Exploiting Emergent Behavior for Inter-Vehicle Communication

David Choffnes and Fabián E. Bustamante
Department of EECS, Northwestern University
{drchoffnes,fabianb}@cs.northwestern.edu

Abstract

We introduce *Virtual Ferry Networking* (VFN), a novel approach to data dissemination services on mobile ad-hoc networks. VFN exploits the emergent patterns of vehicles' mobility to buffer and carry messages when immediately forwarding those messages would fail. Instead of depending on a fixed, small set of vehicles and paths for ferrying messages, VFN allows any vehicle moving along part of a virtual route to become a possible carrier for messages. VFN helps address many of the challenges with supporting distributed applications in challenging ad-hoc vehicular networks with rapidly changing topologies, fast-moving vehicles and signal-weakening obstructions such as bridges and buildings. We discuss the challenges with implementing VFN and present evaluation results from an early prototype.

1 Introduction

Mobile Ad Hoc Networks (MANETs) provide infrastructureless, rapidly deployable, self-configurable network connectivity that can be exploited by diverse application areas ranging from disaster relief to wildlife tracking to city traffic advisory. Recently, technological advances have fueled new research on inter-vehicle communication protocols and systems for Vehicular Ad Hoc Networks (VANETs). VANETs offer a clear potential for impacting society through safer, more efficient and environmentally friendly vehicular transportation. Millions of "equipped" vehicles could provide an easily upgradeable, naturally scalable, and widely distributed network for general applications.

This paper addresses the problem of reliable communication in densely populated VANETs. Reliable message routing in MANETs is challenging not only because the network topology is constantly changing, but also because nodes must use a broadcast medium to probe the topology and send data. The result is that much bandwidth is wasted due to flooding of topology discovery messages and the resulting collisions due to radio interference from simultaneous message transmissions. This problem is significantly worse in VANETs, where vehicles tend to travel at high speed in large groups, often separated by signal-weakening obstructions such as buildings.

In this paper we introduce **Virtual Ferry Networking** (VFN), a mobility-assisted approach for enabling distributed systems in these challenging conditions. Specifically, VFN relies on the aggregated movements of independent vehicles to provide self-configuring virtual paths over which to ferry messages. By using multiple vehicles to ferry messages, VFN provides higher reliability and more flexibility than a single ferry vehicle or fixed ferry path [16]. Any vehicle traveling along part of this path can become a ferry for messages during that period.

2 VFN Overview

Despite their many challenges, VANETs present numerous opportunities for optimizing network communication. For one, vehicles allow generous limits on power consumption and size for system components, enabling always-on network communication and large caches for storing and forwarding messages. Further, as the number of GPS-equipped vehicles continues to grow, VANETs will be able to exploit programmed vehicle routes, and location and bearing information to improve communication efficiency and reliability. More importantly, collections of vehicles are known to favor particular paths based on the number of lanes in a road, traffic-signal behavior and proximity to highways, among other factors. VFN exploits these opportunities and emergent mobility patterns to provide a self-configuring robust backbone topology that enables reliable message delivery with low overhead.

To place our work in the context of a real application, consider the problem of unicasting a message between two vehicles in a VANET. We assume that every vehicle in the system is equipped with a GPS device for positioning, a wireless network interface that provides a unique identifier for the vehicle, and sufficient storage to buffer packets. We also assume that the network uses a geolocation system (e.g., GLS [8]) to map vehicles to locations, that there is sufficiently high density of vehicles to prevent network disconnections and that each vehicle runs our VFN software, which sits between the application layer and the routing layer in the traditional network stack. Finally, we assume that many other vehicles are transmitting messages, leading to contention for the medium when vehicular density is high.

We illustrate our example using a region of downtown Chicago (Fig. 1), where the source, located in the north-

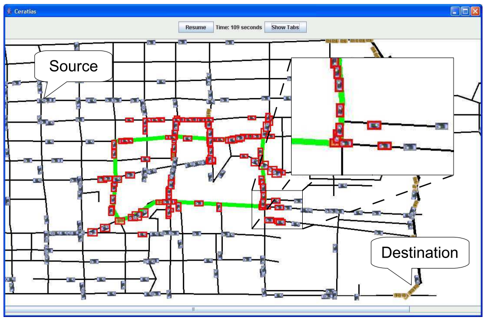

FIGURE 1: A simple VFN topology in downtown Chicago. The thicker, lighter shaded roads represent the virtual path over which messages are ferried, and the vehicles with dark borders are participating in the VFN.

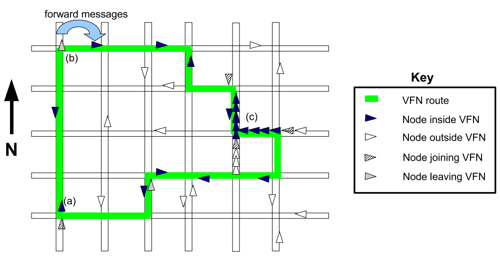

FIGURE 2: Illustration of a simple VFN topology.

west corner of the map, sends a unicast message to a destination, located in the southeast corner. At the source, the routing layer (e.g., GPSR [7] or any other ad-hoc routing protocol) determines the next hop for the message and transmits it. At each successive hop, the routing protocol determines the next hop and immediately forwards the message until it reaches a vehicle in the center of the map, where vehicle density and contention for the broadcast medium are high.

In response to the challenging conditions for broadcast communication, the VFN system has used knowledge of local mobility patterns to determine a virtual ferrying route (VFR), in terms of streets, over which messages can be reliably ferried. When the message enters this region, the VFN layer uses a participating vehicle to *ferry* the message by buffering and carrying it toward its destination. The vehicle carries the message along the VFR until the driver steers the vehicle away from the VFR, at which point the VFN layer broadcasts the carried message to another nearby vehicle still in the VFR.

The participating vehicles coordinate to carry and forward the message toward the area of its final destination. When the message leaves the VFR, it resumes a multi-hop path as determined by the routing layer, which uses broadcasting to forward the message to the destination without carrying it.

We now provide a more formal description of VFN. A VFN system consists of one or more sets of road segments over which vehicles follow an aggressive protocol aimed at improving network reliability, efficiency and connectivity. Each set of road segments defines a circuit that we refer to as a *virtual ferrying route*. All vehicles that are driving on (or near) a portion of a VFR are said to be participating in the VFN. Figure 2 presents a schematic view of VFN with a single VFR in a city grid. The participating vehicles (dark triangles) use message aggregation, awareness of other vehicles' paths, and adaptive message transmission rates to optimize communication by leveraging store-and-forward style communication. By carrying messages when immediately broadcasting them would fail, VFN improves message-delivery reliability and can increase available bandwidth by reducing the number of broadcasts necessary to relay data. VFN relies on the emergent mobility patterns of all vehicles in the network to form virtual routes over which to carry messages, rather than requiring participating vehicles to alter their routes to pick up and/or deliver messages [16, 10, 14, 12, 9]. We expect VFN to be able to operate under a variety of conditions, e.g., by ferrying messages to bridge partitions when the network is sparsely populated and by aggregating/carrying messages to reduce interference when vehicle density is high. Thus, our work is also related to research on partitioned networks that rely on mobile nodes to deliver information opportunistically, for example, with nodes exchanging information when they meet [3, 13, 6, 15].

3 Challenges

We have identified several challenges that must be addressed for the realization of Virtual Ferry Networks, ranging from self-tuning issues in store-and-forward decisions for networks with varying vehicle density, to approaches for self-configuration of dynamic virtual routes. We develop these issues in the following subsections.

3.1 VFN Route Design and Management

The choice of VFRs in VFN is critical to the performance of the network. Due to the lack of control over vehicular mobility patterns, VFN routes will evolve over time, e.g., as a function of the levels of congestion along a path and the routes that vehicles follow.

One approach is for each vehicle to store and maintain a global set of virtual ferrying routes. In a real-world implementation, management and dissemination of VFRs can become a significant issue depending on how often and to what extent routes must change over time. For example, we have found that vehicles can require up to hundreds of KB of data to store lists of VFRs for a large region such as the greater Chicago area. Although this is a modest amount of data, it would incur a great deal of message control overhead if routes change quickly and the entire set of VFRs must be propagated after each route update. Another problem is that not all vehicles will receive the latest VFRs at the same time, allowing the possibility for multiple vehicles to have different views of the VFN region. Thus, one aspect of our work must focus on how to best represent, manage and maintain VFRs in VFN. We are experimenting with approaches based on incremental VFR updates, to reduce control overhead, and expiring leases on existing VFRs, to force vehicles to retrieve consistent views of the VFN region. Leases for VFRs can have vastly different values, such as no expiration for major arteries and fast expirations for roads with transient congestion. Clearly, one must strike a balance between low control message overhead and consistent network views.

Most interestingly, one could employ pheromone-based [4] or collaborative reinforcement learning [5] approaches, for self-configuration of VFRs. In this case, vehicles passively acquire local vehicular-density information and exchange their aggregated views with other vehicles. VFR routes are thus defined in a distributed manner using reinforcement learning. An important question is whether participating vehicles can ever reach a consensus view of VFRs using this approach. A prototype implementations of these ideas leads us to believe that views of VFRs indeed converge in many cases; however, validating the extent to which this occurs in general remains part of our ongoing work.

3.2 Storing and Forwarding

VFN must address the issue of when a vehicle should carry a message and when it should transmit it. The need to store messages for later delivery is clear for the extreme cases of sparse and dense networks. In sparse networks, the average number of vehicles within wireless transmission range of a vehicle is less than one, so many vehicles will be unreachable via multihop wireless routing. In this case, VFN could provide connectivity across partitions in the network by exploiting mobility awareness and storing messages for later delivery. In dense regions of the network, interference due to high rates of message transmission can reduce goodput to zero. Here, a VFN system could temporarily suspend message transmission to carry messages away from a local hot-spot, thereby improving the chance for high goodput. Without this store-and-forward capability, messages passing through such a region would likely be dropped.

In general, VFN must address the issue of how much bandwidth a system consumes by propagating copies of messages. This entails finding the appropriate balance between reliability and bandwidth consumption. We are investigating adaptive algorithms that enable self-tuning of these factors according to network conditions observed by listening on the network interface in permiscuous mode. For example, VFN should buffer messages more aggressively when contention for the medium is high and forward messages opportunistically when contention is low.

Given the exponential growth of storage capacities and the decreasing cost per unit storage, we do not expect the number of messages cached at each vehicle to be an issue. Regardless, for media with large, nonuniform access speeds, it is important to use higher-speed caches and reduce fragmentation of data on the storage medium. To that end, we plan to investigate storage organization and caching policies when the storage load is high.

3.3 Mobility Awareness

An important issue in VFN design is determining how to react when a vehicle joins or leaves the VFN region. For example, when a vehicle is about to leave a VFR, it must decide whether to forward its stored messages to vehicles that will continue in the VFN region for some time. Additionally, it is important to forward stored messages to a vehicle that is not moving away from the messages' destinations. A participating vehicle could use its neighbors' speeds and bearings to inform these decisions.

Message-forwarding reliability is another important issue for VFN. If a group of vehicles is moving toward the same destination, they could cooperate (e.g., by electing a primary message carrier) to ensure reliable message propagation and efficient interactions with other vehicles. If a vehicle unexpectedly leaves the network or is otherwise unable to keep its messages inside the VFN region, an alternate vehicle that has also cached the message can take the post. In fact, as a participating vehicle encounters new ones, it can propagate copies of its messages to those vehicles. Each vehicle in the VFN region can determine which other ones should receive copies of messages based on knowledge of its own route and information about other vehicles' mobility. A VFN system can also modify its behavior to react to changes in a message carrier's mobility. It could detect that a participating vehicle cannot continue along its path at its normal speed (e.g., due to a traffic jam) and take measures to pass the message, e.g., in a multihop manner along the route or to a nearby vehicle moving in a different direction that might facilitate message delivery.

3.4 Other Issues

This is clearly only a limited set of the challenges we face. Another interesting topic that we are investigating is the use of message aggregation for efficient data delivery and the self-tuning of message transmission rates to deal with varying network density. Much like a sensor network can use message aggregation to reduce wireless transmission, VFN vehicles can use aggregation to batch messages, reducing the frequency and increasing the efficiency of channel access. Another challenge is to investigate how a VFN system can adjust the ratio between ferried and broadcasted messages to provide multiple classes of priorities. For example, VFN can be used to provide "quiet zone" in which any high priority message (e.g., safety messages) can pass through the region with little to no resistance in terms of medium contention.

4 Early Evaluation

This section presents results from an early VFN prototype that validate the feasibility of some of the fundamental features of VFN. We first explain our evaluation methodology, then we use our prototype to evaluate some of the advantages of the VFN approach. Lastly, we demonstrate the flexibility of VFN by discussing how we used our prototype system to provide a reliable geostationary message service.

For our simulation-based experiments we rely on SWANS (Scalable Wireless Ad Hoc Network Simulator) [1], a publicly-available, scalable wireless network simulator. SWANS runs atop JiST (Java in Simulation Time), a high-performance discrete event simulation engine that features low memory consumption and fast run times. It supports large numbers of nodes (>1,000,000) and defines an extensible set of highly configurable simulation abstractions to model numerous real-world components for various levels of realism in simulation. In our simulations, we use the STRAW (STreet RAndom Waypoint) mobility model [2], which restricts node movement to streets defined by map data and limits their mobility according to vehicular congestion and simplified traffic control mechanisms. For the experiments that follow, we use STRAW in the region of downtown Chicago, approximately 4000 x 1250 meters, shown in Fig. 1.

4.1 Contention Reduction

We now address the question of whether a VFN system can improve message-transmission reliability by increasing available bandwidth and reducing collisions due to contention for the broadcast medium. The simulation experiment contained 300 vehicles moving according to the STRAW mobility model; namely, by choosing random pairs of origins and destinations and taking the shortest-time path between them. Of those 300 vehicles, 10 transmitted constant bitrate (CBR) UDP traffic using the GPSR routing protocol [7] and 802.11 MAC layer.[1] We then examined the number of dropped packets due to interference when our prototype VFN system was enabled/disabled.

Fig. 3 illustrates our results. The figure clearly shows that for the region 0–300 s, where VFN is *enabled*, few packets are dropped due to interference. For the period between 300 and 400 s, we disable VFN so that messages are routed immediately via GPSR. The resulting increase

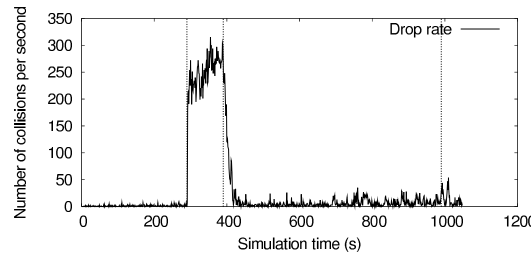

FIGURE 3: Graph depicting the number of packets dropped due to collision when VFN is enabled and temporarily disabled.

in congestion is high and immediate, *strongly suggesting that VFN indeed improves message-transmission reliability by reducing congestion when network contention is high.* When VFN is re-enabled at time 400 s, congestion does not disappear as quickly as it arrived; however, within half a minute, congestion returns to its original levels with VFN enabled. Thus, we additionally see that as dynamic VFN routes change, the system can react to improve performance over relatively small time scales.

4.2 Bridging the Connectivity Gap

Another key capability for VFN is that the network can use mobility patterns to ferry messages across regions where delivery is otherwise impossible. To determine whether a VFN system is successfully ferrying messages, we conducted a pair of experiments: one that uses our early prototype VFN system and one that does not. Everything else about the experiments, including the mobility patterns and message transmission patterns, is fixed. Our goal here was to measure the number of messages that were *dropped* by the system *without* VFN, but were delivered when VFN was active. Our simulation setup was the same as in the previous section except that VFN was either enabled or disabled for the *entire* simulation run. Our results are promising: *45% of the messages dropped by GPSR without VFN were delivered when VFN was active.* Thus, even a simple VFN implementation can exploit emergent vehicular mobility to substantially improve message delivery.

It is important to note that our prototype VFN system is extremely bandwidth-conserving: it will not proactively attempt to forward messages even when bandwidth is available. Thus, the ferries bridge the connectivity gap in the network at the cost of high latency. Although this is acceptable for delay-tolerant networks, it is not acceptable for latency-sensitive applications. We are currently developing new VFN functionality that exploits the available bandwidth from message ferrying to reduce latency in delay-sensitive applications.

4.3 Enabling New Services

Beyond serving as a network core that enhances the reliability of message transmission by reducing contention

[1] For radio transmissions we employ a generic path loss model with exponent 2.8 and shadowing standard deviation 6.0.

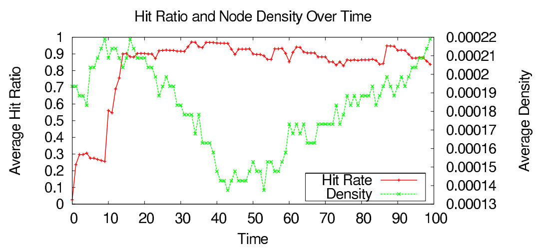

FIGURE 4: Instantaneous hit ratio and vehicle density during a message lifetime.

and helping bridge connectivity gaps, VFN is useful to a variety of applications such as publish/subscribe and geostationary messaging services [11]. To explore this path we implemented a version of geostationary messaging using VFN and evaluated its performance. A geostationary message is essentially a data item that persists in a particular geographic region for a specified period of time. It can naturally support a number of interesting applications, from advertisement to travel-time notifications, safety warnings (e.g., slippery road or crashed vehicle ahead). To enable geostationary messaging with VFN, the corresponding VFR is set to the border of a message's target region, and messages are carried by all vehicles inside and near the region contained in the VFR.

We ran our experiments for periods of 90,000 (simulated) seconds during which time each new message is seeded at a random location with a duration of 100 seconds and dispersion radius of 250 m. Messages were evenly spaced at 100-second intervals, meaning that at most one message was active at any given time and there were 900 messages observed. This large number of data points allows us to examine the effect of density on performance. An interesting metric to study is *hit ratio* — the number of vehicles carrying a message in its target region divided by the total number of vehicles in the same region (i.e. values ≈ 1 are best). Hit ratio measures how successfully our protocol disseminates its messages.

Due to space limitations, Fig. 4 shows the lifetime of only a single message under the microscope. The x-axis represents time during the message lifetime (at 1-second intervals) and the y-axes represent the hit ratio and vehicle density at each second. The figure clearly shows that, after the message has been introduced by a vehicle, there is a brief warmup period during which the message saturates its target region and after which a fairly steady hit ratio is maintained. Note that the hit ratio is fairly stable (90th percentile) even under changing vehicle density, thanks to the VFN system caching the message at vehicles surrounding the target region.

5 Conclusion

We proposed a novel approach to VANET services called Virtual Ferry Networking (VFN), in which groups of vehicles collectively define a ferry network that relies on the aggregation of portions of independent vehicle movements along a certain route. We discussed how VFN exploits emergent vehicular mobility patterns to enable self-configuring store-and-forward routes for reliable message delivery and presented early results that demonstrate the potential VFN to improve application performance. Finally, we discussed numerous remaining important challenges that we must address to bring the VFN approach to fruition.

References

[1] BARR, R. *An efficient, unifying approach to simulation using virtual machines*. PhD thesis, Cornell University, 2004.

[2] CHOFFNES, D. R., AND BUSTAMANTE, F. E. An integrated mobility and traffic model for vehicular wireless networks. In *Proc. of ACM VANET* (2005).

[3] DAVIS, J., FAGG, A., AND LEVINE, B. Wearable computers as packet transport mechanisms in highly-partitioned ad-hoc networks. In *Proc. of International Symposium on Wearable Computing* (2001).

[4] DORIGO, M., AND CARO, G. D. *New Ideas Optimization*. McGraw-Hill, 1999, ch. The ant colony optimization metaheuristic.

[5] DOWLING, J., CURRAN, E., CUNNINGHAM, R., AND CAHILL, V. Using feedback in collaborative reinforcement learning to adaptively optimize MANET routing. In *IEEE Transactions on Systems, Man and Cybernetics* (May 2005), vol. 35.

[6] JUANG, P., OKI, H., WANG, Y., MARTONOSI, M., PEH, L.-S., AND RUBENSTEIN, D. Energy-efficient computing for wildlife tracking: design tradeoffs and early experiences with ZebraNet. In *Proc. of ASPLOS* (2002).

[7] KARP, B., AND KUNG, H. T. GPSR: greedy perimeter stateless routing for wireless networks. In *Proc. of ACM/IEEE MobiCom* (2000).

[8] LI, J., JANNOTTI, J., COUTO, D. S. J. D., KARGER, D. R., AND MORRIS, R. A scalable location service for geographic ad hoc routing. In *Proc. of ACM/IEEE MobiCom* (2000).

[9] LI, Q., AND RUS, D. Sending messages to mobile users in disconnected ad-hoc wirless networks. In *Proc. of ACM/IEEE MobiCom* (2000).

[10] LI, Q., AND RUS, D. Communication in disconnected ad hoc networks using message relay. *J. Parallel Distrib. Comput. 63*, 1 (2003), 75–86.

[11] MAINHÖFER, C., LEINMÜLLER, T., AND SCHOCH, E. Abiding geocast: Time-stable geocast for ad hoc networks. In *Proc. of ACM VANET* (2005).

[12] R. C. SHAH, S. ROY, S. J., AND BRUNETTE, W. Data MULEs: Modeling a three-tier architecture for sparse sensor networks. In *Proc. of SNPA Workshop* (2003).

[13] VAHDAT, A., AND BECKER, D. Epidemic routing for partially connected ad-hoc networks. Tech. rep., Duke University, 2000.

[14] WANG, W., SRINIVASAN, V., AND CHUA, K.-C. Using mobile relays to prolong the lifetime of wireless sensor networks. In *Proc. of ACM/IEEE MobiCom* (2005).

[15] WU, H., FUJIMOTO, R., GUENSLER, R., AND HUNTER, M. MDDV: a mobility-centric data dissemination algorithm for vehicular networks. In *Proc. of ACM VANET* (2004).

[16] ZHAO, W., AMMAR, M., AND ZEGURA, E. A message ferrying approach for data delivery in sparse mobile ad hoc networks. In *Proc. of MobiHoc* (2004).