Proceedings of the IEEE ITSC 2006
2006 IEEE Intelligent Transportation Systems Conference
Toronto, Canada, September 17-20, 2006

Smoothing Vehicular Traffic Flow using Vehicular-based Ad Hoc Networking & Computing Grid (VGrid)

Andrew Chen*, Behrooz Khorashadi, Chen-Nee Chuah, Dipak Ghosal, and Michael Zhang

Abstract—**We propose VGrid: an ad hoc networking and computing grid formed by leveraging inter-vehicle wireless communications. In addition to exchanging data between vehicles, VGrid actively uses pertinent data to perform computations for solving traffic-related problems. In this paper, we examine the problem of smoothing vehicular traffic flow through the use of real-time position and velocity information exchanged over the network. This is accomplished through the application of 1) accident alert messages sent from the site of an accident or other obstruction in the road and 2) dynamically-calculated variable speed limits based on the local density of vehicles. Using a realistic model of highway traffic flow, we show through simulation that VGrid reduces speed variance, corresponding to more homogeneous vehicle behavior in free-flow and obstructed-lane scenarios.**

I. INTRODUCTION

According to a recent report [1], billions of dollars are lost annually due to traffic congestion. Therefore any traffic management measure that reduces traffic congestion by even a small proportion could translate into huge savings. Although the specific triggers of traffic congestion are numerous, one can broadly attribute them to "bottlenecks," or hot spots on the road. Two types of bottlenecks can be found: fixed bottlenecks and dynamic bottlenecks. The former can be points where two or more roads merge together, the number of lanes decreases, grades become steeper, or curves become tighter. Dynamic bottlenecks can be locations of accidents that temporarily block one or more lanes or slowly-moving vehicles in a traffic stream (such as trucks on a steep climb). Accidents are found to cause more than half of congestion hours in this country [1]. As such, measures to reduce the occurrences of traffic accidents are highly desirable from both safety- and congestion-management perspectives. One such measure, variable speed limit control [2], has been demonstrated to hold great potential in stabilizing traffic flow and reducing the number of accidents on the road.

The FCC has recently allocated the 5.85-5.95GHz portion of the spectrum for inter-vehicle communication and vehicle-to-roadside communication, known as dedicated Short Range Communication. This has prompted a recent surge in vehicular research ranging from new design applications, which include driver-vehicle safety applications, infotainment, and mobile in-vehicle Internet services for passengers [3,4]. In addition to these applications, the newly allocated spectrum also provides an opportunity to leverage vehicular ad hoc networks (VANET) to revolutionize the solutions to vehicular traffic flow control.

In this paper, we introduce a new framework called VGrid (Vehicular based Networking and Computing Grid), in which we harness DSRC-enabled vehicles to form a distributed, ad hoc computing grid. VGrid is used not only to monitor real-time vehicular traffic and road conditions, but also to exchange information among vehicles to perform distributed computation to control or smooth traffic flow. This networking/computing capability can enable vehicle-driver safety applications. For example, to reduce "stop-and-go" traffic and unnecessary deceleration or stopping, VGrid can help locate accidents and alleviate congestion through traffic metering and early warning messages.

In normal driving conditions, drivers can see one or at most two cars ahead and around themselves. As a result, a vehicle can only react to its immediate neighbors. If a vehicle stalls or crashes, vehicles far upstream may be forced to quickly decelerate or stop. Due to the "shockwave" effect, vehicles far behind the affected vehicles will slow down as well. However, using VGrid, drivers have the ability to "see" farther down the road and react to incidents early. Instead of waiting to reach a crash site or a heavily congested lane and then changing lanes, upstream vehicles can be alerted far ahead of time giving them much more time to react and switch out of the affected lane. Additionally, VGrid can be used to dynamically implement Variable Speed Limits, which can result in smoother traffic flow when accurately calculated, by leveraging inter-vehicle communication to exchange and propagate individual vehicle speeds to surrounding vehicles.

The remainder of this paper is structured as follows: First we introduce the VGrid framework and the benefits it has over fixed-infrastructure sensing and control systems such as inductive loops and road signs. Second, we introduce two methods by which VGrid can improve driver safety and smooth traffic flow: accident alert, an early warning system which gives drivers significantly greater time to react to roadway obstructions, and variable speed limit (VSL), a dynamic speed limit calculated based on traffic characteristics. Third, we discuss the mobility and network models as well as the parameters utilized in the simulations. Finally, we present the simulation results for both free-flow traffic (no obstructions) as well as roads in which one or two lanes are obstructed. The results show that VGrid reduces speed variance in all scenarios, though with reduced effectiveness at high levels of congestion. We also show that performance increases with the percentage of equipped vehicles.

---

[*] Corresponding author: Andrew Chen (apchen@ucdavis.edu)

1-4244-0094-5/06/$20.00 ©2006 IEEE

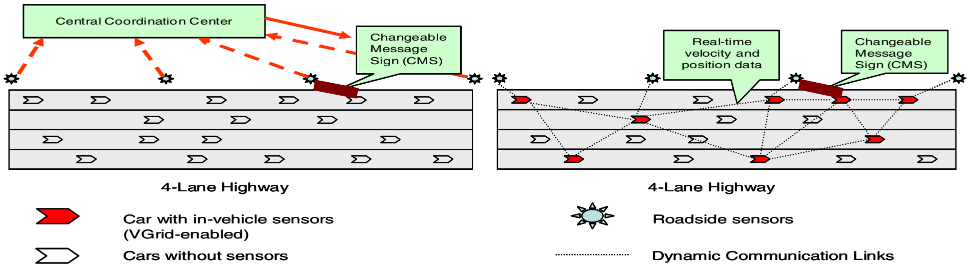

Fig 1: (Left) Fixed ITS system with central controller, roadside sensors, and changeable speed limit sign. (Right) Distributed VGrid architecture with real-time inter-vehicular and vehicle-roadside communications.

II. The VGrid Framework and Architecture

VANET, unlike mobile ad hoc networks (MANET), does not suffer the same constraints that are seen in typical MANET devices such as limited battery and computing power. Vehicles come equipped with significant electrical, and, more recently, computational resources. GPS devices are becoming faster and more accurate as well, with many manufacturers claiming sub-meter accuracy. However, mobile GPS systems at this accuracy are still cost-prohibitive. The highly mobile individual nodes that make up VANET create a system with much more dynamic topology and channel conditions than traditional wireless networks. Another important aspect of VANET is the lack of need for fixed roadside infrastructure.

In VGrid, we describe a fully-distributed traffic monitoring and control system in which individual vehicles share their mobility and sensor information with their neighbors in order to collectively alter their traffic flow patterns in real-time. Unlike traditional grid computers, in VGrid, both the topology and the node membership change with time. When the vehicles forming the ad hoc grid computer are themselves part of the vehicular traffic flow that is being controlled, the system has a characteristic of being a self-referential system. Such a system has the following unique feature: the capacity of the grid computer increases as congestion increases, which is precisely when more capacity is needed (up to a certain point) to accurately determine the vehicular flow characteristics that can smooth and homogenize the traffic flow and thereby optimize the associated delay and throughput.

Applications that utilize VGrid can operate in a number of different scopes, ranging from a single car to a platoon of cars that form a peer space, as discussed in [5]. The *platoon scope* is important since many potential VGrid applications will require knowledge of data contained in a local area of interest. However, for some applications, we may need a larger scope that combines a collection of peer spaces to achieve wider-scale results.

Compared with fixed sensor network or computing grids, VGrid offers several advantages. Because of its mobility, VGrid can be deployed wherever and whenever it is needed in a road network. Because of its ad hoc nature, information can propagate along any direction. Because of its finite information propagation speed, information is disseminated in a staggered fashion, avoiding over-flooding or the so-called herding effect in broadcast based information systems.

III. Application of VGrid to Vehicular Flow Homogenization

The goal of VGrid is to smooth traffic flow through the dissemination of traffic data through a grid network. In order to quantify this, we define a set of metrics we will use to describe traffic flow. The primary metric is the speed variance normalized to the average speed, which indicates the amount of acceleration and deceleration vehicles experience. Without VGrid, drivers lack precise speed and position information about the vehicles around them, which may lead to over- or under-reaction to changes in speed, causing accidents or unneeded slowdowns. If variance is low, traffic flows at a more constant rate and this scenario occurs less often. This homogenization also affects the throughput, the number of vehicles exiting a section of roadway in a fixed time interval, and the latency, the amount of time it takes vehicles to exit a section of roadway, are indicators of the overall performance of the system. Therefore, the goal is to minimize variance and latency and to maximize throughput. In order to accomplish this, all vehicles must periodically broadcast their own GPS location and velocity and listen for these beacons from all other vehicles.

A. *Accident Alert*

One basic application of on-board wireless communications is to alert upstream vehicles of the presence of an obstruction in the road ahead, allowing them to change out of impacted lanes earlier and preventing them from changing into those lanes. This may include accidents, construction zones, and debris. Emergency or construction vehicles as well as nearby automobiles capable of identifying such incidents will broadcast their GPS location to other cars in the area, which will relay this information upstream.

In addition to the message being sent at the location of the accident, we also take into account the queuing of vehicles behind an accident. A vehicle that is below a certain speed threshold and within a threshold distance from either the accident or another vehicle in a queue is considered to be in a queue. When such a queue is formed, the vehicle at the end may be considered to be the location of the accident, allowing drivers to avoid the effects of an accident even if a long queue has already formed.

In this paper, we consider a probabilistic lane changing model where the probability of a vehicle changing out of an impacted lane increases linearly as it approaches the point of the accident, provided there is sufficient space in adjacent lanes.

B. Computing Optimal Variable Speed Limits

Another application we consider is the distributed computation of variable speed limits. Every vehicle periodically broadcasts messages containing its position and velocity. An algorithm based on the local density of vehicles around each vehicle is used to determine the variable speed limit. The fundamental purpose of VSL is to reduce the speed of vehicles in order to limit the sudden acceleration and deceleration that leads to accidents.

For the purposes of this paper, we consider a simplified VSL algorithm where vehicles calculate individual speed limits based on a linear function once the local density around a vehicle passes fixed thresholds. That is, vehicles will move at v_{max} until the local density reaches $density_{min}$, at which point the maximum allowed speed decreases linearly with density until it reaches $density_{min}$. At this point, the speed limit is fixed at a minimum value. The actual values for these parameters are fixed and determined *a priori.*

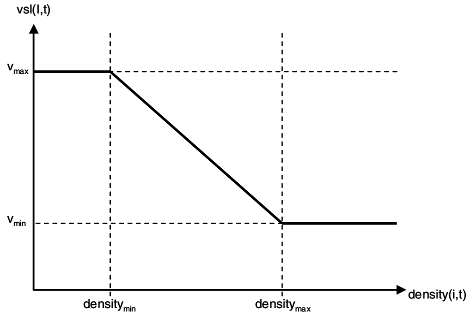

Fig. 2: Example VSL algorithm. Speed limit decreases linearly as density increases within pre-determined

IV. SIMULATION MODEL

A. Mobility Model

In the Nagel and Schreckenberg [6] (N-S model), the road is divided into cells with equal-length $\Delta x=7.5m$ containing 1 vehicle and the simulation granularity is $\Delta t=1s$. For this purpose of this study we have extended the basic N-S model to more accurately reflect real-world traffic though the work in [7] and by adding lane-changing capability.

The basic N-S model considers only free-flow highway traffic with vehicles having a fixed probability of random slowdown per time step. [7] extends this model by adding the following behaviors:

1. At large distances the cars move (apart from fluctuations) with their desired velocity $v_{\max}$.
2. At intermediate distances drivers react to velocity changes of the next vehicle downstream, i.e. to `brake lights'.
3. At small distances the drivers adjust their velocity such that safe driving is possible.
4. The acceleration is delayed for standing vehicles and directly after braking events.

In this new model, $\Delta x=1.5m$ and each vehicle occupies 5 cells. The maximum speed is 20 cells/second, corresponding to 162 km/hour.

We define a vehicle as queued if it is below a certain speed threshold and a within a certain distance from either an accident or another queued vehicle. That is:

1. Its speed is less than v_{min}
2. There is a vehicle in front that is either stalled or in a queue
3. This stalled or queued vehicle is less than d_{min} cells away

In addition, we introduce lane changing behavior into the model. At every time step, each vehicle has a chance to change into another lane with probability $p_{change,}$ where, p_{base} is a fixed lane changing probability, *max_distance* is the maximum distance considered for accident alert, and *stall_distance* is the distance to the accident. For simplicity, we consider *max_distance* to be the maximum wireless transmission range. In the cases where there is no accident or it is beyond *max_distance*, *stall_distance*= *max_distance*. Described mathematically:

$$P_{change} = p_{base} + (1-p_{base})(max\_distance - stall\_distance)/max\_distance$$

In order to change lanes, the vehicle *i* and objective lane must meet a set of conditions. First, the vehicle must have incentive to change, that is, there must be more forward room in the objective lane than the current lane and, if there is an accident in the objective lane, it must be farther away than one in the current lane. Second, there must be enough room in the objective lane behind and in front of the vehicle such that a vehicle *j* will not collide with *i*. These conditions are stated as follows:

(1) Incentive Conditions: $gap(i) < gap_o(i)$
$d_{stall} < d_{o,stall}$

(2) Security Conditions: $gap_o(i) > v(i)$
$gap_{o,back}(i) > v(j)$

To prevent vehicles from continually changing lanes, we introduce a parameter t_{change} that indicates the minimum time between lane changes.

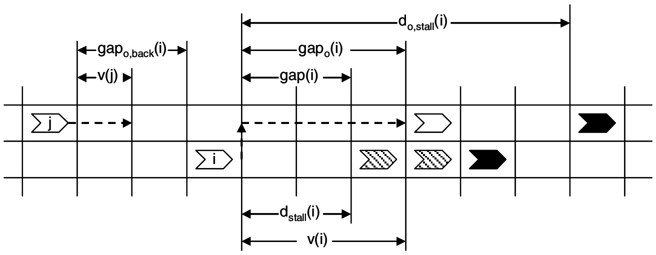

Fig. 3: Lane changing model. Black indicates stalled vehicles. Hatched indicates queued vehicles.

B. Communication Model

Due to the small amount of data versus the high amount of available bandwidth, the simulator utilizes a simplified communication model that ignores higher network functionality in exchange for simulation speed. Messages will only contain a unique ID, GPS location, and speed. Thus, we are able to make the assumption that these small messages can be broadcast in approximately 1ms. Each simulation step (one second) is therefore broken up into 1,000 time slots in which every vehicle is allowed to broadcast. The content of the messages is of interest to all nearby users, thus there is no need to include addressing or routing functionality. We also assume processing time to be negligible.

The main concerns are transmission collisions and transmission range/power. Transmission collisions are defined to occur when two or more vehicles with overlapping transmission ranges broadcast during the same time slot. All vehicles within those overlapping regions do not receive the message. The transmitting vehicles are able to detect this and will back off a random amount of time slots before attempting retransmission. Transmission range is defined as the maximum radius over which vehicles are able to transmit reliably. Phenomena such as fading and interference are not considered in this simplified model. Therefore, as long as a receiver is within range and there are no colliding signals, the receiver will receive the message error-free. For simplicity, we treat the range as a rectangle across all lanes rather than radial from the transmitter since the width of the road is much lower than the transmission distance.

C. Simulation Tool

We developed a Java-based simulation tool based on the Cellular Automaton Traffic Simulators applet developed by Kai Bolay [8] but extended with our simulation model. The mobility model parameters are: $P_b = 0.94$, $P_0 = 0.5$, $P_d = 0.1$, $h = 6$, $gap_{security} = 7$, as described in [7]. The other simulation parameters are listed in Table 1.

V. RESULTS AND DISCUSSION

We consider the effect of wireless transmission range and VGrid on traffic flow over a range of traffic intensities, where intensity is defined as the probability that a vehicle will enter the roadway on every time step. This is a lossy system in that, if a vehicle is scheduled to enter the roadway but cannot, it is dropped from the system. Transmission range is measured in cells and a range of 0 implies that the VGrid algorithm is disabled. The roadway is 1.5km (1000 cells) and vehicles enter the roadway at maximum speed. The primary metric we are concerned with is the normalized speed variance, that is, the variance of velocities of all vehicles normalized to the average velocity of all vehicles. This takes into account the fact that, at slower speeds, there will naturally be lower variance since vehicles cannot move slower than 0 cells/second. Simulations last for 5000s.

PARAMETER	DESCRIPTION	VALUE
v_{max}	Maximum speed limit	20 cells/time step
v_{min}	Minimum speed limit	5 cells/time step
density$_{min}$	Minimum density to begin VSL	0 vehicles/cell
density$_{max}$	Maximum density to stop reducing speed limit	0.5 vehicles/cell
p_{base}	Base probability of changing lanes	0.05
t_{change}	Time between lane changes per vehicle	10 time steps
Queuing Distance	Maximum distance from next vehicle below queuing speed to form a queue	20 cells
Queuing Speed	Maximum speed to be considered to be in a queue	5 cells/time step

Table 1: Simulation parameters

In the 4-lane highway scenario, we look at three accident configurations: no accidents; 1 fixed accident at lane 1, cell 800; 2 fixed accidents at lane 1, cell 800 and lane 2 cell 800. The accidents are fixed for the duration of the simulation. In the no-accident scenario, we look at the effect of VSL alone. In the one- and two-accident scenarios, we look at the effect of both VSL and Accident Alert. We also examine the effect of penetration rate of enabled vehicles. Non-participating vehicles will only be aware of their immediate neighbors.

A. Effect of Transmission Range

From the variance results shown in Figures 4, 5, and 6, we see a marked improvement with the addition of VGrid in all three scenarios. The largest jump is from 0 transmission range (VGrid disabled) to 25 cells (37.5 meters). As we increase the transmission range beyond that, there are diminishing marginal returns. That is, for accident alert, there is little benefit for knowing about an accident from very far away since it may not affect traffic at that distance. For VSL, increased transmission range also offers little benefit since too coarse of an aggregation will ignore local "hot spots" that, again, may have little effect on traffic at a great distance away.

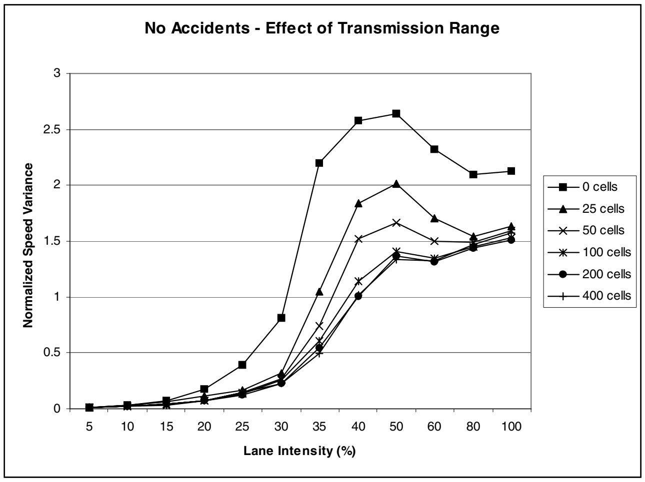

Fig. 4 - Performance of VSL only at varying transmission ranges, from 0 cells (0m) to 400 cells (600m). Intensity is the same across all lanes. Traffic flow becomes smoother (i.e. variance decreases) as transmission range increases.

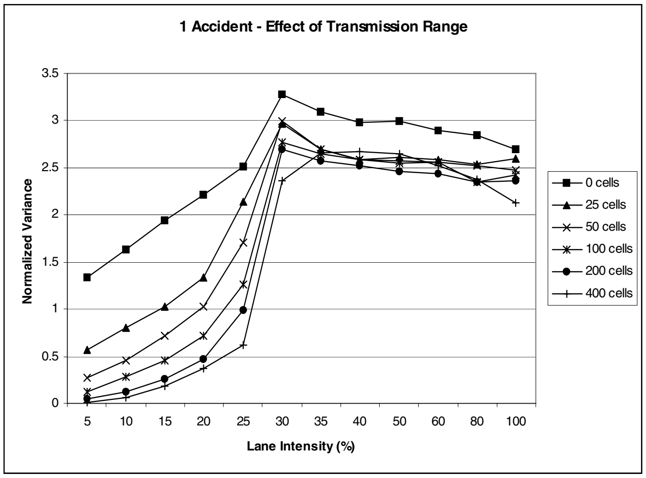

Fig. 5 – The performance of VSL and accident alert with one fixed accident.

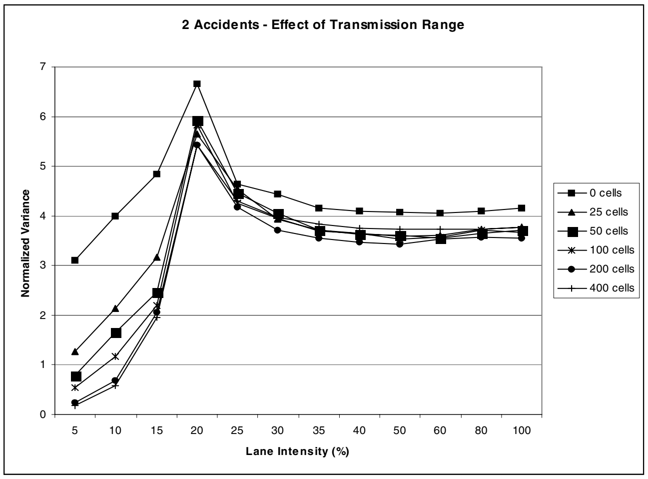

Fig. 6 – The performance of VSL and accident alert with two fixed accidents.

Due to the fact that the vehicles in the mobility model can react and change lanes instantaneously, the simulated performance may not reflect realistic driver behavior but the results still illustrate the utility of Accident Alert and VSL systems in lowering speed variation on the roadway. Real drivers will require significantly more time to process the incoming data and determine the proper course of action.

In the two accident scenarios, there is a noticeable peak in each graph, indicating the point at which the congestion level is the maximum. VGrid has little effect here due to the lack of space for vehicles to maneuver around accidents or slower vehicles. As intensity increases past this peak, we observe a decrease in normalized variance. This occurs because the effective intensity (taking into account vehicles that could not enter the roadway) begins to decrease due to the congestion at the start of the road preventing new vehicles from entering. This is most prominent in Figure 6 due to the extreme bottleneck caused by two accidents limiting the egress rate and causing a pile-up in the two free lanes. This lossy insertion also explains why there is still a spread at high intensities – the total number of vehicles entering the system remains constant once past the peak intensity.

B. Effect of Penetration Rate

We repeat the simulations by varying the penetration rates of VGrid-enabled vehicles. As shown in Figures 7, 8, and 9, there is a clear decrease in variance as the percentage of participating vehicles increases. 100% penetration corresponds to the results for non-zero transmission range presented in the previous section. 0% penetration corresponds to 0 transmission range. The relationship is approximately linear and exhibits similar behavior as penetration rate increases with increasing transmission range. Therefore, even partial deployment of VGrid will have an effect on traffic flow.

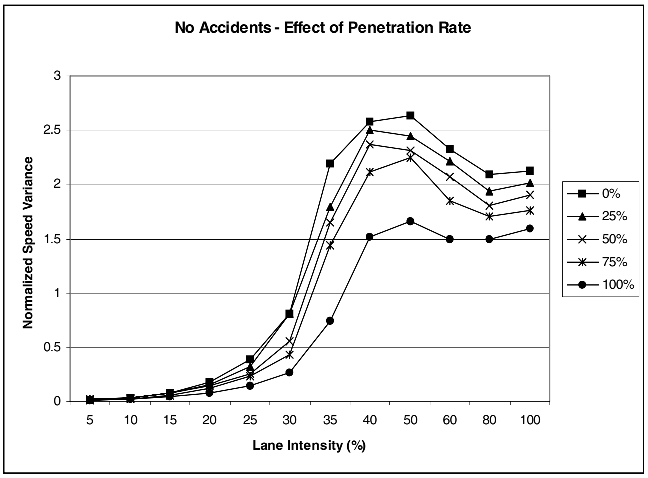

Fig. 7 - The performance of VSL only at varying penetration rates from 0% penetration to 100% penetration.

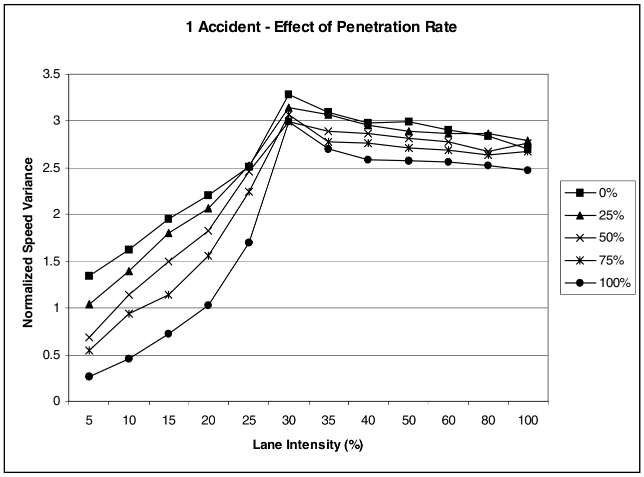

Fig. 8: The performance of VSL and accident alert with 1 fixed accident at varying penetration rates, from 0% penetration to 100% penetration.

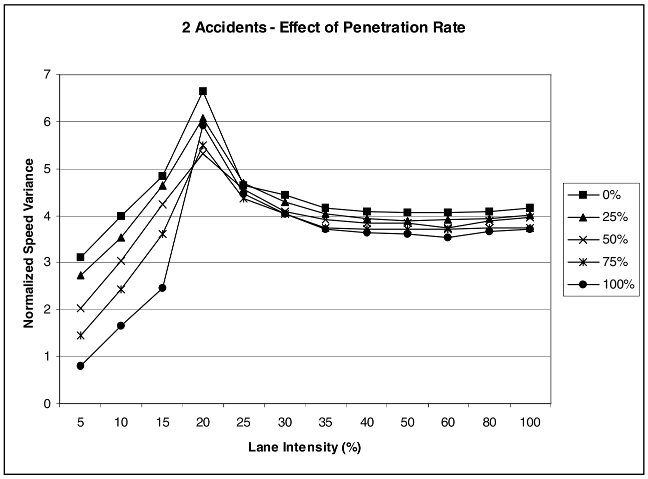

Fig. 9: The performance of VSL and accident alert with 2 fixed accidents at varying penetration rates, from 0% penetration to 100% penetration.

VI. RELATED WORK

There are many ongoing projects dealing with the integration of fixed roadside infrastructure with vehicle communications capabilities The Vehicle Information and Communication System (VICS) combines traffic information from police and highway administrators and broadcasts it via fixed transmission towers in real time. Vehicles with compatible navigation equipment are alerted to congested or regulated roadways kilometers away. [9] Advanced Cruise-Assist Highway System (AHS) utilizes roadside sensors and transceivers to reduce accidents through information propagation, early warning, and safety systems. Its goal is to prevent drivers from making the last-second errors in judgment that lead to accidents. [10]

Other projects focus on ad hoc vehicle-to-vehicle communications in order to facilitate safety as well as comfort. FleetNet is concerned with the development of vehicular wireless network protocols and applications that utilize location awareness and position information. [11] AutoNet uses ad hoc communication to develop a distributed transportation management system. Traffic conditions, vehicle behavior, environmental monitoring, and other safety and performance issues are addressed by the distributed computing platform. [12]

VII. CONCLUSION

There are countless applications for inter-vehicle wireless communications, and as vehicular computational and networking capabilities increase, what once required significant infrastructure, can be achieved in an ad hoc, distributed manner. In this paper, we outlined such a distributed framework for a vehicular computing grid that allows for the monitoring and control of traffic flow. In particular, we look at the homogenization of traffic in order to reduce accidents by lowering the variance of vehicle speeds on a highway. We show that the use of this communication framework allows drivers to make decisions such as lane adjustments and speed control at safer ranges and with greater precision than possible with only human perception.

VIII. ACKNOWLEDGEMENTS

This work is partly supported by National Science Foundation Grant CMS#9984239 and MICRO# 05-038.

REFERENCES

[1] D. Schrank and T. Lomax, "The Urban Mobility Report," Texas Transportation Institute, College Station, TX, 2002.

[2] E. van den Hoogen and S. Smulders, "Control by variable speed signs: Results of the Dutch experiment," *In Proceedings of the 7th International Conference on Road Traffic Monitoring and Control*, pp 145-149, April 1994.

[3] M. Vechler, W. Franz. Mobile Internet Access in FleetNet. http://www.fleetnet.de/.

[4] Drive-thru Internet: IEEE 802.11b for Automobile Users. http://www.drive-thru-internet.org/.

[5] I. Chisalita and N. Shahmehri. A peer-to-peer approach to vehicular communication for the support of traffic safety applications. In *ITSC*. IEEE, 2002.

[6] K. Nagel and M. Schreckenberg,"A Cellular Automaton Model for Freeway Traffic," Journal de Physique, Vol. 2, pp. 2221-2229, 1992.

[7] W. Knospe *et al,* "Towards a realistic microscopic description of highway traffic," *J. Phys.* A: Math, Gen. 33 No 48, Dec 2000.

[8] K. Bolay, Cellular Automaton Traffic Simulators based on the work of Nagel-Schreckenberg (1992), Takayasu (1993), Helbing and Schreckenberg (1999). http://rcswww.urz.tu-dresden.de/~helbing/RoadApplet/

[9] Vehicle Information and Communication System Center. http://www.vics.or.jp

[10] Advanced Cruise-Assist Highway System Resarch Association. http://www.ahsra.or.jp

[11] FleetNet. http://www.et2.tu-harburg.de/fleetnet/

[12] AutoNet: Ad hoc Peer to Peer Information Technology for Traffic Networks. http://www.its.uci.edu/~mcnally/mgm-autonet.html