UV-CAST: An Urban Vehicular Broadcast Protocol

Wantanee Viriyasitavat¹, Fan Bai², and Ozan K. Tonguz¹

¹Carnegie Mellon University, Dept. of Electrical and Computer Engineering, Pittsburgh, PA 15213-3890, USA

²General Motors Corporation, ECI Lab, Warren, MI 48092, USA

E-mail: wviriyas@ece.cmu.edu, fan.bai@gm.com, tonguz@ece.cmu.edu

Abstract—Several vehicular communication applications will involve multicast/broadcast communications where all vehicles in a certain region of interest are the intended recipients of particular messages. While there are several existing broadcast routing protocols for highway vehicular ad hoc networks (VANETs), very few solutions exist for urban VANETs in cities like New York City or Chicago. This paper attempts to fill this gap by proposing a new broadcast routing protocol, namely Urban Vehicular BroadCAST (UV-CAST) protocol, that addresses both the broadcast storm and disconnected network problems in urban VANETs. Key challenges imposed by urban VANETs as well as new mechanisms needed for meeting these challenges are identified and presented. Performance of the proposed UV-CAST protocol is evaluated in terms of network reachability, received distance, and network overhead in ideal Manhattan Street scenarios as well as in real cities, such as Pittsburgh. The overall performance of UV-CAST is excellent.

I. INTRODUCTION

In the past decade, a number of studies and research efforts have focused on developing routing protocols for Vehicular Ad hoc Networks (VANETs). Such protocols can be grouped based on: i) the types of VANET applications they support (e.g., safety, traffic efficiency or commercial applications) [1]; ii) the types of environments and scenarios assumed (e.g., *one-dimensional* highways and/or *multi-dimensional* urban environment; disconnected or well-connected network regimes) [2]; and iii) the infrastructure support required (e.g., whether or not the protocol requires some infrastructure such as Road Side Units (RSUs), repeaters, etc.) [3]. An ideal routing protocol should therefore be able to support a particular VANET application in all possible network regimes and in all scenarios (highway, suburban, and urban settings) regardless of infrastructure support. Since significant percentage of accidents (39.5%, 52.5%, and 54% of accidents in the States of California, Missouri, and Pennsylvania, respectively [4], [5], [6]) occur in urban areas, there is an acute need for the successful dissemination of latency-sensitive safety messages in urban areas.

While some protocols have recently been reported for urban scenarios, they are designed for well-connected urban environments or disconnected scenarios that require infrastructure support; far fewer studies exist on disconnected networks in urban scenarios with no infrastructure support.

In this article, we propose a broadcast routing protocol for urban VANETs. The proposed protocol is completely distributed and it supports both disconnected and well-connected network regimes in urban scenarios with zero infrastructure support. Note that while there are a large number of vehicles

in urban areas, the network of these vehicles may operate in a disconnected network regime, especially during the initial deployment phase of DSRC during which only a *small* fraction of vehicles on the road will be new and DSRC-equipped. Thus, even in a rush hour period with traffic congestion, the network of DSRC-equipped vehicles in an urban environment might be operating in the disconnected network regime as the disconnected network problem will persist until the DSRC penetration rate reaches a certain threshold. The proposed UV-CAST protocol is evaluated via extensive simulations which show that the protocol, with a small increase in overhead, is able to deliver a specific message to all intended vehicles in a timely manner, giving the drivers an opportunity to revise their routes.

The remainder of this article is organized as follows. In the next section, we summarize the existing literature relevant to this subject. We then present key challenges that makes the design of an urban routing protocol complex and unique. The main mechanisms for addressing these challenges and the implementation of the proposed protocol are described in the following section. The simulation setting used and the preliminary evaluation of the proposed UV-CAST protocol are then presented. The key implications of the results are discussed in the following section, and the final section concludes the article.

II. RELATED WORK

It has been shown that routing protocols, designed for highway VANETs [2], cannot be directly mapped or applied to urban VANETs [1], [7] as they do not work well in this new two-dimensional environment. Due to space limit, only routing protocols for urban scenarios will be discussed.

Urban Multi-hop Broadcast (UMB) protocol proposed by Korkmaz et al. is a MAC layer solution for disseminating messages to all vehicles [3]. In this protocol, each vehicle contends for the channel by transmitting a variable-length black-burst; the vehicles with the longest burst end up forwarding the message. Vehicles (or repeaters) at intersections also create additional directional message broadcasts to other road directions. Yi et al. have proposed the StreetCast protocol in [8] which is also a MAC layer protocol that comprises three components: i) relay-node selection (RSUs at intersections choose the optimal relay vehicles); ii) Multicast Request-To-Send (MRTS) handshaking which is used to avoid collisions and hidden terminal problem; and iii) adaptive beacon control which is used to avoid the broadcast storm problem caused by

hello messages at a crowded intersection. However, both UMB and StreetCast protocols assume that the network is always well-connected; no solutions for disconnected networks have been reported.

Costa et al. propose in [9] a Direction-aware Function-Driven Feedback augmented Store and Forward Diffusion (DFD-FSFD) scheme. Each vehicle, upon receiving the message, computes a forwarding probability based on the proposed message propagation function that encodes information about target areas and preferred routes. In the case where the network is disconnected, vehicles store and periodically rebroadcast the message. AckPBSM, an Acknowledged Parameterless Broadcast in Static to highly Mobile protocol, is proposed by F. J. Ros et al. in [1]. This protocol uses the Connecting Dominating Set (CDS) concept for broadcasting in well-connected networks. Message reception acknowledgement, piggybacked in periodic hello messages, is used for relaying the message in a disconnected network. In this scheme, vehicles, upon receiving the messages, have to wait for the new hello message in order to compute their wait time. Hence, the message latency depends on the hello message interval which might cause additional delay when the hello interval is large. Both [1] and [9] are evaluated in ideal Manhattan-like urban scenarios.

In this article, we propose a *fully distributed, lightweight, and zero-infrastructure support* broadcast protocol that can support both well-connected and disconnected network regimes for broadcast applications in urban areas. The proposed protocol utilizes both direct relays through multi-hop transmissions (i.e., spatial relay) and indirect packet relays through the “store-carry-forward” mechanism (i.e., temporal relay). The protocol has been evaluated extensively both in ideal Manhattan-like and real city scenarios.

III. ROUTING CHALLENGES IN URBAN SETTINGS

In this section, we illustrate that the successful dissemination of latency-sensitive messages is a much more challenging requirement to meet in two-dimensional urban areas than in one-dimensional highway scenarios. Such complications arise from the additional dimension of the urban network topology and existence of intersections. Some of the key routing issues/problems that distinguish the protocol design for urban scenarios from the design for highway scenarios are:

1) Low penetration rates in the next ten years

The partial (and low) penetration rate of DSRC technology during the initial deployment stage could exacerbate the disconnected network problem even in dense urban areas. Hence, it is clear that a broadcast protocol in urban areas that does not rely on existing infrastructure has to address first and foremost this inevitable transition problem.

2) Omni-directional message direction and Region of Interest (ROI)

In order to determine the appropriate ROI for a VANET application, one has to consider whether *all the vehicles in a particular geographical location that travel in a particular direction would be interested in the*

broadcast message. This implies that the ROI of a particular application should be determined not only by the geographical area but also by the route itinerary of individual vehicles. Obtaining such information could be cumbersome, if not impossible, due to several privacy issues.

Note also that the ROI depends on the type of applications considered (e.g., post-crash notification, Emergency Electronic Brake Light (EEBL) applications, etc). Hence, ROI and dissemination direction of the message (i.e., *message direction*) may vary and cannot be uniquely determined. Depending on the type of applications and scenarios considered, the message may be directed in several directions (e.g., in the simple Manhattan Street topology shown in Figure 1, the message might be directed in one or several of the 4 directions - North, South, East, West) or it may encompass 360 degrees in general settings.

Fig. 1. Region of Interest (ROI) and direction of the messages in two scenarios: highway (left) and urban (right) scenarios. Green and blue arrows indicate the entry and exit points of the ROI, respectively.

3) Direction change at intersections in 2-D urban scenarios

Due to the possible direction changes of vehicles at intersections, it is not obvious which vehicles should be responsible for storing, carrying and forwarding (i.e., temporally relaying) the message. Unlike highway scenarios, where the temporal relay node is always the farthest vehicle traveling in the direction opposite to the message direction (such a vehicle has the smallest *re-healing time*, i.e., time to encounter new vehicles) [2], the same method may not work in urban scenarios; i.e., only the farthest vehicle criterion might be inadequate, as it will relay the message only to a subregion of a city. Hence, the traditional store-carry-forward (SCF) mechanism (i.e., the selection of SCF-agent vehicles) used in one-dimensional highway protocols might not be an appropriate solution for urban areas.

4) Multiple “enter” and “exit” points to the ROI

Figure 1 depicts another major difference between highway and urban scenarios. While there is only one “entry” and one “exit” locations in the ROI of highway scenarios (as indicated with green and blue arrows, respectively), the ROI in urban scenarios typically has several locations where vehicles can enter and exit from. Because of the multiple “entry” points, *it is no longer realistic to assume that if the message reaches the end of the relevant*

area, the complete section of the road has already been covered. Vehicles that enter into the already-covered area may not receive the message if they arrive at a later time, after the time the message (or the “shock wave”) passes through the area.

5) Connectivity of a vehicle depends on its location

It is clear that the transmission coverage of vehicles in urban areas differs depending on their geographical locations. The transmission coverage of the intersection vehicles may cover more “road area” than that of vehicles between intersections (i.e., non-intersection vehicles). Hence, the intersection vehicles are more likely to have better connectivity, i.e., they have a higher number of neighbors. Thus, an urban routing protocol should utilize this non-uniform transmission coverage characteristic which is unique to the urban environment [10].

IV. UV-CAST: URBAN VEHICULAR BROADCAST PROTOCOL

A. Main mechanisms

Based on the key challenges and insights described in the previous section, we propose that the broadcast routing protocol for urban VANETs should possess the following features:

1) More than one vehicle should be responsible for the store-carry-forward (SCF) task

Due to the omni (or multi-) directionality of message direction and the ROI, if there is only one vehicle responsible for the SCF task, the message will be temporally relayed to the region through which such vehicle passes before it leaves the ROI (i.e., only a subregion of a given ROI will be covered). Therefore, in order to relay messages in many directions, *SCF task should be assigned to more than one vehicle*. This mechanism is crucial, especially during the initial deployment phase of DSRC where only a small fraction of vehicles will be DSRC-equipped.

2) SCF-assigned vehicles should “forward” the message more than once

Due to the possible changes in vehicles’ direction and the fact that the ROI in urban areas has several entry and exit points, it is clear that a vehicle will encounter uninformed neighbors again and again (at different points in time). Thus, vehicles assigned as “agents” for SCF should continue to carry and forward the message even though they have already relayed their messages to the new neighbors. However, such modification may cause a lot of rebroadcasts; some of which will clearly be redundant. In order to avoid such unnecessary rebroadcasts, the routing protocol should restrict the message rebroadcasts as opposed to blindly rebroadcasting the message whenever SCF-agent vehicles meet new neighbors.

One solution to avoid redundant rebroadcast is to use message acknowledgment in periodic hello messages. For example, an additional 4-byte field (called *message id* field) should be added to the message header which

stores id of messages that a vehicle has recently received. With this acknowledgment mechanism, an SCF-assigned vehicle can decide whether it should rebroadcast the message upon receiving hello messages from its neighbors.

3) Intersection-based broadcast storm suppression mechanism

As mentioned previously, intersection vehicles typically have more neighbors (i.e., better network connectivity) than non-intersection vehicles, especially in a network with high traffic density. Message rebroadcasts from the intersection vehicles thus are likely to reach more vehicles within a shorter time, as compared to the case where the same number of messages are rebroadcasted from the non-intersection vehicles. Intersection-based broadcast storm suppression scheme is therefore expected to be more effective and efficient than other non-intersection-based schemes.

B. Protocol Design

In this section, we present the principle of operation of the UV-CAST protocol. As depicted in Figure 2, the UV-CAST protocol comprises two main components which are responsible for well-connected and disconnected network regimes (as indicated with the dotted rectangles). Since UV-CAST is a completely distributed broadcast protocol, each vehicle operates independently and, based only on the local information it has, decides which network regime it belongs to. In other words, if a vehicle is chosen to perform the SCF task, as described in the next subsection, it will operate in the disconnected regime.

Fig. 2. Flowchart describing the operation of the UV-CAST protocol. (Abbreviations: ROI–Region of Interest, SCF–Store-carry-forward mechanism).

1) Selection of Store-carry-forward (SCF) vehicles:

In our implementation, store-carry-forward (SCF) task is assigned to vehicles that have small expected re-healing time (i.e., time before they see new neighbors). However, due to two-dimensional road topology and possible changes of vehicles’ directions at intersections, it is impossible to compute the re-healing time of a vehicle in urban areas based only on its location and direction. Hence, we assign the SCF task to

vehicles on the boundary of the connected component since they, with high probability, have a smaller re-healing time as compared to the non-boundary vehicles. To elaborate on this, let us consider an example depicted in Figure 3. Vehicles in the shaded region belong to the connected component. Thus, it is obvious that the boundary vehicles, as indicated with blue color, are more likely to meet other uninformed vehicles (e.g., outside the connected component) before the non-boundary vehicles (as depicted by red vehicles) do. Therefore, such boundary vehicles are the primary candidates for the SCF task. Algorithm 1 describes the steps used to select boundary vehicles.

Fig. 3. Store-carry-forward (SCF) assignment mechanism of the UV-CAST protocol. Assume that the message is broadcasted from *Src*. Vehicles that are connected to *Src* via spatial relays are in the gray shaded region and thus belong to the connected component. Boundary and non-boundary vehicles are indicated with blue and red dots, respectively. Black squares indicate vehicles that do not belong to the connected components and are disconnected from *Src*.

Algorithm 1 Distributed Gift-wrapping Algorithm for SCF-agent Vehicles Selection [11] (See Figure 4 for illustration).

$\angle(A, S, i) \Leftarrow$ angle between a vector from Vehicle *A* to Vehicle *S* and another vector from Vehicle *A* to Vehicle *i* where $\angle(A, S, i) \in [-\pi, \pi]$.
 $Nbr(A) \Leftarrow$ set of all neighboring vehicles of Vehicle *A*

When *A* receives the message for the first time from Vehicle *S*

for all $i \in Nbr(A) \setminus \{S\}$ **do**

$\theta_i \Leftarrow \angle(A, S, i)$

end for

$\theta_- \Leftarrow \min(\min_i(\theta_i), 0)$

$\theta_+ \Leftarrow \max(\max_i(\theta_i), 0)$

if $|\theta_+| + |\theta_-| < \pi$ **then**

$A \Leftarrow$ SCF task

end if

Fig. 4 illustrates how the boundary vehicles are selected by Algorithm 1. Upon receiving a message for the first time from Vehicle *S*, Vehicle *A* computes the angle θ_i for all of its neighbors (see Figure 4 (left)). Maximum (θ_+) and minimum (θ_-) angles are then identified. In the scenarios given in the middle and right plots of Figure 4, Vehicles *B* and *C* are the neighbors of Vehicle *A* that have the maximum and minimum angles, respectively. If $|\theta_+| + |\theta_-|$ is less than π , then Vehicle *A* is selected as a boundary vehicle and assigned the SCF task (which is the case for the scenario in Figure 4 (middle) but not for the scenario in Figure 4 (right)). The following claim

proves that all of the boundary vehicles will be chosen as SCF-agent vehicles by the distributed gift-wrapping algorithm.

Claim: If Vehicle *A* is a boundary vehicle, then the algorithm will select *A* to be an SCF-agent.

Proof: To prove the above claim is equivalent to proving the following statement: *if A is not chosen to be an SCF-agent, then A is not a boundary vehicle.*

Assume that *A* is not chosen by the algorithm.

Hence, this implies that $|\theta_+| + |\theta_-| \geq \pi$. Let Vehicles *B* and *C* (see Figure 4) be the vehicles that makes the largest and smallest angles with vector between *A* to *S*, respectively. Since $|\theta_+| + |\theta_-| \geq \pi$, Vehicle *A* lies in the interior of a triangle formed by Vehicles *B*, *C*, and *S*, which implies that Vehicle *A* cannot be a boundary vehicle with respect to Vehicle *S*. This completes the proof. ■

However, it is worth mentioning that since the distributed version is only an approximation of the *centralized gift wrapping algorithm* (which assumes global knowledge about the network topology), the algorithm might over-select the SCF-agent vehicles; i.e., since all boundary vehicles are SCF-agent vehicles but not vice versa, a set of SCF-agent vehicles selected by the proposed distributed gift-wrapping algorithm is always a subset of a set that contains all boundary vehicles. This discrepancy indicates the trade-off between accuracy of the algorithm and presence or absence of global knowledge. Since such global knowledge requires excessive message exchanges between vehicles, we use the *distributed gift-wrapping algorithm* which relies solely on local information in the proposed UV-CAST protocol. As will be shown later, one-hop neighbor information is sufficient.

2) **Wait time calculation:** This mechanism corresponds to the shaded rectangular box in Figure 2. Upon receiving a new message from Vehicle *j*, Vehicle *i* computes its wait time τ_i as follows:

$$\tau_i = \begin{cases} \frac{1}{2} \left(1 - \frac{d_{ij}}{R} \right) \tau_{max} & \text{if } i \text{ is at an intersection} \\ \frac{1}{2} \left(2 - \frac{d_{ij}}{R} \right) \tau_{max} & \text{otherwise} \end{cases} \quad (1)$$

where d_{ij} is the distance between Vehicle *i* and Vehicle *j*, *R* is the maximum transmission range, and τ_{max} is the maximum waiting time (which is set to 500 ms in our study). Once the timer expires and Vehicle *i* does not receive any duplicate message, Vehicle *i* rebroadcasts. Otherwise, Vehicle *i* suppresses its rebroadcast. Note that the value of maximum waiting time τ_{max} needs to be carefully chosen. If the value of τ_{max} is too low, most vehicles rebroadcast before they could receive duplicate messages from their neighbors (in which case, they should instead suppress their rebroadcasts). On the other hand, if τ_{max} is too high, then redundant rebroadcasts can be avoided at the expense of an increase in delay. It should be mentioned that the wait time can be computed in different ways and the specific method presented in Equation (1) is only an illustrative example. Sensitivity study on the wait time calculation is an important subject of future study.

Fig. 4. The principle of operation of the proposed distributed gift-wrapping algorithm. While Vehicle A in the scenario shown in the middle figure is selected as a boundary vehicle, Vehicle A is not selected as a boundary vehicle and not assigned with the SCF task in the scenario shown in the right figure. Assume that $NBR(A) = \{S, B, C, D, E\}$ and that A first receives the message from Vehicle S .

V. SIMULATION SETTING AND PERFORMANCE METRICS

A. Network Topology and Mobility Model

In the simulations, we assume a $1 \text{ km} \times 1 \text{ km}$ network topology with 8 evenly-spaced horizontal and vertical streets; each street has two lanes, so the traffic can travel in both directions. All road junctions are equipped with pre-timed signals and the underlying mobility model that governs vehicle movement is based on the work reported in [12].

In addition to Manhattan Street scenario, we also evaluate the proposed UV-CAST protocol in a real city. In this preliminary study, we choose downtown Pittsburgh as an illustrative example. The road topology of downtown Pittsburgh is far more irregular than the conventional Manhattan Street topology. Not all streets are perpendicular and there exist several T-junctions and junctions joining more than 4 streets (e.g., a 5-street junction). Map of the simulated area is shown in Figure 5 with a total of 2.3 km^2 area, 173 intersections, and 97 streets. The underlying mobility model used to simulate realistic vehicle movement is based on the one used in the open-source SUMO simulator [13].

In all scenarios considered, the source of message (i.e., the scene of accident) is located at the intersection approximately at the center of the network. The message source broadcasts a message only once at time $t = 1000$ seconds and the simulation ends two minutes after the source broadcasts the message. The region of interest is assumed to be a $1 \text{ km} \times 1 \text{ km}$ area around the scene of accident. Two transmission ranges are used: 250 and 140 meters for Line-Of-Sight (LOS) and Non-LOS communication ranges, respectively [7]. All the simulations are performed using an open-source *ns-2* simulator [14] where PHY and MAC layers are implemented based on the WAVE standard.

B. Metrics

The following four metrics are used to evaluate the performance of the proposed protocol. While reachability and received distance metrics determine the protocol reliability and effectiveness, transmission and reception overhead metrics quantify the efficiency of the protocol.

- 1) **Network reachability** measures the fraction of vehicles in the region of interest that receive the message. A good protocol must ensure that most, if not all, vehicles

Fig. 5. Road topology of downtown Pittsburgh. Vehicles are indicated with red dots.

intended to receive the message do receive the message before they arrive at the accident scene.

- 2) **Received distance** is the closest Euclidean distance to the accident scene from the arc on the trajectory of a vehicle which was at a point (say A) when the message was broadcast but later receives the message at point B . Note that if a vehicle immediately receives the message at the time the message was broadcast, then the points A and B are identical.

Figure 6 illustrates how the received distance metric is computed. Consider the scenario in Figure 6(a), where the vehicle is at point A when the message was broadcast but receives the message at point B_1 . The received distance in this case is d_{min}^1 (see Figure 6(c)) which is in the minimum of d on the trajectory of the vehicle from A to B_1 . Similarly, if the vehicle receives the message at point B_2 (see Figure 6(b)), then the received distance in this case is d_{min}^2 (see Figure 6(c)).

While the reachability metric indicates whether all the vehicles receive the message, this metric indicates whether vehicles receive the broadcast message just-in-time so that they can reroute and avoid passing through the scene of accident. Hence, the larger the received distance metric, the better the protocol. In addition, observe that the message latency is implicitly captured by this metric.

- 3) **Transmission Overhead** measures the total number of messages *transmitted* into the network by all vehicles.

Fig. 6. Examples illustrate how the received distance metric is computed. The dotted curvy arrows shown in (a) and (b) are the trajectory of the vehicle while A , B_1 , B_2 , and C are the locations of vehicles and d_0 , d_1 , d_{min}^1 , d_2 , and d_{min}^2 are the Euclidean distances between the accident scene and the locations of the vehicles. Figure (c) shows the Euclidean distance between the accident scene and the vehicles at different locations on the trajectory. Clearly Scenario 1 in Figure (a) is better than Scenario 2 in Figure (b) as it allows time to circumvent the accident scene.

This metric is important as it indicates whether or not the message transmission generated by the proposed protocol overwhelms the network; in other words, whether it uses excessive amount of bandwidth.

- 4) **Reception Overhead** measures the average number of duplicate messages *received* at a vehicle. This metric determines whether the protocol can effectively solve or mitigate the broadcast storm problem.

VI. SIMULATION RESULTS

A. Manhattan Street scenarios

Observe from Table I and Figure 7 that the proposed UV-CAST protocol performs quite well in terms of reachability as compared to the optimal scheme (i.e., the reachability obtained when an ideal¹ flooding scheme is assumed). UV-CAST also delivers messages to vehicles in a timely manner: a vehicle, on average, receives the broadcast message when it is at least 700 meters (about 5 blocks in a typical city) away from the scene of accident; hence, there is sufficient time for the drivers to avoid passing through the accident area. Furthermore, the overhead caused by UV-CAST is small and scales well with traffic density; even in a dense network (e.g., a network with 300 veh/km² network density), we have observed that approximately 12% of vehicles rebroadcast. While there are 50 additional packets transmitted to the network, at a density of 300 veh/km², this corresponds to less than 0.2 additional messages per vehicle which is very modest and reasonable. Furthermore, a vehicle, on average, receives less than 3 duplicate messages² (see the red curve in Figure 7).

¹The scheme where we assume no message collision or message loss.

²No packet drop has been observed.

Fig. 7. Performance of the UV-CAST protocol in Manhattan Street scenarios.

This is because, for a given area in a well-connected network, a fixed number of packet retransmissions from vehicles is sufficient to cover the entire area. Even though the number of vehicles increases in a dense network, transmission and reception overhead increase only slightly once the number of retransmissions reaches a certain number and the entire area has been covered. In contrast to existing solutions [1], the proposed UV-CAST protocol requires less than 1 (instead of at least 2 as reported in [1]) transmission per vehicle in order to achieve almost 100% message reachability in the Manhattan scenario. This implies 100% improvement in terms of network overhead which is significant especially in a dense network.

TABLE I
AVERAGE NETWORK REACHABILITY OF THE UV-CAST PROTOCOL AND THE OPTIMAL SCHEME IN THE MANHATTAN STREET SCENARIOS.

Network Density [veh/km ²]	UV-CAST	Optimal scheme
20	83.25%	93.25%
40	97.13%	100%
60	99.08%	100%
80	99.63%	100%
100	99.65%	100%
200	99.72%	100%
400	99.89%	100%

B. Real cities

Performance of the UV-CAST protocol in downtown Pittsburgh is shown in Figure 8. Simulation results (not shown here due to space limits) indicate that *all* vehicles within 1 km x 1 km region of interest receive the message. Observe that regardless of the traffic density, the UV-CAST protocol is able to deliver messages to vehicles on time. A vehicle, on average, receives the message when it is at least 300 meters from the scene of accident (2 – 3 blocks away from the accident). In addition, the reception overhead resulting from the UV-CAST protocol scales well with traffic density: observe in Figure 8 that in a network with 40 veh/km² vehicle density, while there are 80 additional packets transmitted into the network, this corresponds to only about 2 additional messages/vehicle transmitted into the network. Nonetheless, each vehicle receives, on average, less than 3 duplicate messages. Note that the discrepancy in the transmission overhead between the Manhattan street and real city scenarios is mainly due to the irregularity of the road topology.

Fig. 8. Performance of the UV-CAST protocol in downtown Pittsburgh, PA.

C. Future Work

As future work, we plan to further investigate the following issues:

- 1) *There may be several safety messages needed to be disseminated at a given point in time.*

In the scenario where there are several messages to be transmitted, size of such hello messages can be large as they have to include id's of all broadcast messages. Large hello messages are undesirable since these messages are periodically broadcast by each vehicle in the network and therefore can potentially overwhelm the network. Further research is needed to address this issue.

- 2) *A subset of boundary vehicles may be sufficient.*

It should be noted that in some scenarios, especially scenarios with medium and high traffic density, selecting all of the boundary vehicles would be unwise. As traffic density increases, size of connected clusters as well as number of boundary vehicles (or SCF-agent vehicles) also increase. However, an ideal SCF-agent selection algorithm should behave accordingly and in a cognizant manner; i.e., the higher the traffic density, the fewer the SCF-agent vehicles. In other words, a subset (i.e., *not all*) of boundary vehicles should be selected to be SCF-agents. This is an important problem that needs to be addressed.

- 3) *Size and shape of ROI depend on applications*

Size and shape of the ROI depends on the topology and the type of applications considered. While we only consider the post-crash notification application in this article where we assume the ROI to be a square area around the scene of accident, it would be interesting to evaluate the UV-CAST protocol for other VANET applications as well, such as the Emergency Electronic Brake Light (EEBL) (which has a much smaller ROI), Traffic congestion notification, etc.

- 4) *Realistic propagation environments*

While the simulations have shown that UV-CAST has an excellent performance with the specific LOS/NLOS propagation model assumed, performance under different (and more realistic) propagation environments [10], [15] should also be studied. Another interesting issue is to determine the effects that GPS errors may have on the performance of the distributed gift-wrapping algorithm used for SCF-agent selection (the current algorithm assumes a perfect location information of all neighboring vehicles).

VII. CONCLUSION

We have proposed a new broadcast routing protocol, UV-CAST, for *urban VANETs* which assume zero infrastructure support. UV-CAST is a completely distributed broadcast protocol and it can be implemented by using only the local information available to each vehicle in an urban VANET. The protocol is designed by taking into account the two-dimensional road topology in urban settings. In contrast to one-dimensional highway scenarios, routing protocol design in urban areas is a much more challenging task for many reasons: i) direction of vehicles in urban areas may change at intersections while direction of vehicles on highways do not change until they leave the highway; ii) while a message in highway scenarios is disseminated in only one direction, message dissemination direction in urban areas may encompass 360 degrees. Thus, the existing broadcast protocols designed for highway VANETs cannot be applied to urban settings.

The performance of the UV-CAST protocol has been evaluated in terms of reachability, received distance, and network overhead in a regular Manhattan Street scenario as well as in a real city (Pittsburgh, PA). Overall, the results show that the performance of the new UV-CAST protocol is excellent. While the proposed UV-CAST protocol assumes no infrastructure support, it can also utilize infrastructure support whenever it exists. Such infrastructure support could further enhance the performance of the UV-CAST protocol.

REFERENCES

- [1] F. J. Ros, P. M. Ruiz, and I. Stojmenovic, "Reliable and Efficient Broadcasting in Vehicular Ad Hoc Networks," in *Proc. of the IEEE 7th Vehicular Technology Conference Fall (VTC 2009-Spring)*, pp. 1–5, April 2009.
- [2] O. K. Tonguz, N. Wisitpongphan, and F. Bai, "DV-CAST: A Distributed Broadcast Protocol for VANETs," *IEEE Wireless Communications Magazine*, vol. 17, pp. 47–57, April 2010.
- [3] G. Korkmaz, E. Ekici, F. Özgüner, and U. Özgüner, "Urban multi-hop broadcast protocol for inter-vehicle communication systems," in *Proc. of the 1st ACM international workshop on Vehicular ad hoc networks (VANET)*, (New York, NY, USA), pp. 76–85, ACM, 2004.
- [4] Federal Highway Administration, "Highway statistics," July 2004.
- [5] Missouri Department of Transportation, "2008 missouri state highway system: Traffic crash statistics," 2008.
- [6] Pennsylvania County Crashes, "2006 pennsylvania crash facts and statistics," 2006.
- [7] W. Viriyasitavat, O. K. Tonguz, and F. Bai, "Dynamics of network connectivity in urban vehicular networks," *IEEE Journal on Selected Areas of Communications, Special Issue on Vehicular Communications and Networkings*, vol. 29, pp. 515–533, March 2011.
- [8] C.-W. Yi, Y.-T. Chuang, H.-H. Yeh, Y.-C. Tseng, and P.-C. Liu, "Street-cast: An Urban Broadcast Protocol for Vehicular Ad-Hoc Networks," in *Proc. of the IEEE Vehicular Technology Conference Fall (VTC 2010-Fall)*, pp. 1–5, May 2010.
- [9] P. Costa, D. Frey, M. Migliaavacca, and L. Mottola, "Towards lightweight information dissemination in inter-vehicular networks," in *Proc. of the 3rd international workshop on Vehicular ad hoc networks (VANET)*, (New York, NY, USA), pp. 20–29, ACM, 2006.
- [10] R. Bauza, J. Gozalvez, and M. Sepulcre, "Operation and performance of vehicular ad-hoc routing protocols in realistic environment," in *Proc. of the IEEE Vehicular Technology Conference (VTC)*, pp. 1–5, September 2008.
- [11] W. Viriyasitavat, F. Bai, and O. K. Tonguz, "UV-CAST: An Urban Vehicular Broadcast Protocol," in *The IEEE Vehicular Networking Conference (VNC)*, pp. 25–32, December 2010.
- [12] O. K. Tonguz, W. Viriyasitavat, and F. Bai, "Modeling urban traffic: a cellular automata approach," *IEEE Communications Magazine, Topics in Automotive Networking*, vol. 47, pp. 142–150, May 2009.

- [13] D. Krajzewicz, G. Hertkorn, C. Rossel, and P. Wagner, "SUMO (Simulation of Urban MObility); An open-source traffic simulation," in *Proc. of the 4th Middle East Symposium on Simulation and Modelling (MESM)*, pp. 183–187, 2002.
- [14] "Network simulator, ns-2," <http://www.isi.edu/nsnam/ns>.
- [15] M. Boban, T. T. V. Vinhoz, M. Ferreira, J. Barros, and O. K. Tonguz, "Impact of vehicles as obstacles in vehicular ad hoc networks," *IEEE Journal on Selected Areas in Communications*, vol. 29, pp. 15–28, January 2011.

serves or has served as a consultant or expert for several companies, major law firms, and government agencies in the United States, Europe, and Asia

VIII. BIOGRAPHIES

Wantanee Viriyasitavat received both of her B.S. and M.S. degrees in electrical and computer engineering from Carnegie Mellon University (CMU), Pittsburgh, PA in 2006. In the year 2006, she worked as a lecturer in Computer Science Department, Mahidol University, Bangkok, Thailand before she started her PhD study in July 2007 at Carnegie Mellon University under the supervision of Prof. Ozan Tonguz. Since 2007, she has been a Research Assistant at Carnegie Mellon University, where she has joined General Motors collaborative research lab and has been working on designing a routing framework for vehicular ad hoc wireless networks (VANET). Her research interests include traffic mobility modeling, network connectivity analysis, and routing protocol design for wireless ad hoc networks.

Fan Bai (General Motors Global R&D) is a Senior Researcher in the Electrical & Control Integration Lab., Research & Development and Planning, General Motors Corporation, since Sep., 2005. Before joining General Motors research lab, he received the B.S. degree in automation engineering from Tsinghua University, Beijing, China, in 1999, and the M.S.E.E. and Ph.D. degrees in electrical engineering, from University of Southern California, Los Angeles, in 2005. His current research is focused on the discovery of fundamental principles and the analysis and design of protocols/systems for next-generation Vehicular Ad hoc Networks (VANETs), for safety, telematics and infotainment applications. Dr. Bai has published about 40 book chapters, conference and journal papers. In 2006, he received Charles L. McCuen Special Achievement Award from General Motors Corporation in recognition of extraordinary accomplishment in area of vehicle-to-vehicle communications for drive assistance & safety.

Ozan K. Tonguz (tonguz@ece.cmu.edu) is a tenured full professor in the Electrical and Computer Engineering Department of Carnegie Mellon University (CMU), Pittsburgh, PA. He currently leads substantial research efforts at CMU in the broad areas of telecommunications and networking. He has published about 300 papers in IEEE journals and conference proceedings in the areas of wireless networking, optical communications, and computer networks. He is the author (with G. Ferrari) of the book *Ad Hoc Wireless Networks: A Communication Theoretic Perspective* (Wiley, 2006). He co-founded Virtual Traffic Lights (VTL), LLC, a CMU spin-off, in December 2010, which specializes in providing solutions to several transportation problems, such as safety and traffic information systems, using V2V and V2I communications paradigms. His current research interests include vehicular ad hoc networks, wireless ad hoc and sensor networks, self-organizing networks, bioinformatics, and security. He currently