Poster: Accurate Data Aggregation for VANETs

Khaled Ibrahim and Michele C. Weigle
Department of Computer Science,Old Dominion University
Norfolk, VA 23529-0162, USA
{ibrah_k, mweigle}@cs.odu.edu

ABSTRACT

Data aggregation is an important issue for vehicular ad-hoc networks (VANETs). Congestion notification applications are built to warn drivers of traffic slowdowns far enough in advance that the drivers may take alternate routes. Data that is broadcast should be self-contained and fit into a single MAC-layer frame. With dense traffic, aggregation is needed to represent a large number of vehicles in relatively small frame. We present a new technique for aggregating vehicles' data without losing accuracy. Vehicles build a local view based on speed and position reports from neighboring vehicles. This local view, representing vehicles up to 1.6 km ahead, is then aggregated into a single frame and broadcast. Vehicles use received aggregated frames to extend their views even farther.

Categories and Subject Descriptors

C.2.0 [**Computer-Communication Networks**]: General; E.4 [**Coding and Information Theory**]: Data compaction and compression

General Terms

Algorithms, Design

Keywords

Aggregation, Data Compaction, Vehicular Networks

1. BASIC OPERATION

Each vehicle in the system is equipped with GPS, a navigation system, and a wireless communication device with a 300 m range using Dedicated Short Range Communications (DSRC) [3]. Each vehicle is also pre-assigned a public/private key pair and the public key's certificate, used for authentication. Figure 1 is a diagram of how traffic is represented in our system, which we call the vehicle's *view*. The view in front of a vehicle is divided into *clusters*. Each cluster has a width of 4 m, the same as that of a typical highway lane, and a length of 64 m, set for maximal record compression.

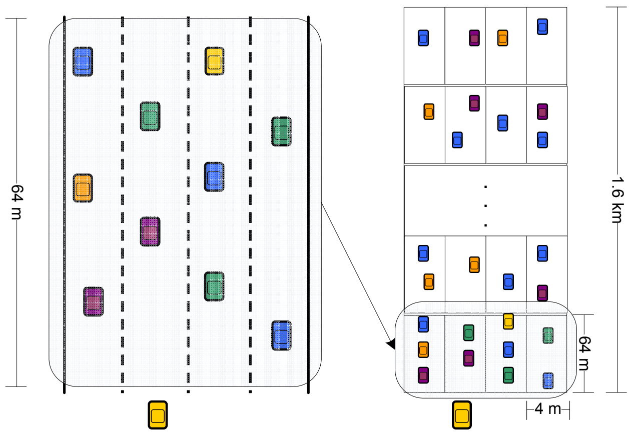

Figure 1: Vehicle's Local View, Divided into 4 m x 64 m Clusters

Vehicles exchange their location and speed by periodically broadcasting *primary data records*. These records are grouped into their corresponding clusters and used to form the vehicle's *local view*. There are 26 "rows" of clusters in a local view, resulting in a visibility of 1.6 km, or 1 mile. Each cluster in the local view is periodically aggregated into an *aggregated cluster record*. All of the aggregated cluster records comprising a vehicle's local view are put into a single frame and broadcast to other vehicles. Upon receiving an aggregated frame, a vehicle can build an *extended view*, which may contain vehicles not present in its local view.

2. LOCAL VIEW

The local view is constructed from primary data records received from neighboring vehicles. The primary data record contains the basic information for a single vehicle *i.e.* latitude, longitude, speed, and timestamp. A vehicle broadcasts a *primary frame* containing its primary data record every 300 ms. This frame consists of a one-bit type field, the position of the sender (broadcaster) of the frame, the primary data record, the vehicle's signature, and a certificate for authentication. In addition to recording the primary data record for a vehicle, a receiver will also store the vehicle's public key with the record as an identifier.

VANET'07, September 10, 2007, Montréal, Québec, Canada.
ACM 978-1-59593-739-1/07/0009.

2.1 Dissemination

Since the transmission range with DSRC is much less than the distance covered in the local view (1.6 km), vehicles will need to forward primary frames in order to propagate the messages backwards farther. Primary frames are not propagated more than 1.6 km behind the original sender. In order to limit the number of broadcast messages used to propagate the frames, we employ the *Smart Broadcast* mechanism proposed by Fasolo *et al.* [1].

2.2 Aging

Since a timestamp is recorded along with the primary data record, old records can be purged from the local view. A record is removed once the vehicle has physically passed the vehicle described by the record.

2.3 Security

Because each primary frame contains a signature and certificate, the local view construction is immune to many security attacks except the false record injection attack.

3. EXTENDED VIEW

Each vehicle can build its local view based on primary records received from other vehicles. In order to extend the view farther, vehicles exchange aggregated records.

3.1 Clustering and Compression

As primary records are received, the vehicles described in those records are grouped into their corresponding clusters, based on their distance from the receiving vehicle. If a received primary record does not belong in one of the clusters comprising the current local view, it is discarded.

The *compact data record* is used to represent a single vehicle within a cluster. The compression is achieved by representing only the differences between the vehicle data and overall cluster data. Before clustering and aggregation, the median speed of all vehicles in a cluster is calculated, and the position of each vehicle is translated into $\{X, Y\}$ coordinates with the aggregating vehicle as the origin.

3.2 Aggregation

Before aggregation, the vehicles in each cluster are described in a compact data record. From this, we form an *aggregated cluster record*, which is a concatenation of the compact data records of the cluster's vehicles.

3.3 Dissemination of Aggregated Records

Once the compression and aggregation has been performed, the cluster aggregated records from the current view are concatenated into a single frame and sent via broadcast. The aggregated frame includes a one-bit packet type (to distinguish it from a primary frame), a timestamp, the position of the sender (broadcaster), the aggregating vehicle's position, and a set of cluster aggregated records for the current view.

Aggregated frames are disseminated in the same manner as primary frames (*i.e.*, using the Fasolo *et al.* contention window adjustment algorithm [1]).

3.4 Building the Extended View

When a vehicle receives an aggregated frame, it first checks to see if the aggregating vehicle's position is within its current view. If so, then there must be an overlap between the receiving vehicle's view and the view contained in the aggregated frame. The receiving vehicle reconstructs the primary data for the vehicles in the aggregated frame. Before placing the vehicles into the view, their positions are adjusted based on the speed they were traveling and the time since the aggregated frame was broadcast. Once adjusted, the receiving vehicle compares the new vehicle data in the intersecting area with the data already in its current view. A number of vehicles will overlap (*i.e.*, be very close together, if not in the same space, and traveling at similar speeds). If over 75% of the vehicles in the intersecting region overlap, then the received view is declared to be consistent with the current view, so the non-intersecting part in the received view can be used to extend the current view. Once the current view has been extended, this view may be extended farther through the receipt of additional aggregated frames.

The comparison between the vehicle's data in the intersecting area can be used for detecting malicious vehicles that try to inject false views in the traffic. Previous approaches to solve this problem [2] would require more overhead and processing. In our approach, these vehicles can be caught at no extra cost, because the comparison of the intersecting views is an essential part of extending the view.

4. CONCLUSION AND FUTURE WORK

We have presented a technique for accurate aggregation of vehicle location and speed data. The local view presents data gathered from primary records, which are sent in signed frames containing a vehicle's position and speed. The local view is grouped into clusters, which are then used to compact and aggregate the local view data. Aggregated data from other vehicles can be used to extend a vehicle's view past its 1.6 km local view. In the best case, this extended view could be as far as 104 km. Since vehicles' positions and speeds are represented as differences from the cluster data rather than combined with other vehicles' data, the accuracy of the aggregated data in our system is very high.

For future work, we plan to investigate how well the *Smart Broadcast* technique [1] works to reduce the number of messages broadcast and to further investigate optimum broadcast intervals and message lifetimes for both primary and aggregated frames. In addition, we plan to develop algorithms for increasing security using the data in the extended frames to detect and isolate vehicles that lie about their position or speed or that intentionally mis-aggregate data.

5. REFERENCES

[1] E. Fasolo, R. Furiato, and A. Zanella. Smart Broadcast Algorithm for Inter-Vehicular Communications. In *Proceedings of the Wireless Personal Multimedia Communications Symposium (WPMC)*, 2005.

[2] P. Golle, D. Greene, and J. Staddon. Detecting and Correcting Malicious Data in VANETs. In *Proceedings of the ACM Workshop on Vehicular Ad Hoc Networks (VANET)*, pages 29–37, 2004.

[3] US Department of Transportation. Standard Specification for Telecommunications and Information Exchange Between Roadside and Vehicle Systems - 5 GHz Band Dedicated Short Range Communications (DSRC) Medium Access Control (MAC) and Physical Layer (PHY) Specifications. ASTM E2213-03, Aug. 2003.