Using Dead Drops to Improve Data Dissemination in Very Sparse Equipped Traffic

Sudarshan S. Chawathe

Department of Computer Science; University of Maine; Orono, ME 04469-5752, USA
chaw@cs.umaine.edu

Abstract— **We study the dissemination of data in a vehicular network in which the density of equipped vehicles, or all vehicles, is very low, and in which there is no centrally connected infrastructure, such as a network of gateways. In this environment, a vehicle is rarely in communication range of another vehicle, so that protocols based on forwarding messages along a chain of vehicles in communication range are not effective. Similarly, the lack of centrally connected network infrastructure precludes protocols that use a fixed network of gateways to transfer messages. Our solution is based on using a collection of stationary, stand-alone dead drops that exchange data with vehicles that pass by them. Although the dead drops do not communicate among themselves or with a central network, their limited ability to store and forward data can improve the connectivity of a vehicular network. We use simulation-based experiments to demonstrate that the use of dead drops results in significant improvements in both the rate of data dissemination and the robustness to frequent changes in the set of participating vehicles.**

I. INTRODUCTION

INTER-VEHICLE COMMUNICATION (IVC) is an important part of the intelligent vehicles landscape. IVC technologies enable a number of applications, including systems for early warning of accidents and other hazards, traffic information systems, and convenience applications such as traveler information systems. Although many such applications may also be implemented without IVC, for instance by using broadcast-based methods, IVC is a promising approach that can be used in conjunction with other alternatives to make more effective use of the available resources, such as bandwidth and fixed communication infrastructure.

The term IVC typically encompasses diverse kinds of communication, ranging from real-time communications for safety-critical applications to best-effort delivery mechanisms for comfort and convenience applications. In this paper, we discuss low-overhead methods for best-effort dissemination of data in a vehicular network. Thus, the target applications are those that can tolerate relatively long delays and non-guaranteed delivery of data. Although these characteristics rule out many applications, several other interesting and useful applications can make use of such a data dissemination mechanism. For example, traffic or tourist information may be delivered in this manner.

The majority of prior work on IVC has focused on urban or suburban environments with moderate to high traffic density, whereas our focus in this paper is on sparse traffic, as may be typical in rural or remote environments. Here, and in the sequel, we use the phrase traffic density to mean the density of *equipped* vehicles, i.e., the average or expected number of equipped vehicles per unit distance of a roadway. While we frame our descriptions for situations in which the raw density of traffic is low, our work is equally applicable in situations with dense traffic but very few equipped vehicles, as is often the case during the early stages of adoption of a new technology.

In sparse traffic, a primary communication challenge is that a vehicle is very rarely in communication range of another vehicle. Thus data propagation using methods that rely on routing or forwarding of messages via a chain of vehicles in communication range is not feasible. Further, if the expected density of traffic is very low, economic considerations make it impracticable to provide extensive fixed infrastructure, such as a system of numerous gateways with which vehicles may communicate, and that are connected to a central network.

Our solution to this problem is based on using *dead drops*, which are, in this context, stationary, *self-contained* devices that exchange data with vehicles within communication range. An important point of distinction here is that a dead drop is not connected to other dead drops or to any central network infrastructure, although such connectivity, if available, can be used to advantage. While, strictly speaking, they are a form of infrastructure, dead drops are low cost and easily deployed. For example, an easily deployable dead drop may be built using commodity hardware such as a small mobile processor, 802.11 radio, flash memory, and a battery, along with freely available software.

Before moving on to details, we note that our focus is on methods for disseminating data, as are typically used to disperse information about the environment sensed by one vehicle to others in the region. Such dissemination methods, often used with methods for selectively pruning and aggregating data, are different from general purpose network mechanisms, that must address issues such as routing. That is, our methods do not provide a general-purpose vehicular network; rather, they enable a certain class of data-driven

---

This work was supported in part by the U.S. National Science Foundation with grant CNS-0426683.

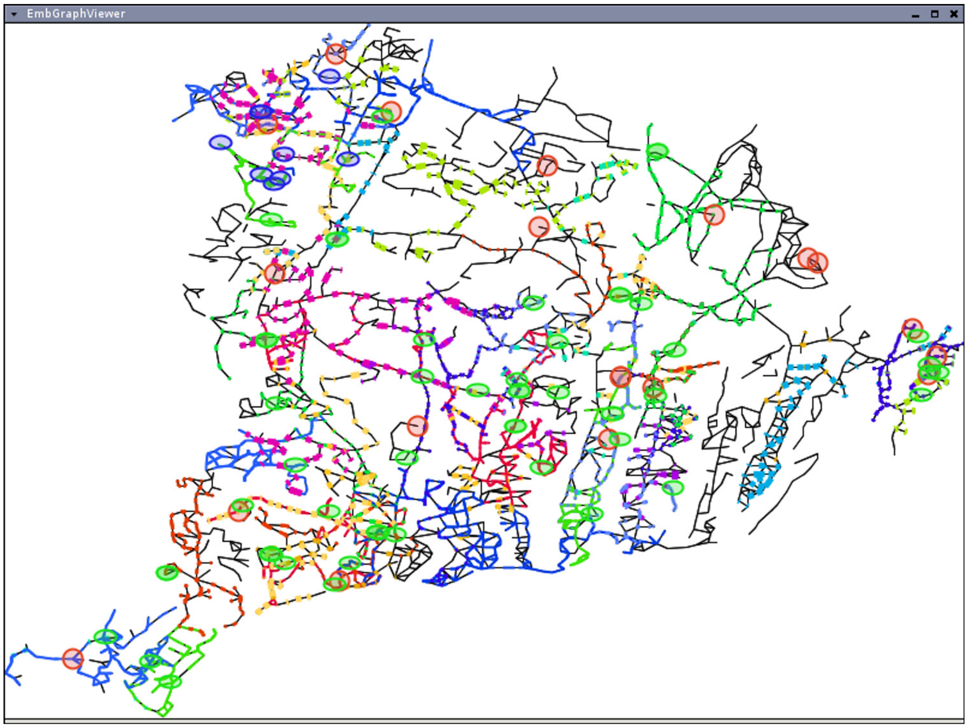

Fig. 1. A screen-shot of the visualization system displaying a typical experimental setup for the town of Falmouth, Massachusetts.

applications that are based on dissemination. Our methods do not depend on the use of a particular low-level communication protocol for data transfer between vehicles and dead drops. Rather, assuming the availability of some such method (discussed elsewhere in the literature), they provide a higher-level data dissemination service.

II. DATA DISSEMINATION WITH DEAD DROPS

Data may be disseminated in sparse traffic as follows: When a vehicle and a dead drop are in communication range, they exchange data. The vehicle then continues along its path with the data on board, repeating the data exchange at all dead drops it encounters. Similarly the dead drop continues to exchange data with all vehicles that pass by it.

This scheme for data dissemination is appealingly simple, but it raises some important questions: First, how many dead drops are needed in a given network of roads, and where should they be located? One may ask for the smallest number of dead drops, along with their placement, that satisfies some network-wide communication goals. We refer to this class of questions as the *static problem*. Second, how effective is this scheme in disseminating data in a timely manner? Intuitively, it is clear that if the traffic density is below some level, data will not move beyond a vehicle's travel path. Similarly, if the number of dead drops is too low, or their placement not advantageous, data may not travel very far. On the other hand, when vehicles and dead drops are very dense, there should be no reachability problems (although other problems, such as congestion of the communication channel, may appear). We refer to this class of questions as the *dynamic problem*.

In prior work [1], we studied some aspects of the static problem. In particular, we formulated the following optimization problem: Given a set of expected vehicle trajectories as input, what is the smallest number of dead drops, and their placement, that enables network-wide communication using the dead drops? This informal problem statement is formalized using a hypergraph representation of the road network and the vehicle trajectories and is mapped the problem of finding a minimum-cost spanning sub-hypergraph. We have shown the problem to be NP-hard, and provided an approximation algorithm that is optimal in the following sense: (1) If t is the maximum number of vehicle trajectories that intersect at a point, our algorithm finds a solution whose cost differs from the cost of an optimal solution by at most a factor of $H(t-1)$, where $H(n) = \sum_{i=1}^{n} 1/i$ is the i'th harmonic number. (2) Furthermore, there can be no algorithm that provides a guarantee that improves on the one provided by this algorithm by more than a constant factor. Although there are other aspects of the static problem that deserve further study, they are not the focus of this paper.

It is the dynamic problem that is the focus of this paper. Our solution to the static problem allows us to disseminate data network-wide using a modest number of dead drops. However, it does not provide any guidance on how effectively the data are disseminated. Informally, we wish to study how quickly data generated in one part of the network arrive at other parts. The answer depends on several parameters, such as the traffic density, its distribution, and the number and locations of dead drops. (We allow for the useful possibility of using more than the minimum number of dead drops required by the static solution in order to improve the speed of dissemination.)

In dense-traffic situations, bandwidth constraints preclude the use of a *flooding* scheme in which each vehicle immediately forwards all data it receives to all vehicles within range. Although such a scheme may appear to achieve very rapid data dissemination, it does not scale to more than a few dozen vehicles as the communication channel is typically overwhelmed by the rapid transmissions. Therefore, much effort has been invested in developing methods to make judicious use of the available bandwidth by limiting the amount of data transmitted. In this regard, sparse-traffic situations are different in two major ways: (1) When data are disseminated using dead drops, data travel mainly at vehicle (mechanical) speeds instead of at network (signal) speeds. In sparse traffic, data travel mainly on board a vehicle instead of hopping from vehicle to vehicle (or fixed station) at electronic speeds. (2) The amount of data to be disseminated is small due to the limited number of vehicles.

As a result of the above, condensing data and limiting their transmissions is not a pressing issue in the sparse-traffic situations we study. Certainly, data must eventually be limited, condensed, and expired. However, in sparse-traffic situations, we may treat this concern separately from the task of disseminating data. The methods we outline may be used in conjunction with a variety of methods for limiting data, in an orthogonal manner. (See Section IV.)

III. EXPERIMENTAL EVALUATION

We implemented a simulation and visualization environment in Java and the experiments were conducted on several PC-class machines running Debian GNU/Linux 4.0 (Etch) and Sun JDK version 1.5. In addition to allowing us to the study the effect of varying parameters such as traffic density and distributions of vehicle paths, our implementation also includes a real-time, interactive visualization component. A representative screen-shot from this component appears in Figure 1. This component allows us to monitor an ongoing simulation in order to check that the results of various parameter settings are intuitively realistic. For example, we used this approach to avoid settings that result in vehicle tracks that are unrealistically repetitive or too short.

Although we use synthetically generated vehicle trajectories, the underlying road network is real, based on data obtained from a number of U.S. agencies. For instance, the network suggested by Figure 1 is built using road data for the town of Falmouth, Massachusetts, from the *MassGIS* collection [2]. For concreteness and brevity, we describe our experimental results for this road network only, although we have similar results for several other towns and neighborhoods. (Although the results are similar, intuitively realistic settings of parameters such as traffic density vary based on the underlying network, depending, for example, on how interconnected the roads are.)

Our main interest is studying the effectiveness and, in particular the dynamic behavior, of data dissemination in sparse traffic conditions, with and without dead drops. To that end, our experiments measure the time required for data from an arbitrarily designated first vehicle to propagate to the others in the simulation. The horizontal axis of all charts in this section represents time, measured in units of simulation stages. (We have limited most charts to 250 stages for legibility; our experiments cover a larger range—see Figure 11). The vertical axis of most charts represents our main metric, which is the fraction of vehicles to which data from the first vehicle have been disseminated. We refer to this number as the *propagation ratio*. Also for legibility, we do not plot error-bars in the charts. Standard deviations of data points discussed are low enough to be hard to discern on these charts, so that all discussed differences are significant. As an example, the average standard deviation of data points in Figure 2 (all three curves) is 0.01.

In most charts, we plot the results for three methods: *VV*, when only vehicle-to-vehicle data dissemination is used; *DVV* when randomly placed dead drops are used to enhance dissemination, and *ODVV*, when the optimized scheme outlined in Section II is used.

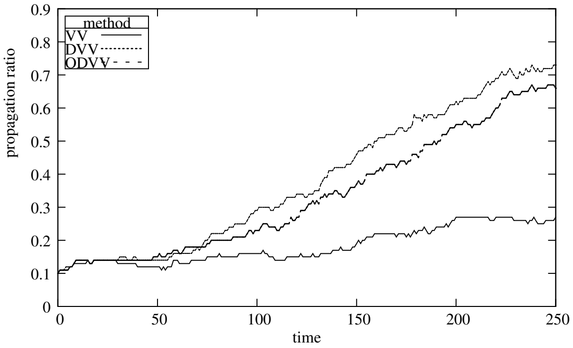

Fig. 2. Propagation ratio over time, very low density traffic ($V = 10$).

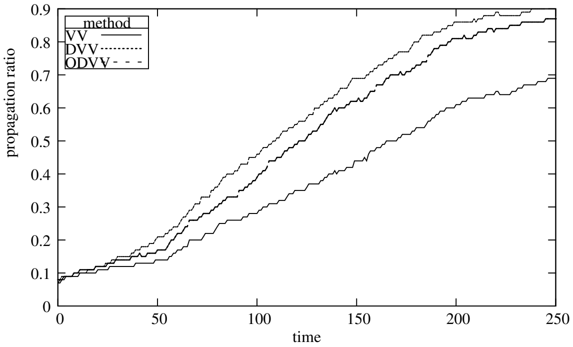

Fig. 3. Propagation ratio over time, low density traffic ($V = 20$).

Our first set of results concerns the effect of varying traffic densities. Recall that, throughout this paper, we use the phrase traffic density to mean the density of equipped vehicles. We measure density for a given road network simply as the average number, V, of concurrently active

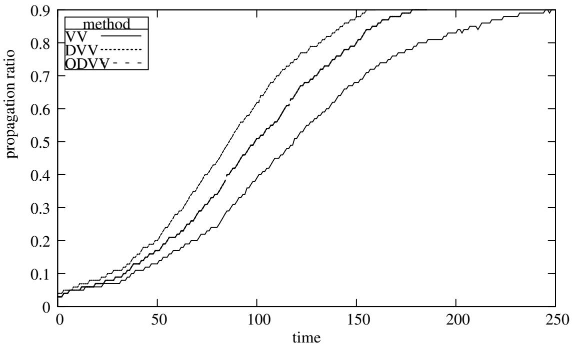

Fig. 4. Propagation ratio over time, moderate density traffic ($V = 40$).

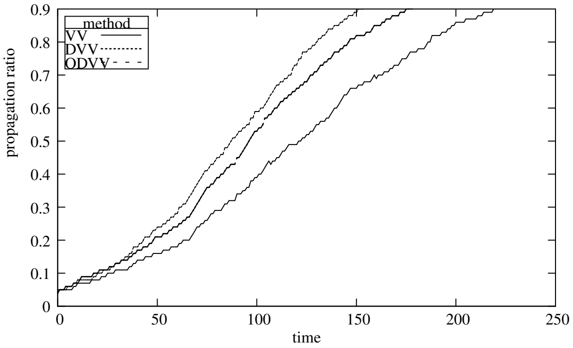

Fig. 5. Propagation ratio over time, high density traffic ($V = 80$).

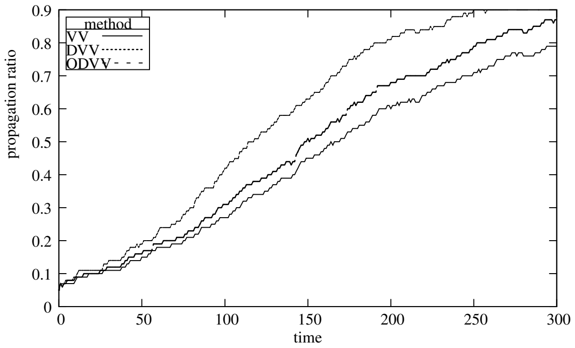

Fig. 6. Using 50 dead drops.

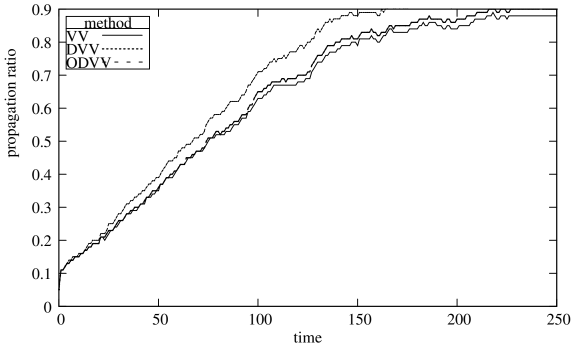

Fig. 7. Using 20 dead drops.

vehicles. For the Falmouth network (Figure 1), Figures 2, 3, 4, and 5 summarize the results for V values of 10, 20, 40, and 80, respectively. Note that the region covered by this road network is approximately 20 km (East-West) × 15 km (North-South). These values of V (especially the first two) thus represent very low traffic densities. In Figure 2, we observe a clear benefit to using dead drops and, further, to using the optimized scheme. For instance, after 150 cycles, VV results in dissemination to only about 15% of the vehicles, compared to 35% and 45% for DVV and ODVV. Further, while propagation ratios for DVV and ODVV continue rising, while that for the VV method remain low, increasing only gradually in this region. (A larger time window is studied later.) The separation between the schemes diminishes as the traffic density increases, progressing from Figure 2 to 3, 4, and 5, agreeing with the intuition that the benefit of dead drops diminishes as a larger number of vehicles are concurrently active in the network.

Our second set of results concerns the effect of varying the number D of dead drops placed in the road network. For the network of Figure 1, Figure 3 above uses $D = 100$. The results for D set to 50 and 20 appear in Figures 6 and 7, respectively, for the same values of other parameters, including $V = 20$ (but note the different maximum abscissa in Figure 6). These values were chosen to be intuitively suitable for the roughly 20 km × 15 km region in the experiment. We note that even a modest number of dead drops for this area provides a noticeable benefit. As may be expected, the incremental benefit of both DVV and ODVV over the baseline VV scheme diminishes with diminishing number of drops.

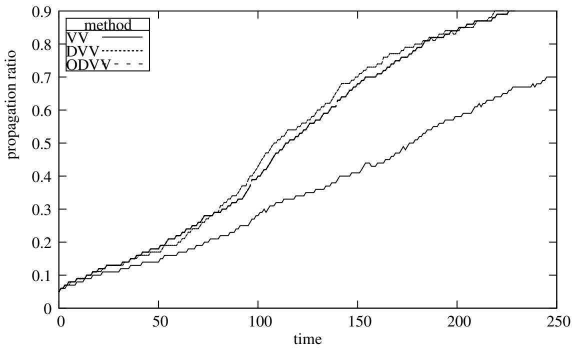

Fig. 8. Vehicles joining and leaving at a 2% rate.

Our third set of results concerns the effect of varying vehicle travel times, as measured by the rate at which vehicles enter and exit the system. In particular, we vary a parameter R where $R = 0$ corresponds to a static set of vehicles while $R = 1$ corresponds to one vehicle entering the system and one vehicle leaving the system at each step of the simulation. Figures 8, 9, and 10 summarize the results for R values of 2, 10, and 20 percent, respectively. In addition, the results of Figure 3 use the value of 5%, with matching values for the other parameters. The results confirm the intuition that if vehicles enter and leave the system more frequently then data dissemination is more difficult because a vehicle with the data may become inactive before it has

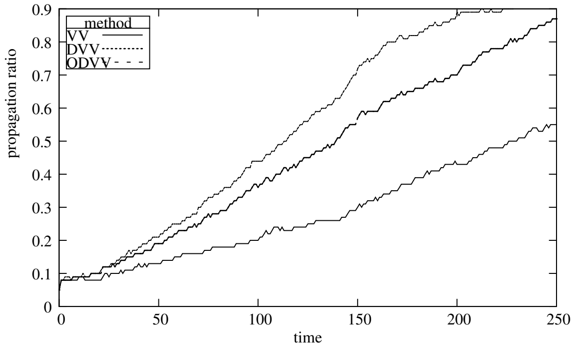

Fig. 9. Vehicles joining and leaving at a 10% rate.

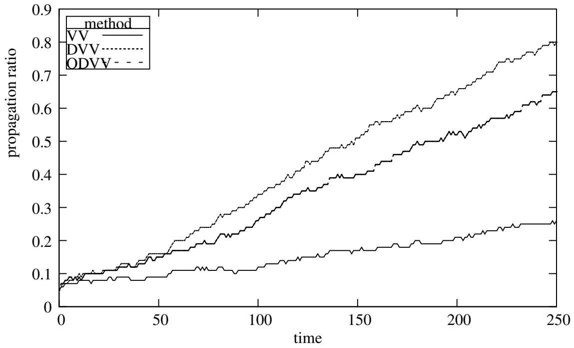

Fig. 10. Vehicles joining and leaving at a 20% rate.

a chance to transfer the data to other vehicles or a dead drop. More notably, the decrease in the rate of dissemination is much more drastic for the VV method than for the methods using dead drops. We observe an increase in the performance separation between ODVV and both DVV and VV. Intuitively, this result suggests that dead drops provide a bigger boost to dissemination when the set of vehicles in the network is more dynamic.

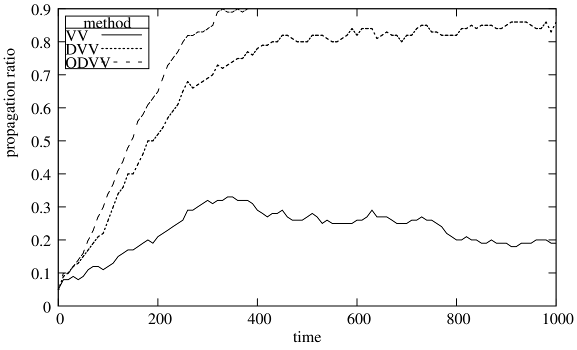

Fig. 11. An extended view of Figure 10

Figure 11 summarizes the results of the same set of experiments as used for Figure 10, but over a larger window of time. (Note the difference between the horizontal axes.) As suggested earlier in this section, while the propagation ratio for ODVV and VV rises to a high value, the ratio for VV remains low.

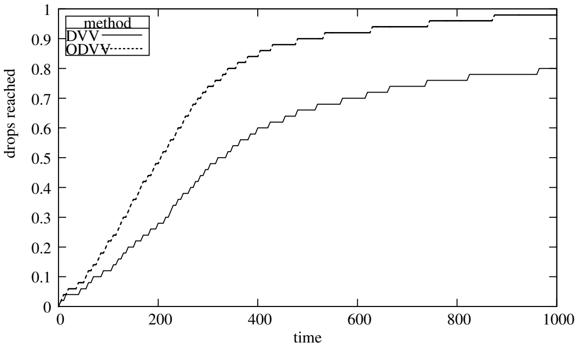

Fig. 12. Data propagation to dead drops.

Finally, Figure 12 summarizes the rates at which data disseminate to the dead drops (not vehicles) with the DVV and ODVV methods. The plot is based on the same set of experiments as was used for Figures 10 and 11. Even in instances that do not differ markedly in the rates at which data disseminate to vehicles, the ODVV scheme provides more effective dissemination to the dead drops. This property is useful because having data at a larger number of dead drops provides more resiliency to widely varying traffic densities.

IV. RELATED WORK

The IVC reference model by Hasegawa et al. [3] provides a good description of the broad range of IVC methods noted in Section I, along with a taxonomy.

Although our methods use no centrally connected infrastructure, the dead drops are, strictly speaking, a form of infrastructure. In contrast, Little and Agarwal [4] have presented a truly infrastructure-less method for propagating data in vehicular networks. However, their method requires a moderate or higher traffic density. In FleetNet [5], clusters of vehicles are connected using gateways that connect to a central network.

Michael [6] has proposed an adaptive layered data structure for reducing bandwidth use in inter-vehicle communications. As data travel farther from the source, higher-resolution layers of this data structure are removed, retaining the lower resolution information at much lower bandwidth. The method proposed by Ghosh et al. [7] has similar motivations and is based on probabilistic message delivery: The probability of message delivery drops as the distance from its origin increases.

Information dissemination in self-organizing networks has been studied by Wischhof, Ebner, and Rohing [8], [9]. They present a data abstraction and dissemination model based on segments, and describe a prototype implementation that uses 802.11 hardware.

As noted in Section I, our focus is on data dissemination, independent of the lower level communication protocols. However, the vehicular environment presents interesting challenges for lower level protocols. Fujimura and Hasegawa [10] have developed MAC protocols specifically for the vehicular environment. In contrast, Goel, Imielinski, and Ozbay [11] report good results using 802.11 networking methods in a vehicular environment. In general, such efforts are complementary to our work because our methods do not make any assumptions regarding the lower level protocols. Nevertheless, it should be interesting to investigate whether performance can be improved by taking advantage of some of the specific characteristics of the lower level protocols in higher level data-dissemination methods.

We have earlier distinguished between methods for dissemination and those for general-purpose networking, including routing protocols. Routing in IVC networks has been extensively studied, but for moderate or higher traffic densities [12], [13], [14]. The idea of using transit vehicles used as probes, as proposed by Cathey and Dailey [15], [16], may be adapted to data collection, as has been done in a sensor-network environment [17], [18]. That is, instead of dead drops, we may use vehicles that move solely to aid data dissemination.

V. CONCLUSION

We have studied the effectiveness of data-dissemination in a vehicular network when the density of equipped vehicles is very low. In these situations, a vehicle is rarely in communication range of other vehicles, making most prior inter-vehicle communication methods inapplicable. Our solution to this problem is based on the idea of using dead-drops, which are stationary communication nodes that exchange data with vehicles that pass by them. We presented a simulation study of the dynamic behavior of data propagation in a vehicular network augmented with such dead drops. For a range of realistic traffic parameters (for sparse equipped traffic), our results indicate that the presence of a modest number of dead drops provides a clear increase in the rate of data dissemination. Further, the schemes using dead drops are more resilient to traffic scenarios in which the set of active vehicles changes rapidly. Without dead drops, much information is lost when the small number of vehicles carrying it exit the system or move to remote parts. With dead drops, the information from such vehicles is likely to remain in the system as other vehicles pass the locations previously traversed by the exiting vehicles.

In continuing work, we are working on more dynamic schemes based on ideas similar to those used by DVV and ODVV. In particular, we are studying how such a system responds to commonly observed changes in traffic patterns, such as morning and evening peaks, and late night lows. To assist in such work, we are continuing the development of our interactive simulation and visualization system. We also hope to apply some of these ideas to pedestrian traffic, using hand-held wireless communication devices [19], [20].

REFERENCES

[1] S. S. Chawathe, "Inter-vehicle data dissemination in sparse equipped traffic," in *Proceedings of the 9th International IEEE Conference on Intelligent Transportation Systems (ITSC)*, Toronto, Canada, Sept. 2006, pp. 273–280.

[2] "MassGIS database," Office of Geographic and Environmental Information (MassGIS), Commonwealth of Massachusetts Executive Office of Environmental Affairs. http://www.mass.gov/mgis/, Dec. 2005.

[3] T. Hasegawa, K. Mizui, H. Fujii, and K. Seki, "A concept reference model for inter-vehicle communications (report 2)," in *Proceedings of the IEEE International Conference on Intelligent Transportation Systems (ITSC)*, Washington, D.C., Oct. 2004, pp. 810–815.

[4] T. D. C. Little and A. Agarwal, "An information propagation scheme for VANETs," in *Proceedings of the 8th International IEEE Conference on Intelligent Transportation Systems (ITSC)*, Vienna, Austria, Sept. 2005, pp. 155–160.

[5] K. Kutzner, J.-J. Tchouto, M. Bechler, L. Wolf, B. Bochow, and T. Luckenbach, "Connecting vehicle scatternets by Internet-connected gateways," in *Proceedings of Multiradio Multimedia Communications (MMC)—Communication Technology for Vehicles*, Dortmund, Germany, Feb. 2003.

[6] L. B. Michael, "Adaptive layered data structure for inter-vehicle communciation in ad-hoc communcation networks," in *Proceedings of the 8th World Congress on Intelligent Transportation Systems*, Sydney, Australia, Oct. 2001.

[7] A. Ghosh, D. Greene, Q. Huang, and J. Liu, "Variable-resolution information dissemination," in *Proceedings of the IEEE Conference on Sensor and Ad Hoc Communications and Networks*, Santa Clara, California, Sept. 2005.

[8] L. Wischhof, A. Ebner, and H. Rohling, "Information dissemination in self-organizing intervehicle networks," *IEEE Transactions on Intelligent Transportation Systems*, vol. 6, no. 1, pp. 90–101, Mar. 2005.

[9] ——, "Self-organizing traffic information system based on car-to-car communication: Prototype implementation," in *Proceedings of the 1st International Workshop on Intelligent Transportation*, Hamburg, Germany, Mar. 2004, pp. 49–53.

[10] K. Fujimura and T. Hasegawa, "A collaborative MAC protocol for inter-vehicle and road to vehicle communications," in *Proceedings of the 7th IEEE International Conference on Intelligent Transportaion Systems (ITSC)*, Washington, D.C., Oct. 2004, pp. 816–821.

[11] S. Goel, T. Imielinski, and K. Ozbay, "Ascertaining viability of WiFi based vehicle-to-vehicle network for traffic information dissemination," in *Proceedings of the 7th IEEE International Conference on Intelligent Transportaion Systems (ITSC)*, Washington, D.C., Oct. 2004, pp. 1086–1091.

[12] H. Füßler, M. Mauve, H. Hartenstein, M. Käsemann, and D. Vollmer, "Location-based routing for vehicular ad-hoc networks," *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 7, no. 1, pp. 47–49, Jan. 2003.

[13] J. Tian and L. Coletti, "Routing approach in CarTALK 2000 project," in *Proceedings of the 12th IST Mobile and Wireless Communications Summit*, Aveiro, Portugal, June 2003.

[14] J. Tian, K. Rothermel, and L. Han, "Spatially aware packet routing for mobile ad hoc inter-vehicle radio networks," in *Proceedings of the 6th IEEE International Conference on Intelligent Transportaion Systems (ITSC)*, Shanghai, China, Oct. 2003.

[15] F. W. Cathey and D. J. Dailey, "Implementation of traffic management using transit probes," in *Proceedings of the IEEE 8th International Conference on Intelligent Transportation Systems*, Vienna, Austria, Sept. 2005, pp. 13–16.

[16] ——, "Estimating corridor travel time by using transit vehicles as probes," *Transportation Research Record*, pp. 60–65, 2003.

[17] R. C. Shah, S. Roy, S. Jain, and W. Brunette, "Data MULEs: Modeling a three-tier architecture for sparse sensor networks," in *Proceedings of the First International Workshop on Sensor Network Protocols and Applications*, Anchorage, Alaska, May 2003, pp. 30–41.

[18] W. Zhao, M. H. Ammar, and E. W. Zegura, "A message ferrying approach for data delivery in sparse mobile ad hoc networks," in *Proceedings of the 5th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, Roppongi Hills, Tokyo, Japan, May 2004, pp. 187–198.

[19] T. Manabe, S. Yamashita, and T. Hasegawa, "On the M-CubITS pedestrian navigation system," in *Proceedings of the 9th International IEEE Conference on Intelligent Transportation Systems (ITSC)*, Toronto, Canada, Sept. 2006, pp. 793–798.

[20] S. Yamashita and T. Hasegawa, "On the M-CubITS pedestrian navigation system using textured paving blocks and its experiments," in *Proceedings of the 8th International IEEE Conference on Intelligent Transportation Systems (ITSC)*, Vienna, Austria, Sept. 2005, pp. 121–125.