

CASCADE: Cluster-based Accurate Syntactic Compression of Aggregated Data in VANETs

Abstract—We present a method for accurate aggregation of highway traffic information in vehicular ad hoc networks (VANETs). Highway congestion notification applications need to disseminate information about traffic conditions to distant vehicles. In dense traffic, aggregation is needed to allow a single frame to carry information about a large number of vehicles. Our technique, CASCADE (Cluster-based Accurate Syntactic Compression of Aggregated Data in VANETs), uses compression to provide aggregation without losing accuracy. As a part of CASCADE, we present probabilistic Inter-Vehicle Geocast, allowing the re-broadcasting of frames to depend upon the surrounding vehicle density. We show that the frame reception rate is improved with this modification. CASCADE provides each vehicle with data that is highly accurate, represents a large area in front of the vehicle, and can be combined with aggregated data from other vehicles to further extend the covered area.

I. INTRODUCTION

Highway traffic congestion is costly. Estimates have shown that millions of hours and billions of dollars are wasted each year because of congested roadways [19]. Advance notification of traffic congestion would allow many drivers to take alternate routes and save time. Vehicular Ad-hoc Networks (VANETs) have been proposed as a means to provide this advance notification of traffic congestion to drivers. A VANET consists of nearby vehicles exchanging information with each other via wireless broadcast. Using the wireless channel efficiently is a challenging problem. Most VANET messages (e.g., speed and location updates) are periodically broadcast by each vehicle. To provide information to non-neighboring vehicles, the messages must be forwarded to vehicles outside the original sender’s broadcast range. The more vehicles participating in the VANET, the larger the number of messages sent, and the higher the probability of wireless collisions. In order to reduce the number of messages that need to be sent, several data aggregation techniques have been proposed [13], [14], [16]. Unfortunately, with these techniques, some accuracy of the data is lost upon aggregation, and data aggregated by one vehicle cannot be combined with data aggregated by another.

In this paper, we present CASCADE (*Cluster-based Accurate Syntactic Compression of Aggregated Data in VANETs*), a new method for accurate aggregation of traffic information in VANETs. The main idea behind

CASCADE is cluster-based compression. In aggregated frames, we represent each vehicle’s location based on its difference from the location of the center of the cluster and its speed based on its difference from the median speed of all vehicles in the cluster. In this way, accurate information can be distributed in a small number of bytes. We also present a modification to Inter-Vehicle Geocast [2], which we call *probabilistic Inter-Vehicle Geocast* (p-IVG), that adapts the re-broadcasting of frames based on the surrounding traffic density.

We show that CASCADE makes efficient use of the wireless channel while providing vehicles with data that is highly accurate, represents a large area in front of the vehicle, and can be combined with aggregated data from other vehicles to further extend the covered area.

II. RELATED WORK

Data aggregation has received much attention in the wireless sensor network community [6], [12], [15], [23], [24], but many of the approaches either assume a static network or require several rounds of communication between nodes to provide security. Both of these requirements are impractical for VANETs.

There has been recent work on data aggregation techniques specifically designed for VANETs. Picconi *et al.* [14] classified aggregation techniques as either *syntactic* or *semantic*. Syntactic aggregation uses a technique to compress or encode the data from multiple vehicles in order to fit the data into a single frame. This results in lower overhead than sending each message individually. In semantic aggregation, the data from individual vehicles is summarized. For instance, instead of reporting the exact position of five vehicles, only the fact that five vehicles exist is reported. The trade-off is a much smaller message in exchange for a loss of precise data.

Nadeem *et al.* [13] present the *TrafficView* system, which uses semantic aggregation. The authors present two techniques for aggregation: *ratio-based* and *cost-based*. In the ratio-based technique, the roadway in front of a vehicle is divided into regions. Data is aggregated based on ratios that have been pre-assigned to each region. Regions farther away from a vehicle are assigned larger aggregation ratios, because precise detail may not be needed over a long range. The resulting view of traffic conditions is, thus, customized for each particular vehicle. For this reason, the produced view may not be useful

for other vehicles unless they use the same aggregation ratios. In the cost-based aggregation technique, data is aggregated based on a cost function that depends on the position of the aggregating vehicle. For this reason, the produced view of the traffic is not useful to any other vehicle unless it is close to the aggregating vehicle.

Lochert *et al.* [11] present a probabilistic technique for aggregating the disseminated data in VANET applications. The proposed technique does not aggregate the actual values but uses a modified Flajolet-Martin sketch as a probabilistic approximation for the values. This technique can be applied to aggregate the data in any non-accuracy-sensitive application (*e.g.*, estimating the number of available parking spaces), but it cannot be used in our target application, which requires the actual vehicle information (position, speed, etc.) to be disseminated and re-aggregated to reach distant vehicles.

One of the issues in aggregating vehicle data is security. To save bandwidth, original digital signatures are typically not included in aggregated messages. The problem is how to prevent a malicious aggregator from injecting false information. There have been a few recent approaches to providing secure aggregation, but the approaches either require the use of a tamper-proof service in the vehicle [14] or use semantic aggregation that removes specific information about each vehicle [16].

III. CASCADE

The goal of CASCADE is to allow a vehicle to obtain an accurate view of upcoming traffic conditions. CASCADE can also support cooperative collision warning applications [4], and so, vehicles exchange their position information with neighboring vehicles several times a second.

A. Assumptions

We assume that each vehicle in the system is equipped with a Global Positioning System (GPS) receiver for obtaining location and time, a navigation system that can map GPS coordinates to a particular roadway and offer routes, and a communications device using Dedicated Short Range Communications (DSRC) [20]. DSRC, with a transmission rate of 6-11 Mbps, is based on the upcoming IEEE 802.11p standard [8], which is a part of the larger IEEE 1609 Family of Standards for Wireless Access in Vehicular Environments (WAVE) [21]. Each vehicle is also pre-assigned a public/private key pair and the public key's certificate, used for authentication. To address privacy concerns, each vehicle may also use multiple pseudonyms to disguise its public keys [7], [17]. In our design of CASCADE, we assume a four-lane highway and ignore vehicles traveling in the opposite direction.

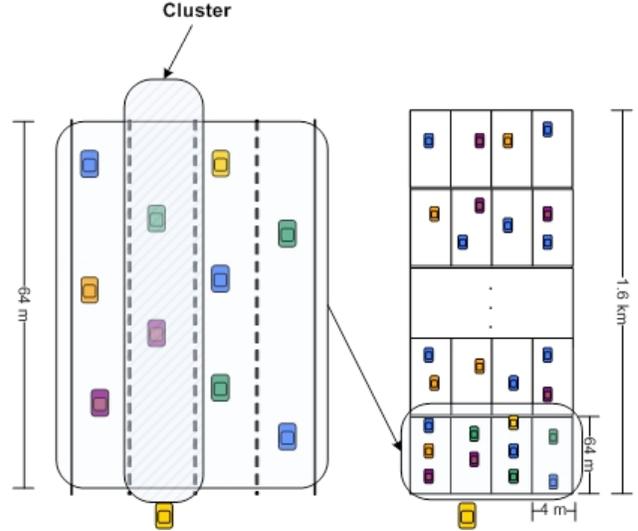


Fig. 1. Vehicle's Local View, Divided into 4 m x 64 m Clusters

B. Overview

Before describing the details of the system, we present a brief high-level overview. Each vehicle periodically broadcasts its position information (including location, speed, acceleration, and heading), which we call a *primary record*. Received primary records are stored in a local database in each vehicle. Those primary records representing vehicles ahead of the current vehicle comprise the *local view*. The local view, as shown in Figure 1, 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 (to be described in Section III-D.1). There are 25 “rows” of clusters in a local view, resulting in a visibility of 1.6 km, or 1 mile. As the local view is longer than the typical DSRC transmission range (about 300 m), primary records may be re-broadcast (to a maximum of 1.6 km behind the original sender).

Each vehicle periodically compresses and aggregates the primary records in its local view into an *aggregated record*. This aggregated record is then broadcast to neighboring vehicles. Received aggregated records may be used to augment the local view by providing information about vehicles beyond the local view, resulting in an *extended view*. Figure 2 shows an example of the local view and extended view.

Although our examples feature a straight, rectangular-shaped road, CASCADE is not limited to such geometries. Figure 3 shows how a local view can be mapped to a curved roadway. As the vehicle enters the curve, more of the vehicles inside and past the curve will be added to its local view.

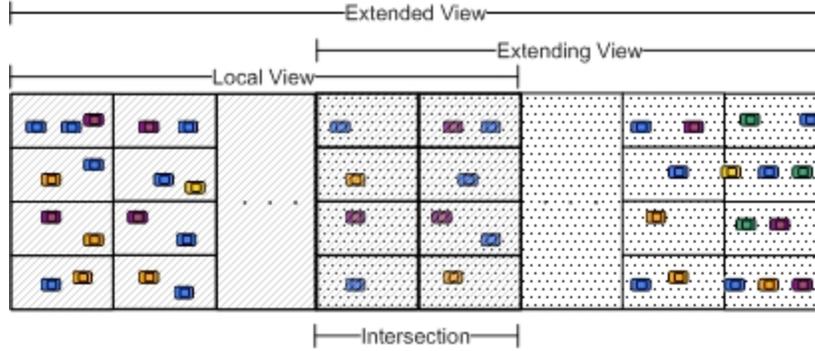


Fig. 2. Vehicle's Local View and Extended View

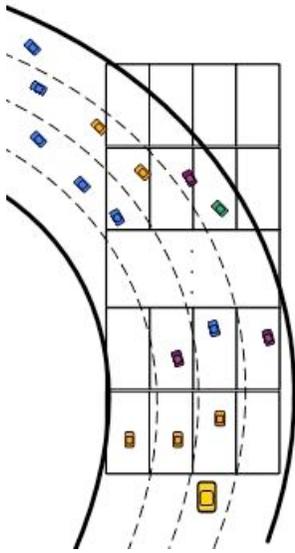


Fig. 3. Vehicle's Local View Mapped onto a Curved Roadway

C. Primary Records

A vehicle's local view is built entirely of received primary records. The primary record contains the basic information for a single vehicle. Each record can be represented in 29 bytes:

- *timestamp* (8 bytes) - time the record was generated
- *location* (16 bytes) - latitude and longitude
- *speed* (1 byte) - in meters/second
- *acceleration* (1 byte) - in meters/second²
- *heading* (1 byte) - in degrees from North (0-360)
- *altitude* (2 bytes) - in meters above sea level

1) *Initial Dissemination*: A vehicle broadcasts a *primary frame* containing its primary record at a random interval between 300-400ms. The primary frame, totaling 1033 bits¹, consists of the following fields:

- *type* (1 bit) - primary or aggregated frame
- *sender's location* (16 bytes) - latitude/longitude
- *primary record* (29 bytes)
- *digital signature* (28 bytes)
- *certificate* (56 bytes)

The primary record is signed by the original vehicle using the Elliptic Curve Digital Signature Algorithm (ECDSA) [9]. The certificate included in the frame contains the original vehicle's public key, signed by the certificate authority. Since the primary record is signed by the original sender, it cannot be tampered with by a re-broadcasting node. In addition, replay attacks are nullified by the presence of the timestamp inside the signed primary record.

A receiving vehicle will record the primary record and use the vehicle's public key as an identifier. Typically, only primary records from vehicles within the receiving vehicle's local view (*i.e.*, vehicles in front of the receiving vehicle) will be stored. But for some applications, such as merging assistance, awareness of vehicles behind or beside the receiving vehicle is important. In these cases, the vehicle would store the primary records of nearby (two clusters' worth, or within 128 m) following vehicles to be used in the application. Again, since these records are from following vehicles, they are not considered part of the local view.

2) *Re-Broadcast*: In order for primary records to reach vehicles farther than about 300 m, the records must be re-broadcast. In order to limit the number of re-broadcast messages used to propagate the frames, we use an adaptation of the Inter-Vehicle Geocast (IVG) algorithm [2] to accommodate the high volume of frames that need to be re-broadcast. In IVG, each node starts a timer for each frame it receives. If the timer expires and the frame associated with this timer has not been re-broadcast by any other node, the node re-broadcasts the frame. The timer value T_x for vehicle x is

$$T_x = T_{max} \cdot \frac{(R^\epsilon - D_{sx}^\epsilon)}{R^\epsilon},$$

¹If the underlying link-layer requires a frame size of full bytes, then the primary frame would be padded to 130 bytes.

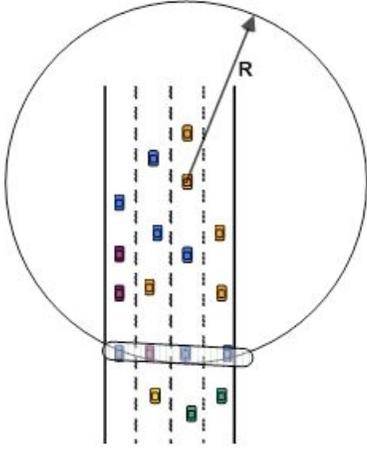


Fig. 4. Nearby Vehicles at the Boundary

where R is the transmission range and D_{sx} is the distance between vehicle x and vehicle s , the sender of the message. The authors of IVG suggest using $\epsilon = 2$ to generate a uniform timer value between $[0, T_{max}]$, where $T_{max} = 200$ ms [2].

By using IVG with its original settings, nodes that are close to each other at the boundary will have very similar, if not equal, timer values. This means the nodes' timers will expire at almost the same time, and the nodes will rebroadcast the frame essentially simultaneously, resulting collisions. This situation is shown in Figure 4, where the highlighted nodes, being a similar distance from the sender, will have similar timer values. To alleviate this problem, we modified IVG by making the re-broadcasting of frames probabilistic based on the surrounding vehicle density. In our *probabilistic-IVG* (p-IVG), when a vehicle receives a frame, it first selects a random number in $[0,1]$. If the selected number is less than $\frac{1}{density}$, the timer is started. If not, then the frame will not be re-broadcast by this vehicle. As the density increases, the number of nodes that will start their timers decreases. Because primary frames will be broadcast by the original sender every 300 ms, there is ample opportunity for fresh frames to be re-broadcast even if all vehicles discard a message during a particular round. In addition to adding the probabilistic timer, we use $\epsilon = 0.5$ in order to produce sparser timer values, because the timer values decrease faster as the distance from the original sender increases.

In addition to using p-IVG to limit the number of re-broadcasts, primary frames have a time-to-live (TTL) value to ensure that only fresh information is disseminated. The TTL of all primary frames in CASCADE is 1 second. If a node receives a frame and the difference between the original sending time and the current time is greater than the primary frame TTL, the frame will

be dropped.

3) *Aging*: The goal of CASCADE is to present highly accurate information about upcoming traffic conditions. So, it is important that old information is purged from the system. As the local view is concerned only with vehicles in front of the current vehicle, primary records are removed from the local view (but not necessarily from a vehicle's database) once the vehicle has physically passed the vehicle described by the record. Additionally, primary records may also be removed from the local view when no updates have been received in 1 second. With vehicles sending new primary frames 3 times a second, receiving no message about a vehicle for 1 second means that 3 messages in a row were not received, which would indicate that the vehicle corresponding to the old record has likely left the area.

D. Aggregated Records

Each vehicle builds its local view based on primary records received from other vehicles. In order to extend the view farther, vehicles exchange aggregated records. In this section, we describe how primary records are grouped into clusters, how clusters are aggregated, and how the aggregated records are disseminated and used to build the extended view.

1) *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. When clustering is done, a vehicle's heading and altitude are taken into account to ensure that vehicles are assigned to the proper cluster.

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 the compact data record is formed, the median speed of all vehicles in a cluster is calculated, the position of the center of the cluster is calculated, and the position of each vehicle is translated into $\{X, Y\}$ coordinates (in integer meters) with the local view origin as the origin, as shown in Figure 5. For this, we assume that the digital map in the vehicle provides the GPS position of the leftmost lane of the roadway.

Each compact data record, totaling 16 bits, contains the following fields:

- ΔX (3 bits) - difference between the vehicle's X coordinate position and the X coordinate for the center of its cluster
- ΔY (6 bits) - the difference between the vehicle's Y coordinate position and the Y coordinate for the center of its cluster
- ΔS (5 bits) - the difference between vehicle's speed and the median speed of the vehicles in the cluster

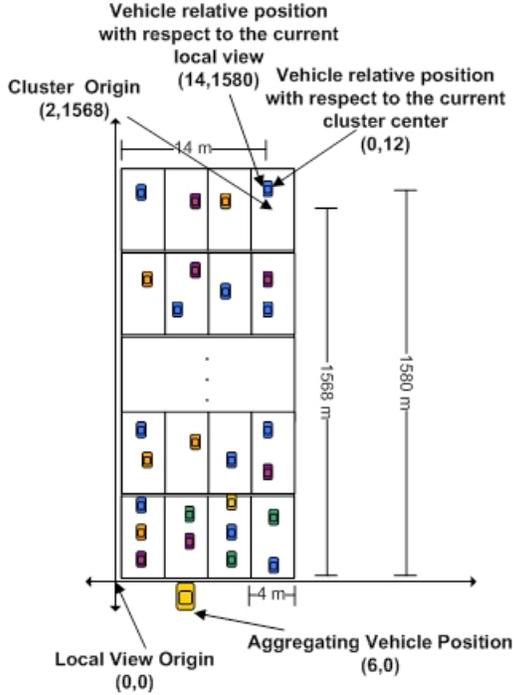


Fig. 5. Converting GPS Coordinates to $\{X,Y\}$ Coordinates

- *Speed Indicator (SI) Flag* (2 bits) - indicates if the vehicle's speed is within the acceptable range for the cluster or not

Since the base value is the center of the cluster (as in Figure 5), ΔX can have a negative value. With sign-magnitude representation, this means that three bits must be used for ΔX , one for the sign and two for the magnitude (with a maximum width difference of 2 m).

Since the cluster length is 64 m, ΔY can be represented with only 6 bits, using 1 bit for the sign and 5 bits for the difference between the vehicle's position and the cluster center (at most 32 m).

The range of acceptable values for ΔS is $[-15 \text{ m/s}, 15 \text{ m/s}]$. If the difference is outside of this range, then the ΔS field will be omitted, and the *SI Flag* will be set. The *SI Flag* can take one of three possible values $\{00, 01, 10\}$:

- 00 - ΔS can be represented in the allowed range $[\min \Delta S, \max \Delta S]$
- 01 - $\Delta S > \max \Delta S$, the vehicle is a *speeder*
- 10 - $\Delta S < \min \Delta S$, the vehicle is a *lagger*

The *SI Flag* has an important application in proactive collision warning and avoidance. Many accidents are due to vehicles that are either traveling much faster than surrounding vehicles (*speeders*) or traveling much slower than surrounding vehicles (*laggers*). If drivers can be alerted to these vehicles in advance, they may be able to avoid accidents.

By using differences in constructing the compact data record, we achieve a compression ratio of at least 88%. The primary data for each vehicle (location and speed) is represented in 136 bits (17 bytes) while the compact data for each vehicle is represented in at most 16 bits. The compression ratio is even higher if the ΔS field is omitted, as in the case of speeders and laggers.

E. Aggregation

Once compression has been completed, we form an *aggregated cluster record*, which is a concatenation of the compact data records of the vehicles in the cluster. Each aggregated cluster record contains the following fields:

- *cluster flag* (1 bit) - indicates if the cluster contains any vehicles
- *cluster median speed* (8 bits) - the median speed of the vehicles in the cluster in meters/second
- *number of vehicles* (4 bits) - the number of vehicles contained in the cluster
- *compact data records* (16 bits each) - concatenation of all of the compact data records for vehicles in this cluster

If there is a cluster that contains no vehicles, its *cluster flag* is set to 0 and no more information about the cluster is contained in the record.

Since a cluster is at most one lane wide and 64 m long, and the average vehicle length is 4 m, there can be at most 16 vehicles in a cluster.

1) *Initial Dissemination*: Every 4 seconds, the compression and aggregation algorithms are run. Once the aggregated cluster records are constructed, they are concatenated into a single frame and sent via broadcast. The aggregated frame includes the following fields:

- *type* (1 bit) - primary or aggregated frame
- *timestamp* (8 bytes)
- *sender's X-coordinate* (4 bits) - meters from the vehicle's local view origin, assuming 4 lanes of traffic
- *sender's location* (16 bytes) - latitude/longitude
- *aggregating vehicle's location* (16 bytes)
- *aggregated cluster records* - up to 104 records
- *digital signature* (28 bytes)
- *certificate* (56 bytes)

Note that if traffic is sparse, the aggregated frame will be much smaller than the maximum 2312 bytes, because empty clusters are represented by a single bit.

The cluster aggregated records are arranged according to their place in the view, starting with the bottom-left cluster, moving from left to right, and then increasing in distance from the aggregating vehicle. Since these records are always arranged in the same manner, there is no need for a cluster ID to be included in the frame or record.

2) *Re-Broadcast*: Aggregated frames are re-broadcast in the same manner as primary frames, using the p-IVG algorithm described in Section III-C.2. Aggregated frames originating from vehicles physically behind the receiving vehicle will be dropped, as will aggregated frames that are older than the TTL. All aggregated frames in CASCADE have a TTL of 5 seconds.

F. 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 (as shown in Figure 2). The receiving vehicle reconstructs the primary data for the vehicles in the received aggregated frame. Before placing the vehicles in 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 can be extended further through the receipt of other aggregated frames. Information about all vehicles not in the local view has a lifetime of only 10 seconds. Vehicles will be periodically receiving new aggregated frames that can be used to re-build the extended view with fresh data.

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. This problem has been addressed before [5], but requires more overhead and processing. In our approach, these vehicles can be caught at no extra cost because the comparison of the intersecting views are an essential part to extending the view. This procedure can be enhanced in future work with a framework to disseminate the identity of lying vehicles and isolate them.

IV. ANALYSIS

Table 1 lists some basic constants that will be used in the analysis.

A. Visibility

The maximum visibility with the local view is 1.6 km. We derive that value here. The maximum visibility is based on the maximum number of clusters, N_C , that can

cluster length, L_C	64 m
cluster width, W_C	4 m
average vehicle length + 1 m cushion, L_V	5 m
lane width, W_L	4 m
size of a single compact data record, S_{CDR}	16 bits
header for an aggregated cluster record, H_{ACR}	13 bits
header for an aggregated frame, H_{AF}	325 bits
encryption trailer for an aggregated frame, E_{AF}	672 bits
IEEE 802.11p frame size, S_F	2312 bytes

TABLE I
CONSTANTS USED IN ANALYSIS

fit into a single MAC layer frame, the maximum length of each cluster, L_C , and the number of lanes on the road, N_L . The number of vehicles per cluster, N_{VPC} , depends on the size of the cluster, the width of a lane, the length of a vehicle, and the distance between vehicles.

$$N_{VPC} = \frac{L_C W_C}{(L_V + D) W_L},$$

where D is the average distance between vehicles. In the case of dense traffic, where D can be as little as 1 m, the maximum number of vehicles per cluster is 10 vehicles.

The number of clusters that can be transmitted in a single frame, N_{CPF} , is determined by the size of the frame, the header and encryption data needed for the aggregated frame, and the size of the aggregated cluster record. Note that the size of the frame is represented in bytes and all of the other quantities are in bits.

$$N_{CPF} = \frac{(S_F * 8) - H_{AF} - E_{AF}}{S_{ACR}}$$

The size of an aggregated cluster record, S_{ACR} , is determined by the number of vehicles in the cluster, the size of the aggregated vehicle record, and the amount of header information needed for the aggregated cluster record.

$$S_{ACR} = (S_{CDR} N_{VPC}) + H_{ACR}$$

So, in dense traffic, which would be the worst case, N_{VPC} is 10, and so S_{ACR} would be 173 bits. So, at most 101 clusters can fit into a single frame.

If we assume that on average, there are 4 lanes on the highway, these 101 clusters would be divided into 25 cluster rows. Since each cluster represents 64 m ahead, the visibility of these 25 cluster rows is 1600 m, or 1.6 km.

For the extended view, we add local views from other vehicles to the base vehicles' local view. To determine the maximum visibility for an extended view, recall that aggregated frames, containing a vehicle's entire local view, have a lifetime limited to 8 seconds after they were originally sent. In dense traffic, frames will be their maximum 2312 bytes long. For DSRC with a 6 Mbps

data rate, the transmission time of such a frame is about 3 ms. Also, in dense traffic, each transmission will reach a vehicle 300 m away before being re-broadcast. The propagation delay for 300 m is about $1 \mu\text{s}$, which we treat as negligible. Assume a conservative 20 ms processing delay at each hop, including medium access delays. With these parameters, a full frame can travel 300 m in about 23 ms.

B. Accuracy

The position and the speed of each vehicle in the view will be represented in its cluster record with the only approximation being the rounding of the value. This means that the accuracy of the vehicle data is very high, especially compared with other aggregation algorithms that represent multiple vehicles with a single position and speed.

C. Knowledge Percentage

The maximum view length has been calculated based on the worst case, which means that all vehicles located within the local view boundary will be represented in the aggregated frame. In other algorithms, such as the ratio-based and cost-based algorithms in *TrafficView*, some vehicle data may be ignored due to lack of space in the aggregated view record, resulting in a decrease in the knowledge percentage.

V. EVALUATION

A. Simulation Setup

We evaluate the feasibility of CASCADE using the SWANS (Scalable Wireless Ad hoc Network Simulation) vehicular network simulator [1], [3]. SWANS has been shown to be scalable and efficient, supporting large numbers of mobile nodes. SWANS fully implements the IEEE 802.11b protocol, which we use as an approximation to IEEE 802.11p, and supports general mobility models, such as Random Walk and Random Waypoint. To produce realistic simulations, we extended SWANS by implementing the IDM (Intelligent Driver Model) vehicular mobility model [18] and the MOBIL (Minimizing Overall Braking decelerations Induced by Lane changes) lane changing model [10].

Table II summarizes our simulation settings. All vehicles in our simulations have a transmission range of 300 m. The roadway used is a four-lane divided highway of length 100 km. Vehicles enter the highway according to a Poisson distribution, and travel at a maximum speed of 30 m/s. The simulation is run for 360 seconds, resulting in a total of 500 vehicles generated. In the 360-second simulation runtime, the maximum distance traveled by any vehicle is 10 km.

transmission range	300 m
highway length	100 km
max distance traveled	10 km
vehicles generated	500
max speed	30 m/s
simulation runtime	360 seconds
high density	90 vehicles/km
medium density	66 vehicles/km
low density	53 vehicles/km

TABLE II
SIMULATION SETTINGS

We evaluate CASCADE with three different traffic density scenarios. In the high density case, there are an average of 90 vehicles/km. In medium density, there are an average of 66 vehicles/km, and in low density, there are an average of 53 vehicles/km. In each of the three scenarios, all the vehicles in the simulation are running an application that implements CASCADE. To show the improvement provided by p-IVG, we ran experiments where CASCADE used either basic flooding, standard IVG, or p-IVG.

B. MAC Delay

In order to investigate the impact of using p-IVG, we measured the amount of MAC-layer delay for each frame (split into primary frames and aggregated frames). This delay represents the time between the MAC-layer receiving the frame for transmission and delivering it to the physical layer. IEEE 802.11b by default does not use the RTS/CTS mechanism for reserving the wireless channel, so collision avoidance based on detecting the channel idle for a certain amount of time is used. If the wireless medium is very busy, the MAC delay will increase because the sender will not be able to detect that the channel is idle for the entire required period. In Figure 6, we show the cumulative distribution functions (CDFs) of the MAC delay experienced by primary frames for the three re-broadcast schemes (flooding, IVG, and p-IVG) at medium traffic density. Primary frames see the smallest MAC delay when p-IVG is used and the largest delay when basic flooding is used. p-IVG is also shown to be an improvement upon standard IVG. This is because p-IVG limits the number of nodes that may be re-broadcasting frames at the same time, thus making the wireless channel more available.

Figure 7 shows the CDFs of the MAC delay experienced by primary frames when using p-IVG for various traffic densities. The maximum MAC delay is still relatively small even at high densities. The key point is that with p-IVG, higher density does not necessarily mean more frames are sent, because the re-broadcasting of frames is affected by the density. At higher densities,

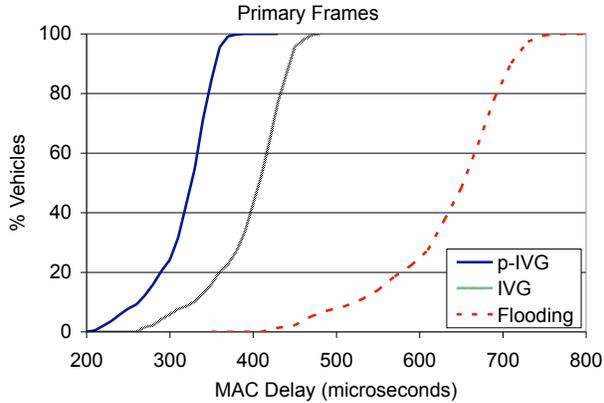


Fig. 6. CDF of MAC Delay for Primary Frames at Medium Density

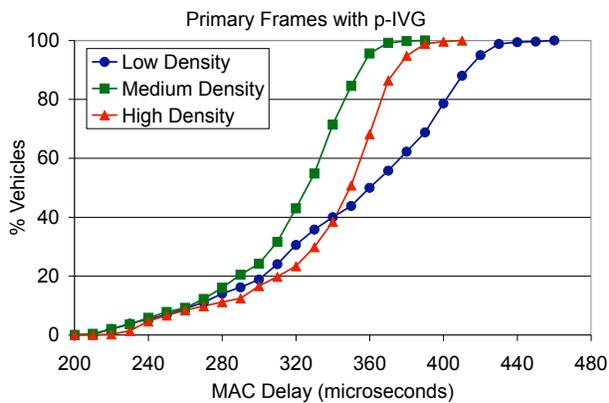


Fig. 7. CDF of MAC Delay for Primary Frames with p-IVG

fewer vehicles will re-broadcast frames.

C. Reception Rate

To assess the impact of wireless collisions, we measured the average reception rate over time. If a frame was transmitted and no other vehicle received it, then the frame was considered to not be received. Either the frame experienced a collision or no other vehicle was within 300 m of the sender. Figure 8 shows the reception rate for p-IVG, IVG, and flooding with medium density traffic, averaged every 10 seconds. The reception rate for IVG and flooding starts high because the number of vehicles in the simulation is still increasing. As more vehicles enter the system, the reception rate decreases. Using p-IVG results in a very high reception rate for the entire simulation.

To further demonstrate the problems we found with standard IVG, we show the reception rate per second when using IVG at medium traffic density (Figure 9). Each drop spike in the reception rate corresponds to a situation where multiple vehicles at the boundary re-broadcast frames at the same time, causing collisions.

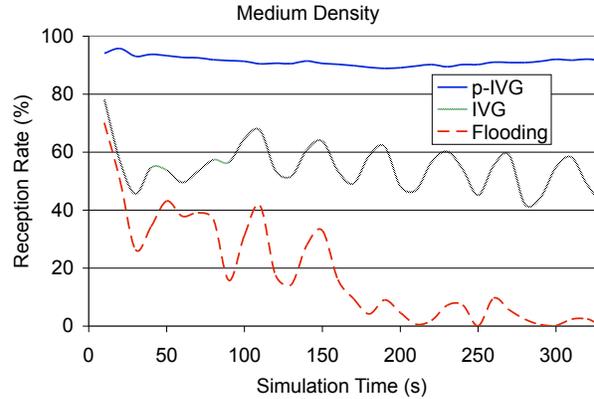


Fig. 8. Reception Rate at Medium Density

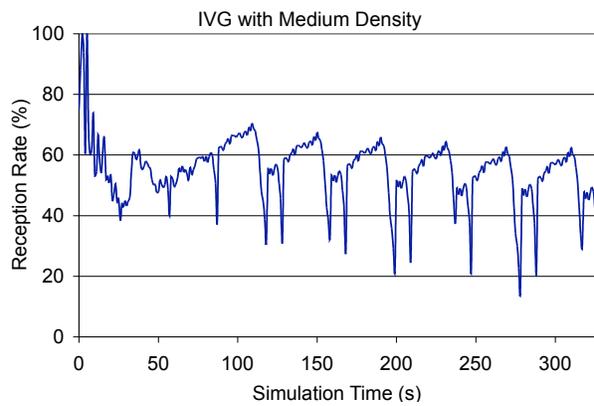


Fig. 9. Reception Rate with IVG at Medium Density

In Figure 10, we show the reception rate for p-IVG with each of the three traffic densities. In our definition, if just one vehicle receives the frame, it is considered to have been received. Higher density provides the best reception rate because there are more vehicles in range, thus more of a chance to receive the frame. The decreasing reception rate in the low density case is an artifact of how vehicles enter the simulation. Vehicles enter the highway with a low speed and gradually increase towards the maximum of 30 m/s. Until the last vehicle reaches the maximum speed, vehicles in front of it will be traveling faster, thus moving out of its transmission range. For low density, it takes longer for 500 vehicles to enter the system than with high density, so it takes longer for the last vehicle to reach its maximum speed.

D. Visibility

The goal of CASCADE is to provide information about upcoming vehicles, so visibility is one of the most important metrics we can measure. We consider visibility for a vehicle to be the distance between the vehicle and farthest vehicle in its extended view. Thus, as a vehicle's

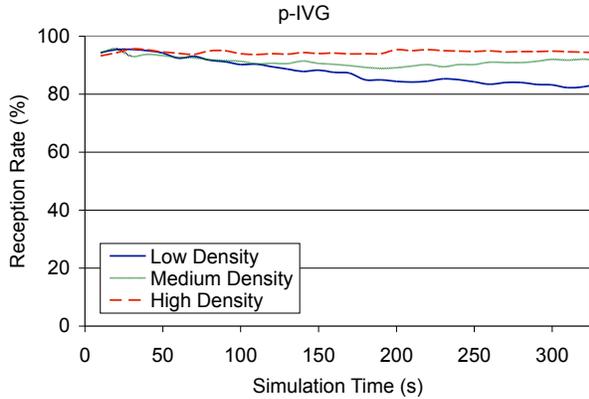


Fig. 10. Reception Rate with p-IVG

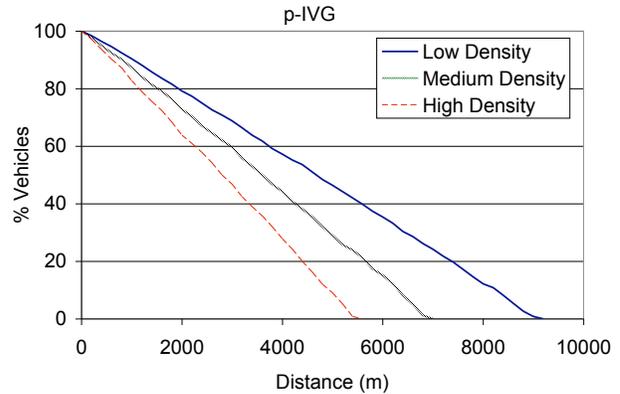


Fig. 12. Visibility with p-IVG

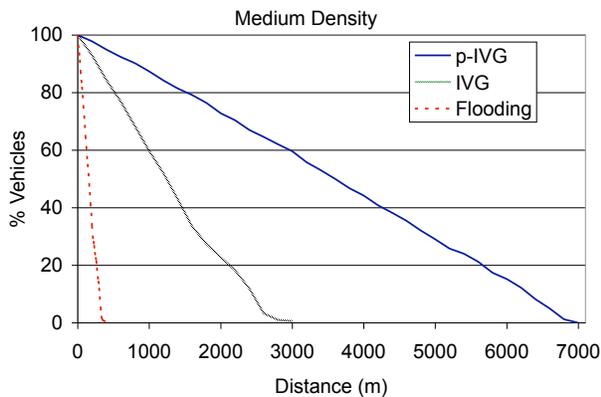


Fig. 11. Visibility at Medium Density

visibility increases, its knowledge about upcoming traffic conditions increases.

Figure 11 shows the percentage of vehicles that have a particular minimum visibility with the three re-broadcasting techniques at medium density traffic. For example, with p-IVG 60% of the vehicles have a visibility of 3000 m or more, while with standard IVG, almost no vehicles have that much visibility. Visibility highlights the importance of the reception rate. If there are many collisions, then information is not able to be disseminated to distant vehicles, so their visibility is reduced.

Figure 12 shows the visibility for p-IVG at high, medium, and low traffic densities. Consider that the optimal density for high visibility is to have exactly one vehicle positioned every 300 m (*i.e.*, at the boundaries). Thus, with low density, there is less of a chance that re-broadcasted aggregated frames experience collisions, so they are able to travel farther. As the density decreases, the chances of having a single vehicle at the boundary increase. Since p-IVG only deals with nearby vehicles re-broadcasting frames at the same time, there may be other transmissions (such as the initial broadcast of primary

or aggregated frames and re-broadcasts from vehicles in range, but not nearby) that may collide with the transmission.

VI. CONCLUSION AND FUTURE WORK

We have presented CASCADE, a technique for accurate aggregation of vehicle data. The local view presents data gathered from primary records, which are sent in signed frames containing a vehicle's position information. 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. 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. We have also adapted Inter-Vehicle Geocast, adding a probabilistic step based on surrounding traffic density to reduce the number of frame collisions.

In future work, 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. Currently, our analysis and compression methods are based on the assumption of a four-lane highway. We plan to develop methods that allow for various sized highways and that take advantage of vehicles traveling in the opposite direction for dissemination and corroboration of vehicle data.

REFERENCES

- [1] JiST/SWANS. <http://jist.ece.cornell.edu>, 2004.
- [2] A. Bachir and A. Benslimane. A multicast protocol in ad hoc networks: Inter-vehicle geocast. In *Proceedings of the IEEE Vehicular Technology Conference - Spring*, pages 2456–2460, Jeju Island, Korea, Apr. 2003.

- [3] R. Barr, Z. Hass, and R. van Renesse. *Handbook on Theoretical and Algorithmic Aspects of Sensor, Ad hoc Wireless, and Peer-to-Peer Networks*, chapter 19: Scalable Wireless Ad Hoc Network Simulation, pages 297–311. CRC Press, 2005.
- [4] T. ElBatt, S. Goel, G. Holland, H. Krishnan, and J. Parikh. Cooperative Collision Warning Using Dedicated Short Range Wireless Communication. In *Proceedings of the ACM Workshop on Vehicular Ad Hoc Networks (VANET)*, Los Angeles, CA, Sept. 2006.
- [5] 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.
- [6] L. Hu and D. Evans. Secure aggregation for wireless networks. In *Proceedings of the Symposium on Applications and the Internet Workshops (SAINT)*, page 384, Orlando, FL, Jan. 2003.
- [7] J.-P. Hubaux, S. Capkun, and J. Luo. The security and privacy of smart vehicles. *IEEE Security and Privacy Magazine*, 2(3):49–55, 2004.
- [8] IEEE. Status of Project IEEE 802.11p. http://grouper.ieee.org/groups/802/11/Reports/tgp_update.htm, Nov. 2006.
- [9] D. Johnson, A. Menezes, and S. Vanstone. The elliptic curve digital signature algorithm (ECDSA). *International Journal of Information Security*, 1(1):36–63, Aug. 2001.
- [10] A. Kesting, M. Treiber, and D. Helbing. MOBIL: General lane-changing model for car-following models. In *Proceedings of the Transportation Research Board Annual Meeting*, Washington, DC, Jan. 2007.
- [11] C. Lochert, B. Scheuermann, and M. Mauve. Probabilistic aggregation for data dissemination in VANETs. In *Proceedings of the ACM Workshop on Vehicular Ad Hoc Networks (VANET)*, pages 1–8, Montreal, Canada, Sept. 2007.
- [12] S. Madden, M. J. Franklin, J. M. Hellerstein, and W. Hong. TAG: a Tiny AGgregation service for ad-hoc sensor networks. *ACM SIGOPS Operating Systems Review*, 36(SI):131–146, 2002.
- [13] T. Nadeem, S. Dashtinezhad, C. Liao, and L. Iftode. TrafficView: Traffic Data Dissemination using Car-to-Car Communication. *ACM Mobile Computing and Communications Review (MC2R), Special Issue on Mobile Data Management*, 8(3):6–19, July 2004.
- [14] F. Picconi, N. Ravi, M. Gruteser, and L. Iftode. Probabilistic Validation of Aggregated Data in Vehicular Ad-hoc Networks. In *Proceedings of the ACM Workshop on Vehicular Ad Hoc Networks (VANET)*, Los Angeles, CA, Sept. 2006.
- [15] B. Przydatek, D. Song, and A. Perrig. SIA: Secure information aggregation in sensor networks. In *Proceedings of the International Conference on Embedded Networked Sensor Systems (SenSys)*, pages 255–265, Los Angeles, California, Nov. 2003.
- [16] M. Raya, A. Aziz, and J.-P. Hubaux. Efficient Secure Aggregation in VANETs. In *Proceedings of the ACM Workshop on Vehicular Ad Hoc Networks (VANET)*, pages 67–75, Los Angeles, CA, Sept. 2006.
- [17] M. Raya and J.-P. Hubaux. The security of vehicular ad hoc networks. In *Proceedings of the ACM Workshop on Security of Ad Hoc and Sensor Networks*, Nov. 2005.
- [18] M. Treiber, A. Hennecke, and D. Helbing. Congested traffic states in empirical observations and microscopic simulations. *Physical Review E*, 62(2):1805, 2000.
- [19] US Department of Transportation. Incident Management: Detection, Verification, and Traffic Management. Intelligent Transportation Systems Field Operational Test Cross-Cutting Study, Sept. 1998.
- [20] 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.
- [21] US Department of Transportation and IEEE. IEEE 1609 - Family of Standards for Wireless Access in Vehicular Environments (WAVE). Intelligent Transportation Systems Standards Fact Sheet, http://www.standards.its.dot.gov/fact_sheet.asp?f=80, Jan. 2006.
- [22] US Department of Transportation, National Highway Traffic Safety Administration. Identify Intelligent Vehicle Safety Applications Enabled by DSRC. DOT HS 809 859, Vehicle Safety Communications Project Task 3 Final Report, Mar. 2005.
- [23] D. Wagner. Resilient aggregation in sensor networks. In *Proceedings of the ACM Workshop on Security of Ad Hoc and Sensor Networks (SASN)*, pages 78–87, Washington DC, Oct. 2004.
- [24] Y. Yang, X. Wang, S. Zhu, and G. Cao. SDAP: A secure hop-by-hop data aggregation protocol for sensor networks. In *Proceedings of the ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 356–367, Florence, Italy, May 2006.