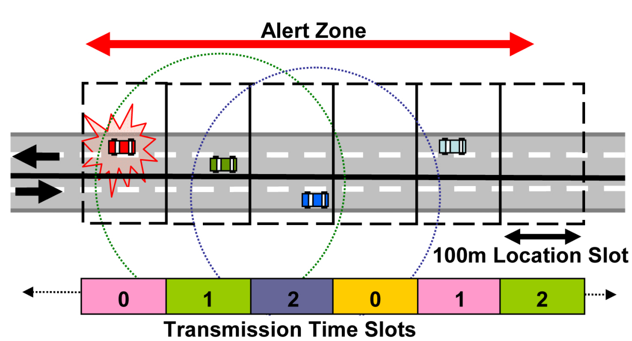

Fig. 2. LDMA for bounded message latency in the Alert Zone

We describe the LDMA protocol in Section III. We describe our realistic simulation environment in Section IV and our on-road testbed in Section V. Section VI presents the relative performance of LDMA followed by the conclusion.

II. Related Work

Link layer medium access and routing have been active areas of research in the ad hoc networking community. A primary problem in highly-mobile vehicular networks is the Broadcast Storm problem [3]. Uncoordinated broadcast-based communication suffers from poor performance in end-to-end throughput, latency and coverage because broadcasts are unreliable (messages are not explicitly acknowledged) and re-broadcasts are highly-correlated in time and space. This results in high link utilization due to contention and collisions and cannot provide a tight bound on the end-to-end latency. [3], [8] and several follow-on papers propose adaptive broadcast protocols with a variety of probabilistic schemes where the rebroadcast rate is based on node location, node degree, relative distances between nodes, etc. There are three primary problems with several of the probabilistic schemes: (a) It is difficult to select the best operating point (i.e. rebroadcast probability) without relative and neighborhood information; (b) The trade-off between latency and link utilization is non-linear, where for an incremental reduction in latency, there is a corresponding larger increase in link utilization; (c) The bounds on the end-to-end latency are very loose. Bar-Yehuda et al show that the time required for a contention-free communication across multiple hops by uncoordinated procedures is exponential in the time required by randomized ones [12]. These reasons make probabilistic message broadcast suppression schemes less attractive for time-critical Safety Alerts.

While several ad hoc routing protocols have been proposed [9], it has been shown in [10] that end-to-end connected protocols such as DSR [9] perform poorly in the vehicular networking context. The increased mobility rates with high relative speeds of 300kmph and large number of nodes over 5,000 vehicles/km^2 cause ad hoc protocols to suffer high overhead and deliver low throughput. For networks where every node's position information is available there have been several proposals for location-based and position-based protocols [11]. While our focus is on the link layer scheduling, we adopt a way-point based routing similar to the grid routing scheme mentioned in [11].

While the use of temporal and spatial scheduling is not new, LDMA provides an initial description of a practical protocol for bounded latency in the context of vehicular networks. We compare the performance of LDMA to the traditional probabilistic scheduling schemes which have a linear trade-off between reliability and end-to-end delay.

III. Location Division Multiple Access

We now describe the temporal and spatial scheduling employed by LDMA and the use of the control channel to reprogram the temporal and spatial parameters for topologically-customized schedules.

A. Spatial Regions and Cells and Temporal Slots

The task of overlaying a map with a scheduled grid may be decomposed into two sub-problems: one of efficiently dividing the street topology into spatial *active regions* followed by scheduled-slot assignment to individual active regions. We describe an active region by the region coordinates, resolution (spatial slot size) and the slot schedule as shown in Fig. 4. We assign active region resolution in a hierarchal manner from the smallest square region, *Level-0* with a side length of 50m up to the largest assigned region, *Level-11*, with a side length of 102.400km. Each lower-level region is a quarter of the area of the square region a level above. This way we can efficiently represent active regions by a tree with four siblings at each level and represent cell sizes of 50m, 100m, 200m, 400m and so on, to county-sized regions of 100km+. We define the smallest active region size to be a *Cell*.

By encoding regions with different levels, we are able to assign multiple smaller active regions (higher spatial resolution) for dense urban and suburban areas as well as sparser regions for less dense rural areas. Thus urban active regions, with higher expected vehicular density, are also assigned a higher grid resolution with smaller grid dimensions so that vehicles in rush hour may be scheduled to transmit in a tightly packed pipeline. In our simulation, we are also able to assign arbitrarily large grids with user-defined resolutions that my vary based on the time of the day or on a slow feedback of the current traffic conditions.

Once a slot schedule is assigned to an active region, we are able to pipeline the data as the slots are colored based on a k-hop coloring rule so that concurrent transmission are separate by at least their interference range (i.e. k-hops). As shown in the Pittsburgh city map in Fig. 3, the map is overlaid with a grid of 100m Cells in GrooveNet and assigned a schedule of $\{A, C, B\}$where $\{A, B, C\}$ correspond to slot numbers 0, 1 and

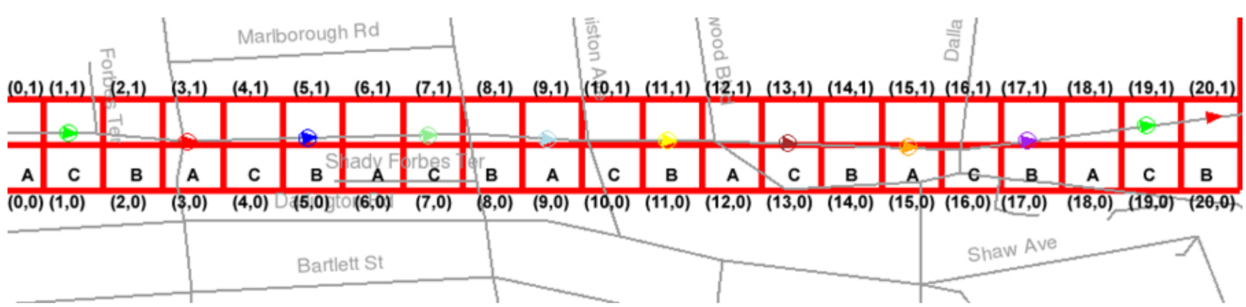

Fig. 3. Pipelined LDMA with 100m spatial Cell resolution and 3-slot temporal schedule. The schedule provides a lower delay in the eastward direction.

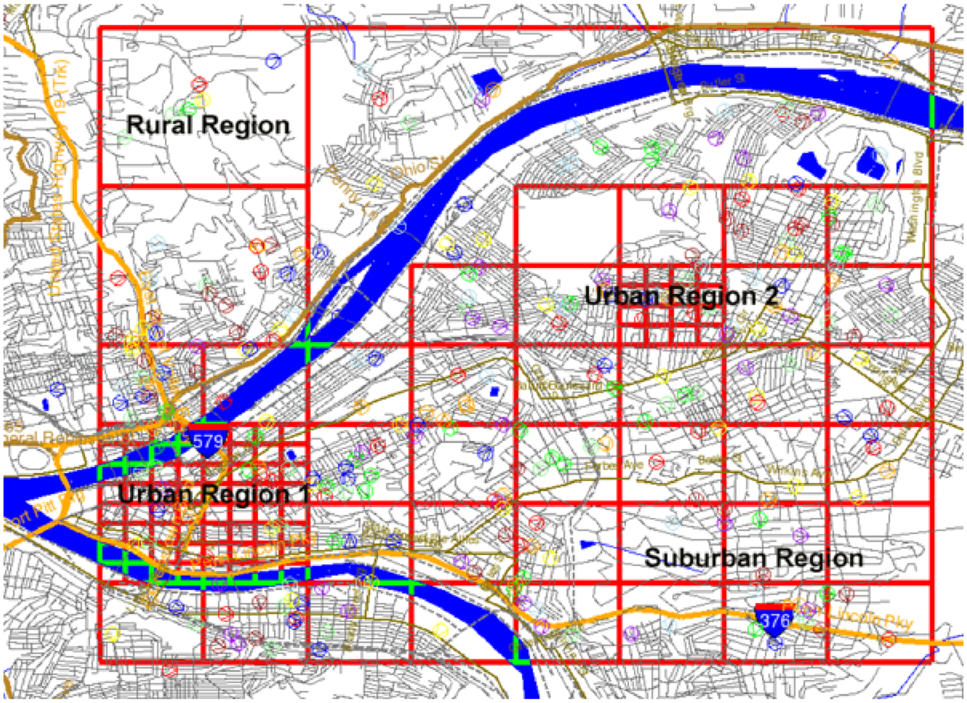

Fig. 4. LDMA active regions with different grid resolutions and schedules across Allegheny County in Pennsylvania, USA

2 respectively. The traffic is moving westward and the DSRC communication range is assumed to be 300m. The west-most vehicle is experiencing an accident in Cell (1, 1) and transmits the alert message in slot C. The alert is forwarded to all approaching vehicles in the subsequent slot with a cumulative end-to-end delay of one-half slot duration per-hop. In this example, if slot C is the current active slot, the message broadcast from (1, 1) is received by Cells (2,1) and (3, 1). In the next time slot, slot A is active and the message is forwarded eastward. Thus we are able to send the message twice as fast with the $\{A, C, B\}$ schedule than we would have with $\{A, B, C\}$. We also notice that westward-bound messages follow the $\{A, B, C\}$ schedule and travel at half the speed of eastward-bound messages. This is an example of the use of slot schedule assignment with a preference in a particular direction. For 2-dimensions, we have devised several schemes to schedule large grids in a distributed manner in [13]. While multiple vehicles may reside in a particular spatial slot, we ensure that only a single vehicle forwards the message as there is a small jitter duration at the beginning of each slot. This way, if an active transmission is overheard in the current active slot (in terms of a raised noise floor), all vehicles remain silent during the time slot.

B. LDMA Schedule Re-programmability

In order to adapt the slot schedule and spatial-cell resolution to changing vehicle densities, traffic patterns or in response to traffic incidences, we employ an out-of-band control channel. Furthermore, we can assign different schedules and cell resolutions to different areas as shown in Fig. 4. Urban and downtown areas are assigned higher resolution cells while suburban and rural areas are assigned lower resolution cells. While smaller cells (e.g. 50m) reduce the probability of collisions due to concurrent transmissions in the same cell, they result in a longer schedule. For a target end-to-end delay, it is therefore necessary to balance the size of the slots and the number of slots for a given vehicular density.

We use the FM/RDBS data channel via a regional radio station to communicate the different boundary coordinates of each active region, their associated slot resolution, schedule and time of activity. For example, during the morning rush hour, there is significantly more traffic from the suburbs to the downtown in most US cities. We can therefore specify schedules with lower latency in the direction of the suburbs along major highways entering the city. The schedule will be active during the morning rush hours as there is a greater chance of accidents along more congested roads. Similarly, schedules with lower delay in the opposite direction can be activated in the evenings. As most vehicles are equipped with an FM/RDBS receiver, they can receive the scheduled LDMA parameters and operate with the currently active schedule when they are in the range of the regional radio station. Fig. 5 illustrates a typical FM radio channel with the left and right audio channels for stereo sound at the lower frequencies. The RDBS signal nests into the 57 kHz position between the stereo multiplex and the 67 and 92 kHz sub-carrier channels. Through the use of the Open Data Channel (ODC) [2], we can define our own application-specific data stream for the LDMA protocol. The RDBS sub-carrier offers a 1 Kbps raw data stream which is used to periodically communicate the LDMA

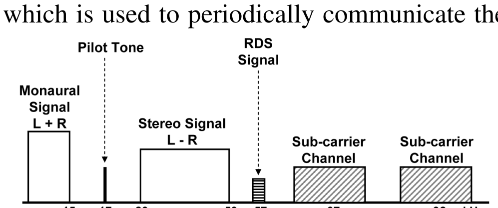

Fig. 5. Radio Data Broadcast System sub-carrier in the FM radio band

configuration information for the different active regions. For example, the local Classical radio channel (WQED) broadcasts a 20kW signal across a 40-mile radius around Pittsburgh and can send active region updates once every 10 seconds. With frequent updates, the active regions can be re-configured to suit traffic patterns at different times of the day based on historical data and also react to feedback from current traffic incidences and congestion.

We employ three slot scheduling strategies: pipelined lines for highways, 2D-grids for sub-urban areas and radials for urban intersections. These scheduling schemes and their associated end-to-end delay properties are described in detail in the associated technical report [14].

IV. Realistic VANET Simulation

We use the GrooveNet network virtualization platform for both simulation and on-road evaluation of LDMA. All vehicles travel along a street map topology and realistic mobility, trip and communication models. We have implemented a car-following, traffic light and several adaptive broadcast communication models in GrooveNet. In order to correctly represent vehicle interaction, GrooveNet includes simple car-following, traffic lights, lane changing and simulated GPS models. Three types of simulated nodes are supported: vehicles which are capable of multi-hopping data over one or more DSRC channels, fixed infrastructure nodes and mobile gateways capable of vehicle-to-vehicle and vehicle-to-infrastructure communication. GrooveNet supports multiple message types such as GPS messages, which may be broadcast periodically to inform neighbors of a vehicle's current position, and vehicle emergency and warning-event messages with priorities.

GrooveNet supports multiple network interfaces for real vehicle-to-vehicle and vehicle-to-infrastructure communication such as: a 5.9GHz DSRC interface, IEEE 802.11a/b/g, 1xRTT and EVDO cellular interfaces. Communication may be established over TCP or UDP sockets. All real vehicles communicate with DSRC or 802.11 with each other and in addition, mobile gateways communicate with infrastructure nodes over the cellular interface. GrooveNet is able to support hybrid (i.e. communication between simulated vehicles and real vehicles on the road) simulations where simulated vehicle position, direction and messages are broadcast over the cellular interface from one or more infrastructure nodes. Real vehicles communicate with only those simulated vehicles which are within its transmission range. GrooveNet is able to connect to the vehicle's on-board computer and read OBD-II diagnostic codes which can trigger alert or warning messages.

We implemented LDMA in the simulator such that the same model can be used both in simulation-mode and also on-road with real vehicles. We are currently able to schedule arbitrary-sized 1-D and 2-D grids with a variety of custom schedules. Each LDMA slot is 10ms long to allow for a 6.5ms maximum-sized IEEE 802.11p message (2312 octets) at a minimum rate of 3Mbps [1]. We currently use a 3ms guard time between slots. Each slot has a 500μs initial back-off window to suppress concurrent transmission in the same cell. Our implementation currently permits at most one packet transmission in a time slot, regardless of the packet size.

TABLE III

GPS/PPS Time Synchronization Accuracy (time in ms)

Time Server	IP address	Relay	Offset	Jitter
+time1.apple.com	17.82.254.14	98.624	1360.95	833.978
+ac-ntp0.net.cmu.edu	128.2.1.20	20.720	1.124	1.438
127.127.1.0	.LOCL.	0.000	0.000	0.001
*192.168.0.04	.GPS.	1.942	0.890	0.171

V. LDMA System Implementation

We are able to extract the GPS PPS signal to achieve sub-200μs local time synchronization accuracy and a spatial accuracy of 3m for a vehicle moving at 45m/s. We use the Linux 2.6 based Gumstix single-board computers [15] to obtain local time synchronization via the GPS/PPS signal. The hardware platform is able to do dead-reckoning in the downtown areas by feeding the GPSstix add-on module with the vehicle's odometer, direction and steering CAN-bus outputs. For stationary vehicles, we are able to obtain pairwise synchronization with sub-2ms accuracy. We are currently in the process of deploying LDMA across several real vehicles and are yet to determine the maximum jitter across a fleet of fast moving vehicles. Table III shows the time offset and jitter of our GPS receiver compared to two stratum-2 Network Time Protocol (NTP) servers. We observe the remote server (i.e. time1.apple.com) has a larger relay time and a jitter that is two orders of magnitude larger than the local GPS time. A physically closer NTP server (ac-ntp0.net.cmu.edu) has a jitter on the order of 1-2ms but like the other NTP server it is not accessible on the road. The PPS signal offers a small relay time and a very low jitter of 171μs.

In order to incorporate realistic link and physical layer models in GrooveNet, we drove a pair of vehicles over 40 miles along a national highway. We logged the SNR, packet error rate, speed, position and number of visible GPS satellites. In Fig. 6, we recorded the receive SNR from two vehicles at a third vehicle driving in-between. While the SNR was quite stable along the highway, the variation increased near traffic intersections and in urban areas. We are able to playback and incorporate this data in the physical layer model to get more realistic results of packet error rate and message delays across multiple hops. This test helped us determine our LDMA cell sizes to 100m.

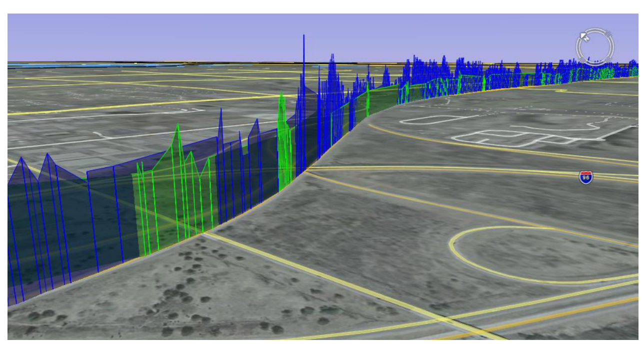

Fig. 6. Variation of SNR of two vehicles on a test drive along I-96 in Michigan, U.S.A.

VI. LDMA Performance

In order to evaluate the the end-to-ed latency and link utilization of LDMA, we compare it with four adaptive rebroadcast schemes proposed in [3]. We implemented probabilistic, location-based, distance-based and neighbor-based adaptive rebroadcast link layer models in GrooveNet. In order to keep the topological parameters to a minimum, we routed a fleet of vehicles along a single highway (1-dimension). The vehicles were separated such that we could achieve various vehicle densities while maintaining a connected network. When a traffic incident message is broadcast from the source vehicle, each scheme throttles the rebroadcast rate based on its position, distance from the event, number of active neighbors, etc. Our link layer employs a packet collision model where no packet is received successfully if more than one reception occurs concurrently. The LDMA scheme forwards the packet with a small jitter (e.g. 500us) when a vehicle is in an active slot. After the first re-broadcast, subsequent rebroadcasts were scheduled at 1 sec intervals. This results in a pipelined set of transmissions with a minimum end-to-end delay of 10ms per vehicle hop. As mentioned in Section III, the highest speed a message can achieve is 200m/20ms or 10,000m/s with the $\{A, C, B\}$ schedule.

The distance-based rebroadcast rate-control selects the rebroadcast rate as a function of the current vehicle's distance from the event location. The rebroadcast rate is high near the event (e.g. 6.5ms) as it is most relevant to alert drivers near the incident. The rate decreases linearly with distance to a maximum interval (e.g. 5 sec) determined by the maximum relative vehicle speed. With the position-based rebroadcast suppression a vehicle does not transmit during an interval (e.g. 50ms) if it overhears a broadcast from a vehicle further away from the event. This is relevant because the message has already propagated down the road beyond the vehicle's current location. Finally, the neighbor-based duty cycle throttling scheme increases the rebroadcast interval exponentially based on the number of rebroadcasting neighbors the vehicle overhears. We used the binary exponential back-off with a minimum window of 500us. [3] shows that additional spatial coverage is $<0.05\%$ when message is heard from >4 neighbors. The periodic scheme rebroadcast messages every 200ms.

We evaluated the above four rebroadcast schemes by observing the trade-off between message delay and the rebroadcast duty cycle function (link utilization). We choose a highway near Pittsburgh with a vehicle density of 25 vehicle/km^2. A message was broadcast from an event vehicle and the delay and message receive rate (messages/(vehicle$\times$sec)) was recorded. In Table IV, we observe that the LDMA Rebroadcast scheme has the smallest delay and a moderate (controllable) message receive rate due to the time synchronized scheduled operation. Among the adaptive schemes, the neighbor-based rebroadcast scheme provided the best trade-off between flooding the network and end-to-end message delay. For a more detailed performance analysis please refer to the LDMA technical report [14].

TABLE IV
Comparison of End-to-end Delay (in sec) and Link Utilization of VANET Rebroadcast Schemes

Distance	Periodic	LDMA	Location	Distance	Neighbor
0.5km	0.666	0.040	1.906	1.30	0.559
1km	2.03	0.100	2.76	3.58	3.26
1.5km	2.92	0.152	3.93	4.22	3.81
Rate	4.39	3.8	0.67	2.67	0.31

VII. Conclusion

In order to broadcast time-critical Safety Alerts across multiple hops in a vehicular network with bounded end-to-end delay, we presented LDMA, a Location Division Multiple Access protocol. LDMA employs tightly-coupled time synchronization to temporally and spatially schedule regions in a map. By pipelining communication we can achieve low-delay and controlled link utilization compared to adaptive broadcast link protocols. We have simulated LDMA on the GrooveNet vehicular network virtualization platform and implemented GPS/PPS-based time synchronization for real vehicles. We are able to achieve a sub-200μs local synchronization accuracy. In addition to the scheduled protocol, we employ an out-of-band control channel to specify the spatial cell resolution and temporal schedules to vehicles in different regions. With this we are able to adapt LDMA schedules based on vehicle density, traffic patterns and as a response to observed traffic incidences. Based on our current evaluation, LDMA's globally synchronized approach to achieve bounded end-to-end delay is a promising direction for time-critical VANET protocols.

References

[1] IEEE P1609, Wireless Access in Vehicular Environments (WAVE) - Architecture for Intelligent Transportation Systems.

[2] S. Wright. The Broadcaster's Guide to RDS. Focal Press, 1997.

[3] S.-Y. Ni. et al., The Broadcast Storm Problem in a Mobile Ad hoc Network. *ACM MobiCom*, 1999.

[4] M. Green. 'How Long Does it take to Stop?' Methodological Analysis of Driver Perception-Brake Times. *Transportation Human Factors*, vol. 2, no. 3, pp. 195-216. 2000.

[5] 'Driver Handbook', *California Department of Motor Vehicles.* 2007.

[6] H. Balakrishnan et al. The Distance-2 Matching Problem and its Relationship to the MAC-layer Capacity of Ad hoc Networks. *IEEE JSAC*, vol. 22, no. 6, pp. 1069–1079. 2004.

[7] R. Mangharam, et al. GrooveNet: A Hybrid Simulator for Vehicle-to-Vehicle Networks, *Second International Workshop on Vehicle-to-Vehicle Communications (IEEE V2VCOM),* San Jose, USA. July 2006.

[8] Holger Fubler, et al. Contention-Based Forwarding for Mobile Ad-Hoc Networks. *Elsevier's Ad Hoc Networks*, 1 (4), pp. 351-369, Nov. 2003.

[9] E. Royer and C.-K. Toh. A Review of Current Routing Protocols for Ad Hoc Mobile Wireless Networks. *IEEE Personal Communications Magazine*, vol. 6, no. 2, pp. 46-55, 1999.

[10] J. Blum, et al. Challenges of Inter-vehicle Ad Hoc Networks. *IEEE Trans. on Intel. Transportation Systems*, 5(4):347–351. 2004.

[11] M. Mauve, et al. A Survey of Position-based Routing in Mobile Ad hoc Networks. *IEEE Network*, 15 (6), pp. 30-39, 2001.

[12] R. Bar-Yehuda, O. Goldreich, A. Itai. On the Time-Complexity of Broadcast in Radio Networks: An Exponential Gap Between Determinism and Randomization. *PODC* pp. 98-108, 1987.

[13] R. Mangharam and R. Rajkumar, MAX: A Maximal Transmission Concurrency MAC for Wireless Networks with Regular Structure, *IEEE Broadnets*, San Jose, CA, 2006.

[14] R. Mangharam, R. Rajkumar, M. Hamilton, P. Mudalige and F. Bai. LDMA: Location Division Multiple Access for Bounded-latency Broadcasts in Vehicular Networks, *CMU Tech Report* CMU-07-02, 2007.

[15] Gumstix Linux-based single-board computers. http://www.gumstix.com