Cooperative Certificate Revocation List Distribution Methods in VANETs

Michael Nowatkowski, Chris McManus, Jennie Wolfgang, and Henry Owen, III

School of Electrical and Computer Engineering
Georgia Institute of Technology, Atlanta, Georgia, USA
{nowatkom,cmcmanus3,jwolfgang3,owen}@gatech.edu

Abstract. This paper discusses two new methods for distributing certificate revocation lists (CRL) in a vehicular ad hoc network environment using cooperative methods. The main purpose for using cooperative methods is to attempt to reduce the number of collisions in the dedicated short range communication (DSRC) broadcast environment. The reduced number of collisions will increase the effective throughput of the medium. The first method uses a polling scheme to determine which nodes possess the most number of CRL file pieces. The second method takes advantage of the multiple service channels available in DSRC. Both methods use a form of coding that reduces the impact of the piece problem. Both methods are compared to the Code Torrent method of file distribution.

Keywords: Certificate revocation list, vehicular ad hoc networks, network security.

1 Introduction

Vehicular ad hoc networks (VANETs) are a subset of mobile ad hoc networks (MANETs) composed of network-equipped vehicles and infrastructure points. Infrastructure points provide a connection to network services, similar to an access point in traditional wireless networks. The infrastructure points are referred to as road-side units (RSUs). Vehicles will have an on-board unit (OBU) installed to communicate and store information. VANETs enable vehicles to communicate directly with other vehicles using vehicle-to-vehicle (V2V) communications and with RSUs using vehicle-to-infrastructure (V2I) communications. While sharing some of the same limitations of traditional MANETs, such as lack of infrastructure and limited communications range, VANETs have several dissimilarities that make them a much different research area. VANETs are hosted on vehicles, so power and space for radio, storage, and processing units are not an issue; however, privacy becomes a very important issue. Researchers have discussed several issues pertaining to the trade-offs between privacy (confidentiality) and authenticity [1].

While consumer VANETs have not yet been deployed, the idea has been discussed by research groups, government agencies, and vehicle manufacturers for several years. In the United States, the FCC designated a 75 MHz band in the 5.9 GHz range

J. Zheng et al. (Eds.): ADHOCNETS 2009, LNICST 28, pp. 652–665, 2010.

for this purpose, named Dedicated Short Range Communication (DSRC). Several other countries have also apportioned frequencies for VANETs. Currently, test beds in the United States and Europe are fielding prototypes and testing some of the initial protocols. The IEEE P1609 working group has developed and issued a series of Trial-Use Standards for Wireless Access in Vehicular Environments (WAVE). The IEEE 1609.2 standard is written specifically for security applications [2].

For security and safety reasons, messages must be authenticated to ensure that they were sent by a legitimate member of the VANET. This is especially critical for safety-related messages. Public key certificates are used for message authentication to prevent attackers outside of the network from causing harm. Certificates normally have a time period for which they are valid, described by a start time and an end time, or simply a "lifetime." A pseudonym is a short-lifetime certificate that does not contain identity-linking information. Pseudonyms are requested from a certificate authority (CA) using a longer lifetime certificate. Changing pseudonyms periodically greatly increases end-user anonymity while still maintaining a reliable means of authentication.

When a node's certificate is identified for revocation, the currently used certificate must be revoked along with every pseudonym stored in the OBU. This assumes that whatever caused the current certificate to be revoked will cause future uses of certificates by the same node to also trigger a revocation. Examples that would cause this event include a malfunction in the vehicle's sensors causing erroneous warning messages to other vehicles, or malicious activity by a given vehicle. By revoking all of the pseudonyms, further potential damage is avoided. The information regarding which certificates are no longer valid, i.e., revoked, is sent out in a certificate revocation list (CRL). The size of the CRL is directly proportional to the revocation rate, the number of nodes in the system, and, for VANETs, the number of pseudonyms used by each vehicle.

Using the fields described in [2] and the lengths found in Annex D.3.3 of the IEEE 1609.2 standard for other structures, the size of a CRL is 230 bytes plus up to an additional 14 bytes per revoked certificate. The certificate ID is found by generating the SHA-256 hash of the certificate and then taking the lower-order 10 bytes of the hash output. The expiry date is an additional 4 bytes, which results in 14 bytes per certificate. Up to $2^{64} - 1$ revoked certificates can be included in a single CRL. CRL file sizes can range from megabytes to several hundred megabytes, depending on the distribution frequency and the rate of certificate revocation.

This paper contributes two new methods for distributing CRL files in a VANET. The methods are Most Pieces Broadcast (MPB) and Generation per Channel (GPC). Both of these methods attempt to reduce the wireless channel contention by limiting the number of nodes broadcasting, as well as making use of the multiple channels available in DSRC.

The rest of the paper is organized as follows. In section 2 the layer 1 and 2 properties of DSRC and WAVE are explained. Section 3 is an overview of erasure coding and network coding, which is used in the existing Code Torrent file distribution method. Other methods of CRL distribution and the existing Code Torrent method are covered in sections 4 and 5 respectively. Section 6 contains details of the proposed new CRL distribution techniques, Most Pieces Broadcast and Generation per

Channel. Section 7 discusses the simulation set up, while section 8 shows the results of the simulation. Finally, section 9 concludes the paper.

2 Overview of DSRC/WAVE

Dedicated short range communication (DSRC) is the physical layer used with VANETs. The term Wireless Access in Vehicular Environment (WAVE) is also used to describe the communications in VANETs. It is a 75 MHz band in the 5.9 GHz frequency range. Seven non-overlapping channels are available in DSRC, as shown in Fig. 1. Two different classes of channels are described for use in DSRC/WAVE. The first class is the control channel, referred to as CCH, which is a single channel "reserved for short, high-priority application and system control messages [3]." The other class of channel is the service channel, or SCH, which has six different 10 MHz channels that support a wider range of applications and data transfer. Each node in the VANET must monitor the CCH during time periods designated as control channel intervals. The time period for an entire CCH Interval and SCH Interval is called a Sync Interval (see Fig. 2). Between control channel intervals, nodes may switch to participate on a SCH for applications such as file downloads. A WAVE announcement action frame is an announcement that is broadcast on the CCH to inform nodes of available content on the SCHs. It is used to initiate a WAVE basic service set (WBSS) to allow multiple nodes to communicate on an SCH.

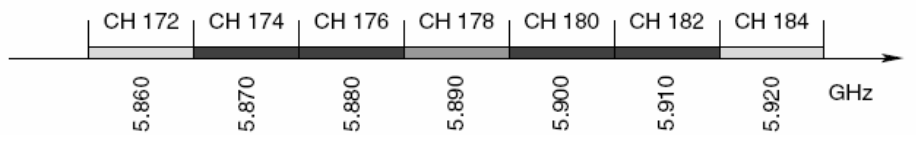

Fig. 1. DSRC Channels. CH 178 is the control channel, the others are the service channels [3]

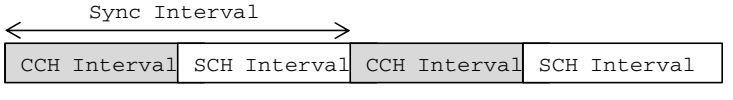

Fig. 2. WAVE Sync Interval [3]

Packet collisions and radio reception faults are the primary causes for nodes not receiving data sent over a wireless medium. DSRC uses carrier sense multiple access/collision avoidance (CSMA/CA) to reduce the number of collisions and to allow fair access to the medium. Each node using CSMA/CA must first sense the medium to determine if the medium is idle or busy. When the medium is idle, nodes wait for a fixed arbitration inter-frame space (AIFS) time plus a random time between zero and the minimum contention window value (CWmin) before sending data. Packet prioritization is scheduled using Enhanced Distributed Channel Access (EDCA) based on the IEEE Standard 802.11e, described in [4] and [5]. EDCA specifies four different access categories (ACs) for different priorities of data. The CCH and SCH have slightly different parameters for the different ACs. Table 1 shows the values of the

contention window timers and the AIFS timers for the CCH and SCH over the different ACs. The times are calculated using 15 slots for aCWmin, 511 slots for aCWmax, and 16 microseconds for the slot time, as described in [3].

Table 1. EDCA Parameters in DSRC, time in microseconds [3]

		Control Channel (CCH)			Service Channel (SCH)		
ACI	AC	CWmin	CWmax	AIFS	CWmin	CWmax	AIFS
1	Background	240	8176	144	240	8176	112
0	Best Effort	112	240	96	240	8176	48
2	Video	48	112	48	112	240	32
3	Voice	48	112	32	48	112	32

Broadcast mode sends data to the network broadcast address instead of to a specific destination address. The ready-to-send (RTS) and clear-to-send (CTS) handshake used for point-to-point data transmission is not used in broadcast mode. The RTS/CTS handshake is an attempt to reduce the number of collisions resulting from the hidden terminal problem. Also, acknowledgements (ACKs) are not sent by receiving nodes when broadcast mode is used. Therefore, since there are no RTS, CTS, or ACKs when using broadcast mode, channel overhead is reduced since fewer transmissions are needed to complete a data transfer. This allows for a potentially greater throughput of data. However, there is no confirmation that the data was successfully received by any destination, and there is a higher risk of collisions due to hidden terminals [6]. Beacons are transmitted by OBUs at regular intervals, either three beacons per second or ten beacons per second, depending on the traffic density. The beacons broadcast information about the location, speed, and heading of the vehicles. This information is used for safety-related and network-related uses.

WAVE supports two different network layer protocols, IPv6 and the WAVE short message protocol (WSMP). IP traffic is not allowed on the CCH, but WSMP traffic is allowed on both the CCH and the SCHs. The WSMP does not use a MAC address or IP address to identify the source or destination. Instead, WSMP uses an Application Class IDentifier (ACID) and an Application Context Mark (ACM) to identify the application class and the instance of the application class, respectively. This helps to increase the level of user-anonymity since the MAC address and IP address could be used to identify nodes and their presence in the VANET. A WAVE Short Message (WSM) is a message format used for sending messages using the WSMP. WSMs may be sent on either the CCH or the SCHs. According to [2], a CRL shall be transmitted as a WSM.

3 Overview of Network Coding versus Erasure Coding

Many VANET file-sharing models discuss the implementation of Network Coding or Erasure Coding as a means of making data dissemination throughout a network more efficient and timely [7-10]. These coding techniques reduce the impact of the piece problem that exists in file sharing techniques that simply split a file into multiple pieces, such as Bit Torrent. To download a copy of the original file using Bit Torrent, the destination node must download a copy of every specific piece of the file from

peers. The piece problem occurs when there is difficulty downloading one or more pieces due to the lack of availability of the missing pieces at other peer nodes. The missing pieces prevent the destination from completing the download. Network Coding and Erasure Coding techniques mitigate the piece problem by coding the file in a manner such that the file can be reconstructed from a set of pieces, rather than having to download every specific piece.

Erasure Coding is a method of breaking down a file into N pieces and then encoding those N pieces into M > N pieces in such a way that the original file can be reconstructed with a subset of those M pieces. This makes data easily available to nodes within a network that is using a peer-to-peer form of file sharing but has the drawback of requiring additional memory space for stored files. The rate of a file can be expressed as R = N/M, meaning that a code with a rate of R=1/4 has four times as many pieces after encoding than before, theoretically representing a four-fold increase in storage costs. The original file can be reconstructed from any T encoded pieces. In what are known as Optimal Erasure Codes, T is equal to N. Otherwise T is commonly greater than N but always less than M. Once the M pieces are encoded, the pieces are not further encoded by other nodes. Nodes using Erasure Coding share exact copies of pieces that have been received.

Network Coding also breaks a file into many pieces; however, file pieces received are combined with other stored pieces of the same file to generate "new" coded frames to share with other nodes. The defining principal behind Network Coding is the encoding of information at intermediate nodes within a network. Due to the overhead involved with generating a coded file, large files may be split into smaller sized files. These smaller files, referred to as generations, are then coded independently [9].

4 Overview of Previous CRL Distribution Methods

Papadimitratos, Mezzour, and Hubaux in [7], state that the problem of distributing certificate revocation lists in VANETs has not yet been solved. To the best of their knowledge, the efforts investigating security in vehicular communications have not examined “the fundamental problem of how to distribute the certificate revocation lists across a large-scale and multi-domain system as the vehicular communications systems [7].” The approach they present is one that assumes fixed infrastructure consisting of roadside units spaced every 1 to 3 kilometers. Assuming the use of Erasure-type codes and an allocated bandwidth of 1 kbps, and a 200 KB CRL consisting of 500 bytes of security overhead plus 4 bytes per revoked certificate, they claim “all vehicles can obtain the latest CRL within a delay of 30 or 40 minutes of drive, e.g. the duration of a commute [7].” Only vehicle-to-infrastructure communication is used in this method.

The idea of using application layer peer-to-peer technology to propagate CRLs in VANETs has already been proposed and a rough comparison to a fixed RSU infrastructure has been completed. Laberteaux, Haas and Hu in [10] note that “previous work assumed that CRLs will be distributed by broadcasting updates from roadside units.” They point out the problem that a large number of RSUs would be required and that some vehicles would never encounter a roadside unit. They propose an epidemic car-to-car method to communicate the CRLs. Their results show a “better

performance for a single deployed RSU than the performance of 325 RSUs without epidemic car-to-car passing of CRLs." Their simulation model consisted of a time-contact model that did not take radio properties or file transfer protocols into consideration.

Nandan et al. in [11] present details about SPAWN, which is a swarming protocol intended for VANETs. They assume each node in the network is running the ad hoc on-demand (AODV) routing protocol, communicating via ad-hoc mode (V2V) as well as infrastructure mode (V2I), and that RSUs are approximately 5 to 10 miles apart. Nodes "gossip" with each other using UDP messages containing information about files they are willing to share as well as files they are interested in downloading. Once a peer has been discovered that wants to share the same file, they conduct a TCP file transfer, possibly over multiple hops. They experimented with three different file piece selection strategies: 'First Available, Rarest First, and Rarest Closest." Nodes estimate the distances of a "particular peer by looking at the gossip message of the peer, and then calculating the number of nodes that have stamped the packet from the relevant field." They also assume a popularity index for files to be 20% to 80% instead of the 100% needed in a CRL scheme. The popularity index is defined as the percentage of nodes in the network that are interested in downloading the file. This is proportional to the number of nodes that propagate file pieces.

For the purposes of general file swapping in VANETs, Lee et al. in [12] advocate "peer-to-peer file swarming in which users out of access point range can still download parts of files from others." Car Torrent is based upon SPAWN. Clients use k-hop limited scope broadcasting (known as gossiping) and a piece selection strategy to "optimally download files from one another."

5 Overview of Code Torrent

Like SPAWN and Car Torrent, Code Torrent's main purpose is for file distribution in VANETs; however, Code Torrent approaches data dissemination from a single-hop perspective by only allowing peers to share with their immediate neighbors, thereby eliminating the need for a routing mechanism [13]. In the Code Torrent scheme, participating nodes (seeds) broadcast information regarding the files that they have and can share to their immediate neighbors. All nodes promiscuously listen to the broadcasts that are sent by their neighbors. If a node receives information that one of its neighbors has a file that is of particular interest to it, the node will broadcast a request for that file to all of its neighbors. Any neighbor node that receives the request and that has all or part of the requested file to share then responds with a "coded frame" containing parts of the requested file. The interested node will continue to request coded frames from its neighbors until it has enough to construct the entire file. This is very similar to the gossiping used in SPAWN and Car Torrent.

Code Torrent uses random linear coding (Network Coding) to construct the coded frames. Each coded frame contains a random selection of various parts of the requested file and the encoding vector that the sender used to encode it, as well as the file identification and transaction identification. If a file is separated into N pieces, then the receiving node needs to collect $M > N$ coded frames with linearly independent encoding vectors to ensure that the entire file has been collected [12].

Performance analysis of Code Torrent and Car Torrent shows that as the number of nodes increases, like in a metropolitan area, Car Torrent performance gradually degrades because of all the gossip messages and the underlying routing protocol. However, delay to obtain files in Code Torrent decreases as mobility (and congestion) increases. Code Torrent performs better in congestion situations than Car Torrent, but still suffers from all nodes attempting to access the medium. This results in nodes waiting until the medium is idle for several slot times before broadcasting. The other condition that occurs is the hidden terminal problem, which results in collisions.

5.1 Using Code Torrent for CRL Distribution

While the original intent of Code Torrent was for sharing files that only a subset of the network was interested in, the method may also work very well for files such as the CRL that every node in the VANET needs. The Code Torrent developers refer to the percentage of nodes interested in downloading the file as the file's "popularity." The CRL will have a popularity of 100% since every node will want the most recent CRL to protect them from malicious users and malfunctioning equipment, as well as to increase the overall security and safety of the VANET.

The basic Code Torrent algorithm for CRL distribution is:

```
1. seed nodes broadcast a beacon or WAVE service
announcement with CRL file description during the CCH
interval

2. nodes that want the file broadcast a request for the
file; requests are sent until the file is reconstructed
(i.e., the node receives enough linearly independent
coded frames)

3. every node hearing the request broadcasts a new
coded frame of the file
```

6 Development of Methods to Reduce Channel Contention

While Code Torrent makes improvements over other previous methods by using broadcast mode to avoid routing difficulties and network coding to mitigate the piece problem, it still suffers under heavy load. This is due to the method having every OBU broadcast requests and relevant coded frames. To reduce the contention for the wireless channel, we propose two methods to attempt to reduce the number of OBUs contending for the channel. Reducing the number of OBUs contending for the channel will increase the number of coded frames successfully received by the OBUs in the VANET.

6.1 Description of "Most Pieces Broadcast" (MPB) Method

The best situation for reducing contention is to limit the number of broadcasting nodes to a single node. At this point there is no contention for the channel, so the throughput will be the highest possible within the constraints of the channel. The

Most Pieces Broadcast method creates a situation where there is a single node broadcasting within the given radio broadcast range of selected nodes. The hidden terminal problem still exists for those OBUs that are within radio range of more than one selected node, so collisions still occur, but contention for the channel is reduced significantly. MPB will work in both V2I and V2V conditions with or without the presence of RSUs.

MPB takes advantage of the CCH in DSRC to accomplish the node selection. During the CCH interval, nodes exchange beacon packets. MPB adds information to the beacon packet to identify the CRL and the number of pieces, i.e., coded frames, that the node possesses. Each CRL is uniquely identified by the pair of fields containing the CA identifier (8 bytes) and the CRL serial (4 bytes). Methods for reducing the number of bytes required for the CA identifier and CRL serial, such as hashing or using only the lower two bytes of each field, could be used to reduce the number of bytes added to the beacons. The number of pieces can be represented with 2 bytes, allowing up to 65535 pieces to be accounted for. This requires a total of 6 bytes added to the beacon to advertise the number of CRL pieces possessed by each node.

Beacons that include CRL piece information are used to build a list at each OBU that contains values larger than the number of pieces possessed by the OBU receiving the beacon. If an RSU is within radio range of the OBU, the RSU is automatically placed at the top of the list since the RSU will always have the complete CRL. If there are no numbers on the OBU's list, that means the OBU has the most number of pieces within its listening range, so it will become the broadcaster during the SCH interval. During the SCH interval, only OBUs that have been "selected" will broadcast CRL pieces. If an OBU has other numbers on its list, this means that other OBUs within its radio range have more CRL pieces, so it will listen for broadcast pieces. These OBUs will store any new pieces they receive during the SCH interval and update their pieces count to send out with their next beacon on the CCH. If no pieces are received for a time period equal to (CWmin + AIFS) times the number of entries on the list, the OBU will begin to broadcast. For SCH AC0, CWmin + AIFS is 288 microseconds, as seen in Table 1.

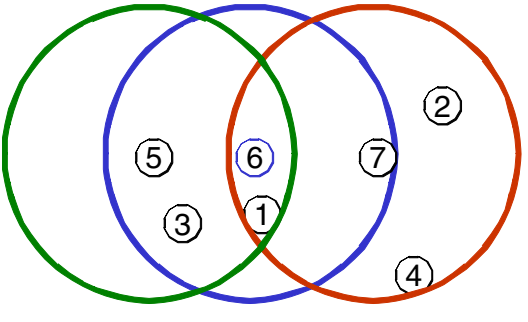

Node (pieces)	List entries	Wait time (us)
1	7,6,5,3	1152
2	7	288
3	6,5	576
4	7	288
5	6	288
6	7	288
7	none	0

Fig. 3. Example of Most Pieces Broadcast file piece list

The MPB method is illustrated with the help of Fig. 3. The numbers inside the diamonds are the number of CRL file pieces possessed by that node. Node 7 has the most CRL file pieces, so there are no nodes on its list; therefore, node 7 will begin broadcasting CRL pieces as soon as the SCH interval begins. Node 6 puts node 7 on its list, so node 6 will not broadcast during the SCH interval. Node 5 puts node 6 on its list, so it will not broadcast during the SCH interval. During the SCH interval, node 6 will receive pieces from node 7, but node 5 will not receive any pieces since node 6 is silent. Therefore, node 5 will wait 288 microseconds and then begin to broadcast. Nodes 6 and 1 have the opportunity to receive pieces from both node 7 and node 5 at this point, although the possibility of receiving many collisions also exists.

The basic MPB algorithm for CRL distribution is:

```
1. all nodes broadcast a beacon on CCH with CRL file
description and number of pieces possessed

2. nodes keep a list of other nodes that possess more
pieces than self (if RSU is present it automatically
goes to the top of the list)

3. if no nodes are on the list then broadcast pieces
for the remaining SCH interval;

   else wait for a fraction of SCH interval equal to
288 microseconds times the number of entries on the
list;

4. if no pieces are received then broadcast pieces for
remainder of SCH interval;

   else wait for remainder of SCH interval; //node is
receiving pieces

5. return to 1.
```

6.2 Description of “Generation Per Channel" (GPC) Method

While MPB significantly reduces the possible number of collisions, there is also additional overhead on the CCH due to adding 6 bytes to each beacon to distribute CRL piece information. To reduce the additional overhead and to allow for more nodes to participate by broadcasting their pieces, while still reducing the overall number of nodes contending for the medium, a method that takes advantage of the multiple DSRC channels was developed. Generation per Channel uses multiple service channels to distribute network coding generations of the CRL. Network coding generations are described in section 3. Four service channels were used in the simulations here. Announcements on the CCH inform nodes which generations are available on which SCHs. This method has the benefits of pure Code Torrent, but the number of nodes contending for the medium is divided by the number of SCHs used.

GPC could also use the MPB method on multiple SCHs, but this is not analyzed at this time. It is expected that for a single SCH interval of 50 milliseconds, GPC using

the MPB method will perform the same as MPB on a single channel; however, when a full CRL distribution is analyzed in future work, GPC using the MPB method may show different results than either MPB on a single channel or Code Torrent on multiple channels.

A drawback of this method is that nodes must spend some time on every channel to receive the entire CRL. If the CRL was split such that each generation could be used independently, this would help this situation by allowing for use of the partial CRL.

The basic GPC algorithm for CRL distribution is:

```
1. the CRL is split into 2-6 "generations" (each
generation is a different part of the file)

2. service announcements are broadcast on the CCH to
make nodes aware of which generation is on which
channel

3. each generation is distributed on a different SCH

4. nodes go to different SCHs to get the parts of the
file (could use either Code Torrent or MPB)
```

7 Simulation Setup

Simulation of Code Torrent, Most Pieces Broadcasts, and Generation per Channel was conducted using the Georgia Tech Network Simulator (GTNetS) [14]. The model consisted of five broadcasting OBUs during MPB plus a varying number of mobile OBUs in a wireless environment (see Fig. 4). This simulates a length of road with mobile OBUs moving on the road. Values of 50, 250, 500, and 1000 additional OBUs were selected to obtain initial results to compare the three techniques. The density of OBUs ranges from about 24 OBUs per kilometer up to about 471 OBUs per kilometer in the scenarios.

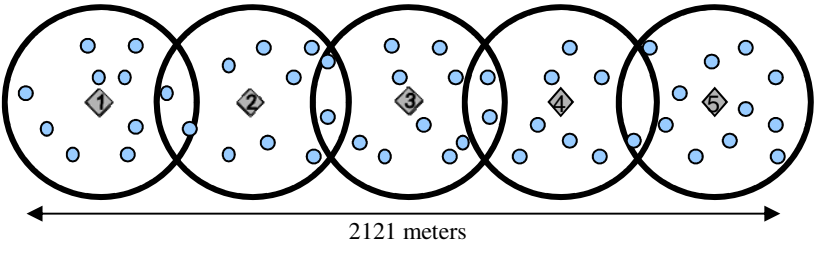

Fig. 4. Simulation layout with 5 broadcasting nodes and 50 OBUs

Parameters from the IEEE 802.11 standard were changed in GTNetS to the specific parameters for a DSRC service channel with an access category index (ACI) of zero. These parameters are shown in Table 2. The simulation time was a single SCH interval of 50 milliseconds. The OBUs were mobile in the simulation, but they did not

Table 2. DSRC Parameters Used in Simulation

Parameter	Value
Slot time	16 microseconds
Channel data rate	6 megabits per second
AIFS	48 microseconds
CWmin	240 microseconds
Radio Range	300 meters
Packet size	500 bytes

travel far enough to impact the preliminary results presented here. Four service channels were used for GPC.

8 Simulation Results and Analysis

The performance metrics used to compare the methods are the average number of packets received per OBU and the average number of collisions detected per OBU. The packets received per OBU is the total number of packets received by all OBUs during the 50 millisecond simulation period divided by the total number of OBUs. The average number of collisions per OBU was derived in the same way. The results from the average of ten runs per scenario are shown in Fig. 5 and Fig. 6.

Based on the number of packets received per OBU in Fig. 5, the effective data received is shown below in Fig. 7. The MPB method produces a near-constant rate regardless of the vehicle density since there is only one broadcaster in all density scenarios. Code Torrent and GPC show the same pattern of results, which is expected since GPC uses the Code Torrent algorithm, but with one-fourth the number of OBUs contending for the channel. Reducing the number of nodes broadcasting greatly improves the effective throughput to each individual OBU.

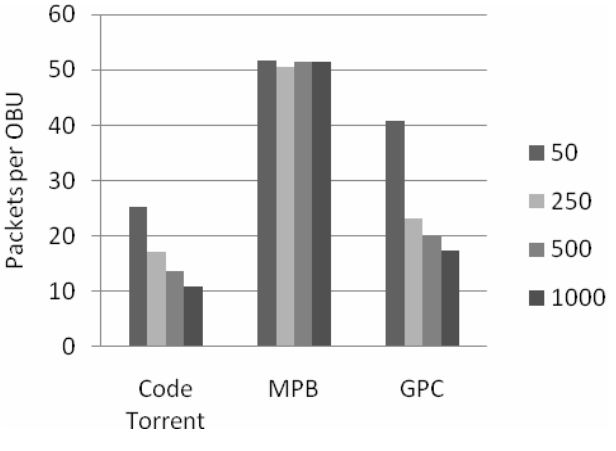

Fig. 5. Average number of packets received per OBU during a 50 ms SCH interval with 50, 250, 500, and 1000 OBUs

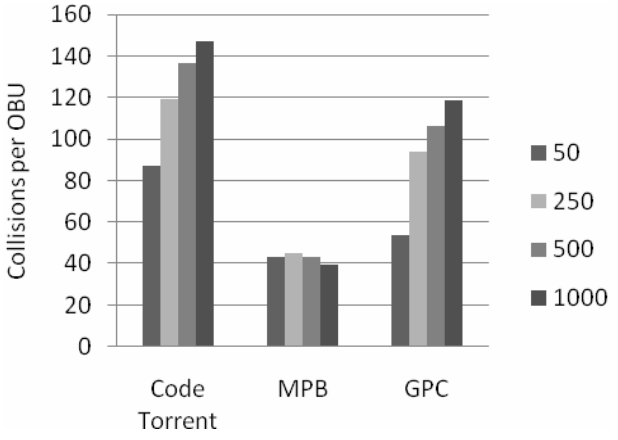

Fig. 6. Average number of collisions per OBU during a 50 ms SCH interval with 50, 250, 500, and 1000 OBUs

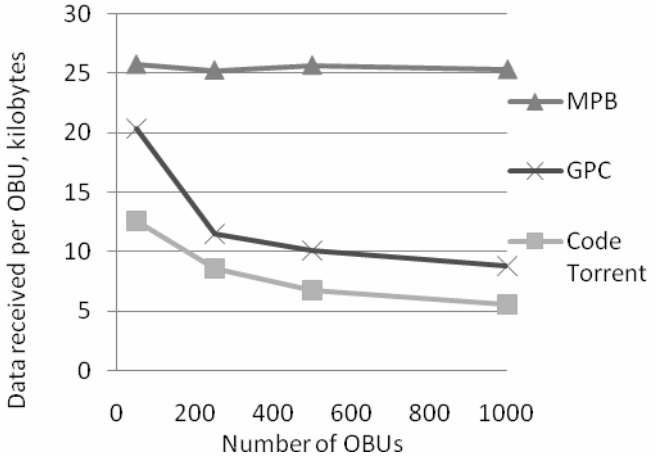

Fig. 7. Average data received per OBU during a 50 ms SCH interval with 50, 250, 500, and 1000 OBUs

9 Conclusion

Reducing the number of broadcasters in a wireless environment greatly improves the efficiency of the channel. Reducing the time needed to send updated CRLs to VANET participants will increase the overall security and safety of the VANET. Using Most Pieces Broadcast (MPB) showed that a near-constant distribution rate can be maintained regardless of the vehicle density. The benefit of MPB becomes greater as the node density increases. The MPB method requires all nodes on the selected service channel to remain silent except for the selected nodes with the most number of

pieces. To achieve this level of cooperation among the nodes would require a dedicated channel to ensure that only nodes following the MPB algorithm were on that channel, otherwise, other nodes would contend for the channel for their own purposes and reduce the effectiveness of MPB.

Code Torrent has many characteristics that make it work well in a VANET environment, but performance degrades sharply as the OBU density increases. Generation per channel (GPC) uses the Code Torrent benefits, but reduces the number of nodes contending for the channel by spreading the nodes out over several of the DSRC service channels.

While MPB and GPC both use vehicle-to-vehicle and vehicle-to-infrastructure communication to distribute CRL files, infrastructure is needed to start the process. CRLs are created by the certificate authorities and distributed to the VANET through the roadside units (RSUs). Determination of which DSRC channels to use for the CRL file distribution should be made by the infrastructure to prevent unnecessary coordination by the OBUs. Once the channel determination is made, both for MPB and the multiple channels used in GPC, those same channels should be used even outside of RSU range.

Future work includes developing simulation models to analyze the entire CRL distribution process to compare the three methods in more detail. From this study, actual CRL distribution times will be determined along with the overhead involved in the three methods. The GPC method will also be studied more in depth to determine if there are benefits to limiting the number of broadcasting nodes on each channel by combining the GPC method with the MPB method. Additionally, the effects of using different access categories for prioritization may be studied.

Acknowledgements

Special thanks to Revathi Balakrishnan for her great assistance with GTNetS, and to Kevin Fairbanks for writing the data extraction script.

References

1. Parno, B., Perrig, A.: Challenges in Securing Vehicular Networks. In: HotNets-IV (2005)
2. IEEE Trial-Use Standard for Wireless Access in Vehicular Environments (Wave) - Security Services for Applications and Management Messages. IEEE Standard 1609.2-2006 (2006)
3. IEEE: IEEE Trial-Use Standard for Wireless Access in Vehicular Environments (Wave) - Multi-Channel Operation. IEEE Std 1609.4-2006 (2006)
4. IEEE: IEEE Standard for Information Technology - Telecommunications and Information Exchange between Systems - Local and Metropolitan Area Networks - Specific Requirements Part 11: Wireless Lan Medium Access Control (Mac) and Physical Layer (Phy) Specifications Amendment 8: Medium Access Control (Mac) Quality of Service Enhancements. IEEE Std 802.11e-2005 (Amendment to IEEE Std 802.11, 1999 Edition (Reaff 2003)) (2005)
5. Unapproved Draft Standard for Information Technology- Telecommunications and Information Exchange between Systems- Local and Metropolitan Area Network- Specific Requirements Part 11: Wireless Lan Medium Access Control (Mac) and Physical Layer (Phy) Specifications. IEEE Unapproved Draft Std P802.11-REVma/D9.0 (2007)

6. Torrent-Moreno, M., Corroy, S., Schmidt-Eisenlohr, F., Hartenstein, H.: IEEE 802.11-Based One-Hop Broadcast Communications: Understanding Transmission Success and Failure under Different Radio Propagation Environments. In: Proceedings of the 9th ACM international symposium on Modeling analysis and simulation of wireless and mobile systems. ACM, Terromolinos (2006)
7. Papadimitratos, P., Mezzour, G., Hubaux, J.-P.: Certificate Revocation List Distribution in Vehicular Communication Systems. In: The Fifth ACM International Workshop on Vehicular Internetworking (VANET 2008), San Francisco, CA, USA (2008)
8. Min, Y., Yuanyuan, Y.: Peer-to-Peer File Sharing Based on Network Coding. In: The 28th International Conference on Distributed Computing Systems, pp. 168–175 (2008)
9. Fujimura, A., Oh, S.Y., Gerla, M.: Network Coding Vs. Erasure Coding: Reliable Multicast in Ad Hoc Networks. In: Military Communications Conference, MILCOM 2008. IEEE, Los Alamitos (2008)
10. Laberteaux, K.P., Haas, J.J., Hu, Y.-C.: Security Certificate Revocation List Distribution for Vanet. In: The Fifth ACM International Workshop on Vehicular Internetworking (VANET 2008), San Francisco, CA, USA (2008)
11. Nandan, A., Das, S., Pau, G., Gerla, M., Sanadidi, M.Y.: Co-Operative Downloading in Vehicular Ad-Hoc Wireless Networks. In: Second Annual Conference on Wireless On-demand Network Systems and Services (2005)
12. Lee, K.C., Seung-Hoon, L., Ryan, C., Uichin, L., Gerla, M.: First Experience with Car Torrent in a Real Vehicular Ad Hoc Network Testbed. In: 2007 Mobile Networking for Vehicular Environments (2007)
13. Lee, U., Park, J.-S., Yeh, J., Pau, G., Gerla, M.: Code Torrent: Content Distribution Using Network Coding in Vanet. In: Proceedings of the 1st International Workshop on Decentralized Resource Sharing in Mobile Computing and Networking. ACM, Los Angeles (2006)
14. GTNetS, Georgia Tech Network Simulator, `http://www.ece.gatech.edu/research/labs/MANIACS/GTNetS/`