

Efficient Secure Aggregation in VANETs

Maxim Raya, Adel Aziz and Jean-Pierre Hubaux

Laboratory for computer Communications and Applications (LCA)
School of Computer and Communication Sciences
EPFL, Switzerland
{maxim.raya, adel.aziz@epfl.ch, jean-pierre.hubaux}@epfl.ch

ABSTRACT

In VANETs, better communication efficiency can be achieved by sacrificing security and vice versa. But VANETs cannot get started without either of them. In this paper, we propose a set of mechanisms that can actually reconcile these two contradictory requirements. The main idea is to use message aggregation and group communication. The first class of solutions is based on asymmetric cryptographic primitives, the second class uses symmetric ones, and the third one mixes the two. We have also evaluated the performance potential of one technique and arrived at the conclusion that aggregation in VANETs increases not only efficiency but also security.

Categories and Subject Descriptors

C.2.0 [Computer-Communication Networks]: General—Security and protection; C.2.1 [Computer-Communication Networks]: Network Architecture and Design—Network communications, Wireless communication.

General Terms

Algorithms, Performance, Security

Keywords

Vehicular networks, Security, Efficiency, Onion signature, Aggregation, Group communication

1. INTRODUCTION

The recent academic and industrial research on VANETs has reached the maturity to consider security as a fundamental building block of any deployable architecture. Several existing works confirm this development. Yet, all VANET security solutions are subject to the same founded criticism: overhead. In fact, the most reasonable choice for a VANET security architecture is a PKI-supported asymmetric authentication, in addition to other functions, such

as anonymity. But in this scheme, every message would have to be signed in order for the receiver to authenticate it. Although cryptographers have greatly improved the efficiency of asymmetric algorithms, notably ECC (Elliptic Curve Cryptography), these still are resource-hungry in terms of computation and communication. This leads us to the obvious question: can VANET security be more efficient? This is the question we will try to answer in this paper.

Most VANET application designers attempt to minimize costs, sometimes even suggesting to scrap security totally. On one hand, this can be understood if we consider that the percentage of attackers will probably be very small. On the other hand, leaving open breaches in huge networks like VANETs can lead to devastating results even if there is only one determined and skillful attacker. This means that both efficiency and security are essential, though seemingly contradictory, conditions for the success of VANETs. The problem we address in this paper is hence finding a tradeoff between the two. This can be achieved by exploiting several properties of VANETs that include geographically constrained paths, vehicle density and high mobility; we will further discuss these properties in a later section.

In this paper, we explore the approach of *secure message aggregation*, the long-time trademark of resource constrained sensor networks. Roughly speaking, instead of letting the de facto flooding approach take care of message dissemination in a VANET, this is delegated only to selected vehicles who share a similar view of their environment. We will describe several algorithms for achieving this and compare them with each other. We will also introduce the concept of *onion signature*, which can be considered the counterpart of onion routing [4]. Relying on realistic simulations, we have come to the conclusion that VANET security can be more efficient when using our aggregation mechanisms.

A useful by-product of secure aggregation is the increase in *information dependability*. In fact, grouping several messages provides the receiver with more evidence concerning a given event. Our simulation results show this effect.

Another aspect that we address is *secure group*¹ *formation*, in itself an open problem in VANET research. Hence we do not claim to provide a complete solution, but rather a feasible option that takes security into consideration.

The paper is organized as follows. Section 2 overviews related work. Section 3 describes the system model and addresses relevant secure group issues. Section 4 presents the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

VANET'06, September 29, 2006, Los Angeles, California, USA.
Copyright 2006 ACM 1-59593-540-1/06/0009 ...\$5.00.

¹In this paper, we use the term *group* in a networking rather than distributed systems sense. Hence, it can be used interchangeably with the term *cluster*.

secure aggregation mechanisms. Section 5 studies one of the proposed techniques using simulations. Section 6 concludes the paper.

2. STATE OF THE ART

The research on VANET security is still developing. Most existing efforts on the industrial [2], as well as the academic [8, 9, 11, 17], side focus on describing the problem statement and proposing the outline of a general solution for VANET security. To provide vehicle authentication, all these works commonly agree on the need for a PKI (Public Key Infrastructure) and the use of digital signatures. Fewer papers focus on specific issues such as the detection and correction of malicious data [5]. The topic of secure aggregation in VANETs has not been addressed so far, except for a brief mention in [9] although it was introduced in a sensor networking sense (e.g., vehicles computing the count of encountered vehicles). Hence, our paper is the first to study in detail this topic in VANETs.

The closest reference in literature to secure aggregation in networks can be found in sensor networking papers. In [7], Hu and Evans propose using delayed aggregation (at the second hop rather than the first) and delayed authentication (by delaying key disclosure) to counter the threat of false data in the network. Their assumptions of a static network with pre-established shared secrets (between sensor nodes and the base station), as well as the key idea of delaying authentication, make their work unsuitable for VANETs. The focus of [10] by Przydatek et al. is also on mitigating the effects of false aggregation results (the so-called *stealthy attack*) by using an *aggregate-commit-prove* mechanism that involves interactive proofs between the aggregators and the home server. Their work also introduces the *efficiency vs. accuracy* tradeoff. But again, the assumption of a static network and the use of interactive protocols hamper the use of their techniques in VANETs. In a similar network setting, Yang et al. [16] introduce secure hop-by-hop aggregation by using *divide-and-conquer* and *commit-and-attest* mechanisms; thus, aggregates can be obtained from multiple subgroups rather than the whole network, reducing the effect of false data injection attacks in some of these subgroups. Last but not least, Wagner [15] also seeks to achieve *approximate integrity* of data through statistical methods, such as outlier elimination. This makes aggregation functions *resilient* to small changes in sensor observations by attackers. This approach can be complementary to the techniques introduced in the following sections, especially to resolve the group agreement problem described in Section 3.4.2.

3. SYSTEM MODEL

In the following, we present several aspects related to the core mechanisms introduced in the next sections. These include relevant VANET properties, geographic routing, group formation, and the attacker model. Finally, we use these elements to describe the problem statement.

3.1 Network Model

In this paper, we address only safety related applications. Each vehicle broadcasts messages to its immediate neighborhood. In addition to vehicles, the network may include roadside base stations but these are not pervasive. All entities are equipped with positioning devices, such as a GPS.

Security provision in VANETs is foreseen mainly by the means of digital signatures. With the existence of a vehicular PKI, each vehicle will possess a set of public/private key pairs that it will use to sign broadcasted safety messages. This ensures that other vehicles will be able to authenticate a received message if it includes a digital signature and the corresponding certificate issued by a CA (Certification Authority). For the sake of comparison, we will dub this mechanism the *basic scheme* throughout the rest of the paper.

3.2 Efficiency-Propitious Properties of VANETs

VANETs consist of large numbers of vehicles moving at high speeds over a continent-size network of roads. Most vehicles are private, which means a lack of a central on-line coordinating entity. All this may look like a nightmare for VANET application designers. But when it comes to the aggregation mechanisms discussed in this paper, these properties turn out to be very helpful. In fact, the higher the density of vehicles, the more accurate the aggregate information. In addition, VANET safety messages are mainly sent to all vehicles in a given geographic region rather than to specific vehicles. In this case, the predefined road topology makes it easier to route these messages. And the mobility of vehicles in both directions can also optimize message delivery.

3.3 Attacker Model

To avoid reinventing the wheel, we refer the reader to other works [9, 11] for a full discussion of the attacker model. In the context of this work, we focus on the assumptions and properties that are directly related to the aggregation mechanisms introduced later.

Similarly to sensor networks [7, 10, 15, 16], the major threat that can target specifically VANET aggregation mechanisms is that of *false information dissemination*. In fact, with a PKI and digital signatures in place, message authentication is not a direct issue here. Also, availability problems (due to jamming) are not aggravated and can actually be alleviated by aggregation due to the reduction of channel congestion. But the fact that aggregation reduces the number of messages (and not the amount of information) can allow cheaters to insert false data into the network. Therefore, we have to make the following single assumption:

Any group of vehicles should contain a majority of honest nodes under normal density conditions.

The definition of groups will follow shortly in Section 3.4. Normal density conditions refer to typical scenarios on roads: vehicles driving within at most few tens of meters of each other. This assumption allows us to rely on the existence of honest group members able to rectify the false data disseminated by attackers. This is also in line with the data correctness requirement introduced in Section 3.5.

3.4 Group Aspects

Our algorithms revolve around the core idea of information relating between groups of vehicles rather than individual vehicles. This, of course, does not concern the physical transmission of data but the data flow in the network. More precisely, vehicles are arranged into groups. Within each group, one or more vehicles, automatically determined by their positions, transmit the data aggregated in that group to neighboring groups. This is illustrated in Figure 1.

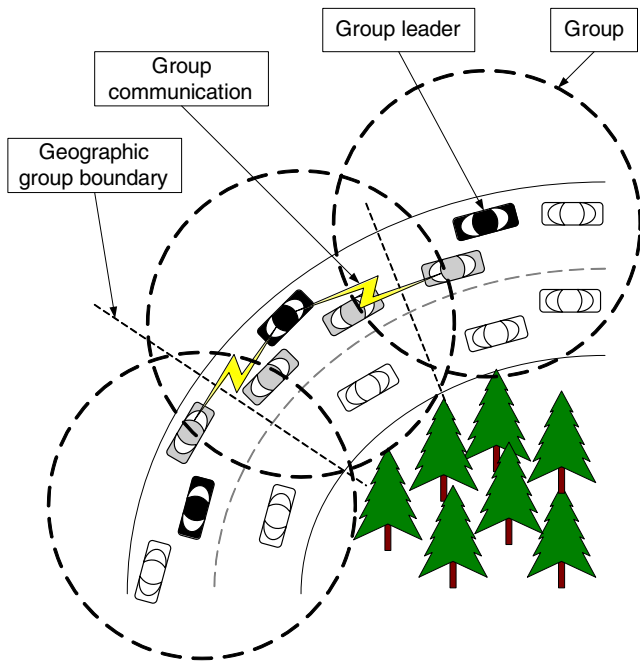


Figure 1: Efficient aggregation by means of overlapping groups. Communication between the two outer groups is possible because at least the leader of the center group is in reach of relaying vehicles (in grey) in both outer groups.

The area of group formation and management is one of the most important and at the same time complicated topics in VANET research. Groups have many intuitive applications in VANET settings, especially platooning-like applications [12]. From the security standpoint, a recent work [14] has also suggested using groups to increase the anonymity of vehicle to infrastructure communications. But there are two major problems that need to be tackled when addressing group aspects in VANETs: *group formation* and *intra-group agreement*. In this work we will focus on the first and due to the lack of space we will give only some hints concerning the second problem.

3.4.1 Group Formation

There can be many ways to form groups in VANET applications. For example, all public transport buses can be members of a *preset group*. This is the easiest and most efficient way of group formation, but it requires prior knowledge of group members, as well as a common authority over them. This is not the case when individual drivers on a highway decide to join a platoon in order to improve their driving experience. This necessitates *on-the-fly group* formation where a group leader² is elected and group membership is managed dynamically. This latter category of groups is the most useful functionally due to its flexibility, but it is also the most difficult to form due to a multitude of issues, such as group leader election, group overlap (e.g., how to decide which group to join if a vehicle is within the boundaries of two overlapping groups), and the related security hurdles.

²A *group leader* can also be called a *clusterhead*.

In order to escape the rigidity of preset groups and the complexity of on-the-fly groups while retaining, at least partially, the efficiency of the first and the flexibility of the latter, we have sought a hybrid solution. The result is *location-based groups*. In fact, for safety applications, which are the focus of this paper, it is essential to know *where*, and not *who*, the neighbors of a vehicle are. As mentioned earlier, messages are mostly destined to geographic regions rather than individual vehicles. For example, if there is sliding terrain behind a curve, all vehicles entering the curve should be informed. Hence the intuitive idea of sending messages from groups of vehicles in one location to groups of vehicles in another. This brings us to the group formation primitive we use in this paper: the map (more precisely, the roads) is dissected into small area cells that actually define the groups. A vehicle will automatically know to which group it belongs by comparing its GPS position to a preloaded dissection of the area map into cells. The group leader, the vehicle closest to the center of the cell, is determined dynamically. Cells, and hence groups, overlap in such a way that any vehicle moving from one cell to the next remains in transmission range of both group leaders. This means that the cell size depends on the transmission range of vehicles. Using the typical DSRC (Dedicated Short Range Communications) [1] range of 300 m, we have set the cell length in our simulations to 400 m, which proved to be a suitable value. Further improvements on cell size calculations could be possible, which we leave to future work. Figure 1 illustrates this concept, as well as some details.

A seemingly difficult - but in fact straightforward - process associated with location-based group formation is that of group leader election. As mentioned in the previous paragraph, the group leader is the vehicle closest to the group's cell center. Because cells are predetermined, the center location is also known to all vehicles in the cell. In addition, by leveraging on the periodic safety message broadcasts (at most each 300 ms [1]) that include a vehicle's position, each vehicle is aware of the positions of its neighbors within a tolerable margin of error (few meters) due to the imprecision of GPS. Thus, a group leader election takes place within a delay of at most 300 ms. If there are several vehicles close enough to the center such that the error margin does not allow a clear-cut decision, the vehicle with the lowest ID among these will be elected as group leader. We should note here that vehicles do not broadcast their actual IDs but rather pseudonyms for privacy purposes.

By using location-based groups, we can reap two major benefits:

- **Efficiency:** A vehicle will automatically know to which group it belongs. Hence, group formation will not require any additional communication overhead or delay.
- **Routing:** As most routing in safety applications is geographic, determining which groups should relay messages is straightforward.

To achieve the above advantages, almost the only costs involved in this type of group formation is the preloading of map dissections into vehicles. But this can be easily included on the vehicle navigation maps that will probably be an integral part of each vehicle when VANET communications hit the market.

3.4.2 Group Agreement

In order for information to be generated and propagated by groups rather than individual vehicles, all vehicles in a group should share a similar view of their environment. Any kind of group agreement protocol would be expensive in terms of communication overhead and delay, without mentioning security. Hence we adopt a simpler yet effective approach: each vehicle *locally* processes all events, either directly observed or reported by other vehicles, before making a decision concerning that event. By using this approach, we make the following reasonable assumptions:

- Most vehicles in one cell receive messages with similar information from other cells. This would be the case if the cell size is comparable to the transmission range (the respective values that we use in this paper are 400 m and 300 m).
- Under normal traffic density (defined in Section 3.3), any event happening in a cell is observed by several vehicles. This means that there are alternative sources of information that can be crosschecked for consistency verification.
- Most honest vehicles observing the same event report similar observations. Possible errors are due only to differences in on-board sensor readings.

In addition, we leverage on the assumption that attackers are a minority among network members, which means that there is a majority of correct observers.

Under the above conditions, the majority of vehicles in a cell share a similar view of the environment within tolerable margins of error. We leave the details of this algorithm for future work - due to the lack of space. A good existing example of such an approach can be found in [5].

3.4.3 Secure VANET Group Protocol

We end the discussion on group aspects by introducing a simple protocol for symmetric group key establishment in VANETs: SVGP (Secure VANET Group Protocol). It is inspired notably from the Group Key Management Protocol (GKMP) [6] but with geographically defined groups. This protocol will be later used in Sections 4.2 and 4.3.

As explained before, roads are divided into cells that define groups, with the group leader being the vehicle closest to the cell center. Leveraging on periodic broadcasts of certified public keys, the group leader (L in this example) distributes the group key K to members A , B , and C encrypted with their respective public keys as follows:

$$L \rightarrow * : \{K\}_{PuK_A}, \{K\}_{PuK_B}, \{K\}_{PuK_C}, \\ Sig_{PrK_L}[\{K\}_{PuK_A} | \{K\}_{PuK_B} | \{K\}_{PuK_C}]$$

Subsequent message broadcasts will include only a HMAC in addition to the message itself:

$$L \rightarrow * : m, HMAC_K(m)$$

When a new vehicle D enters the cell, it receives the group key from the current group leader:

$$L \rightarrow D : \{K\}_{PuK_D}, Sig_{PrK_L}[\{K\}_{PuK_D}]$$

When a vehicle leaves the cell, nothing needs to be done. In fact, the creation of secure groups will only contribute to reducing the security overhead and not defining different security levels among VANET members. Similarly to digital signatures, the use of secure groups protects the network from outsiders (entities that do not possess certified public/private key pairs). Hence, while renewing or transferring existing keys during member joins is still necessary, member leaves should not necessarily entail an update of the group key.

Special attention needs to be paid to exchanges on cell boundaries when a vehicle switches from one group to another. In order to make this operation smooth, cell dimensions should be smaller than the diameter of the transmission range disk. For example, if the transmission range is 300 m (the disk diameter is 600 m), we can choose a cell size of 400 m. Hence, at the cell boundaries, a vehicle will receive messages from the leaders of both its previous and new groups.

There are several other functional details of this protocol that need to be worked out. But our purpose is to make it as simple as possible in order to demonstrate its usage in secure message aggregation. Hence, we describe this protocol only as an example; other mechanisms are also possible. It is also important to note that shared group keys are limited in space and time due to the small fixed cells and high vehicle mobility. This makes the vulnerability windows, opened by the compromise of these keys, small.

3.5 Efficient Security by Aggregation

The basic scheme described in Section 3.1 actually summarizes the way VANET messaging security is envisioned thus far. Its main advantage is simplicity and the inherent robustness, which is very important for critical safety applications. But the overhead generated by this simple scheme leaves ample space for improvement. In fact, there is currently a gap between efficiency and security in VANETs, with one of the two usually achieved at the expense of the other. In this work, we make an attempt at bridging this gap by exploiting several properties of VANETs and proposing fundamentally different solutions. Moreover, we will try to keep a level of security at least equivalent to the one provided by the basic scheme. The main focus of this paper will be *message aggregation* and *group communication* as sketched in Figure 1. As Section 5 will show, our objectives are achievable.

Before attacking the solution space, it is important to keep in mind the requirements of any viable solution. A full list of such properties can be found in several papers [9, 11]. In the following, we list the main properties we seek to achieve in our solutions, the others being orthogonal to the proposed mechanisms.

- **Channel Efficiency:** The security overhead due to signing safety messages is high. In fact, the security material can even sometimes exceed the rest of the message payload [11]. This in turn results in channel congestion that hinders the transmission of safety messages. Despite this drawback, security cannot be totally overlooked because a non-secure network is more dangerous than an overloaded network, especially when it comes to life-critical applications. Hence, a primary objective of our work is to increase channel efficiency by reducing the security overhead.

- **Low Delay:** Another important VANET parameter to be strictly respected is low delay. Because an accident can be avoided within fractions of a second, reducing the safety message delivery delay is another priority.
- **Data Correctness:** Data received by vehicles should be checked for correctness. Otherwise, cheaters will be able to disseminate false information in the network, thus endangering its safety. A typical and simple mechanism for ensuring data correctness is by cross-checking several sources of information. It is interesting to observe that aggregation mechanisms inherently provide this property, as we will explain later.
- **Non-repudiation:** Last but not least is the non-repudiation property, or the capability to trace back attackers even after the attacks take place. This creates a prevention mechanism against potential cheaters because they can be pursued by justice.

4. SECURE AGGREGATION MECHANISMS

In this section we present three major classes of aggregation techniques, each one representing a different way to achieve better efficiency. The main feature of these techniques is that only selected vehicles (the relaying vehicles in Figure 1) take care of message aggregation and dissemination. Efficiency is achieved because vehicles overhearing the source of the aggregated message will restrain from broadcasting messages with similar information.

4.1 Combined Signatures

In the basic scheme, each vehicle broadcasts a signed safety message. This creates considerable security overhead, especially in terms of message size and signature generation delay. There may be an additional delay resulting from data verification algorithms running on the receiving vehicles. One of the ways to cope with false data in the network is to crosscheck data concerning an event by comparing messages received from several sources. The downside of data crosschecking is that a vehicle needs to wait for several messages with information on the same event, which creates a certain delay.

Bearing in mind the above drawbacks of independent safety messages, we have sought to combine the signatures generated by a group of vehicles reporting the same event. Thus, all the overhead will be grouped in one message instead of being spread over several, resulting in a more efficient channel usage. In addition, once a vehicle receives such a combined message, it can skip the data verification process because the combined signatures imply that all the involved signers agree on the content of the message. There are several types of signature combinations, each with its own benefits and drawbacks, especially in terms of overhead. The formats of the three types of combined signatures are shown in Figure 2. It should be noted that this aggregation technique makes use of only asymmetric cryptography, hence the need for including the public key certificates of all signers in the corresponding message.

4.1.1 Concatenated Signatures

In this case, a vehicle that receives a message with correct information (from the receiver’s perspective) appends




Figure 2: Three different types of combined signatures. n is the total number of signers.

its signature to the existing signatures and rebroadcasts the new message. As signatures are appended to the message independently of each other, they will also be verified independently. Hence, there is no need for signers to verify the other signatures. Thus, concatenating signatures generates the same security overhead in terms of signature size and generation as the basic scheme. But its network overhead is considerably smaller as the results in Section 5 show. It also overcomes the basic scheme in terms of data verification delay because a destination vehicle receives data explicitly approved by several signers. This form of *source aggregation* (supporting data is gathered at the source) results in smaller data verification delay than *destination aggregation* (the receiver collects messages from different sources and then crosschecks them), as we will show in Section 5.4. Another advantage of signature concatenation is that an invalid signature does not affect the whole message, in contrast to the next scheme.

4.1.2 Onion Signatures

A key challenge in this paper was to reduce the signature size. Although this is not possible for a single message under a given cryptosystem, we have achieved the reduction of the total overhead associated with given information (usually expressed in several messages with similar content). Exploiting the fact that signature sizes are constant because a message is hashed before being signed, we explored the possibility of *oversigning* a message. Instead of simply appending its signature, a vehicle signs the signature of the previous transmitter. Before retransmitting the new message, it should also include the last signature, i.e., the one it received, so that the vehicle at the next hop can verify the new signature. With this approach, no matter how many times a message is oversigned, the ultimate result will always be the safety message with two signatures (the new and the previous ones). Since this technique is similar to the message reencryption process in onion routing [4], we coined it *onion signature*.

Obviously, the drastic improvement in signature size comes at a cost. As each vehicle contributes to the combined signature by adding an onion layer signature (i.e., oversigning), an invalid signature at any layer invalidates the final combined signature. To reduce the processing costs of signature verification at the receiver, each signer has to verify the last signature before oversigning. If the signature is invalid, it has to discard the existing signature and restart the onion signature generation. Obviously, verifying signatures at each hop increases the overall computation overhead, in addition to delaying the delivery of the combined signature to the destination. Moreover, a false message with a valid

signature cannot be detected directly. But this attack can be thwarted by the possibility of punishment as the non-repudiation property of digital signatures allows the CA to determine the attacker.

4.1.3 Hybrid Signatures

Although concatenated signatures, in terms of computation overhead, are more efficient than onion signatures, the opposite is true with respect to communication overhead. This gives several choices in terms of the hardware and software platforms to use. To accommodate an even wider spectrum of equipment configurations, a hybrid solution that combines features of both approaches is possible. A *hybrid signature* would consist of several concatenated onion signatures, each of a given depth. The signature depth represents the number of layers it includes. The number and depth of onion signatures can be varied according to the targeted communication and computation overheads.

4.2 Overlapping Groups

Despite the advantages achieved by the schemes in the previous section, they are still in the realm of asymmetric cryptography, which remains expensive. To make a quantitative leap, we designed a mechanism based on symmetric cryptography. The core idea is to use *overlapping groups*, each group having its own symmetric key. Vehicles in the intersection of two groups know the keys of both groups and hence are able to assure the bridge for data flow between the two groups. In order to further clarify this mechanism, we need several assumptions and tools, some of which have been introduced in Section 3:

- The position of a vehicle can be securely determined or verified. This implies the need for a secure position verification primitive.
- Groups are established with each group having its own symmetric key. We have presented in Section 3.4 the different aspects related to this issue, namely group formation, agreement, and key establishment.
- There is a majority of honest vehicles, upon which data destination vehicles can rely to receive correct information. Because all communication between groups uses symmetric mechanisms, the non-repudiation property is lost. Although cheating attackers cannot be discouraged by the possibility of later punishment, their effect will still be countered by the honest majority.

Figure 1 illustrates how information flows between overlapping groups. Because groups are geographically predefined, vehicles in the intersection of two groups are aware of their status as relaying nodes. The SVGP protocol introduced in Section 3.4.3 provides them with the keys of both groups.

Obviously, the main advantage of data flow between overlapping groups is the reduced communication and computation overhead due to symmetric cryptography. The drawbacks are the need for secure position verification, the overhead of group aspects, and the loss of the non-repudiation property. Secure position verification will be needed anyway in VANETs and the lack of non-repudiation can be countered by the presence of an honest majority, thus overlapping groups can still be a viable solution to build secure and efficient VANETs.

4.3 Dynamic Group Key Creation

A spontaneous question that stems from the above two mechanisms is: Can we achieve the low overhead of symmetric cryptography without losing the non-repudiation property of digital signatures? The answer is *yes* and the means is *dynamic group key creation*. Before delving into the details of this mechanism, we have to make one assumption: the existence of sporadic infrastructure coverage. Although still on the to-do list, this infrastructure coverage will probably soon become reality.

In contrast to overlapping groups where memberships are geographically determined, in this section we consider dynamic groups created by vehicles sharing the same driving pattern over an extended period of time. A good example of such groups is a platoon of vehicles formed on the highway. We leave the exact mechanism of group formation out of scope of this work because it highly depends on the specific VANET application and is not directly related to our main subject.

The key idea is that once the leader and members of the group are identified, the leader creates a *key request* message with the format and content as illustrated in Figure 3. And once in the range of a roadside base station, the leader transmits the message to the CA through the base station (Figure 4). The CA will use the information it receives to generate an asymmetric group key pair and broadcast it to all the group members. The key pair will be encrypted with the symmetric group key included in the key request message (this key can be generated using SVGP described in Section 3.4.3). In addition, the CA assigns to each group member a unique ID for non-repudiation purposes.

Once the asymmetric group key is established, any group member can send a message signed on behalf of the group. This implies that a group agreement process (Section 3.4.2) has taken place before the message is sent. As with the digital signatures in the basic scheme, the digital signature of the group is accompanied by its certificate issued by the CA in order to allow the receivers to verify the signature. The message also includes the unique ID assigned by the CA to the group member that sent the message. Normally, this group member is the vehicle closest to the destination.

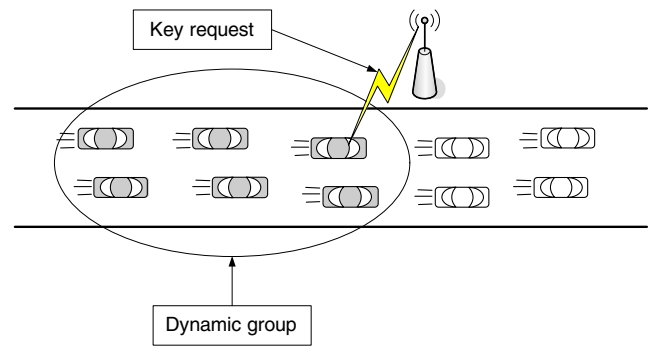


Figure 4: Dynamic group key creation.

As in the previous mechanism, the case of a cheating group member can be mitigated by making more than one vehicle send the message and relying on the presence of an honest majority. In contrast to the overlapping groups, dynamic group creation conserves the non-repudiation property. The

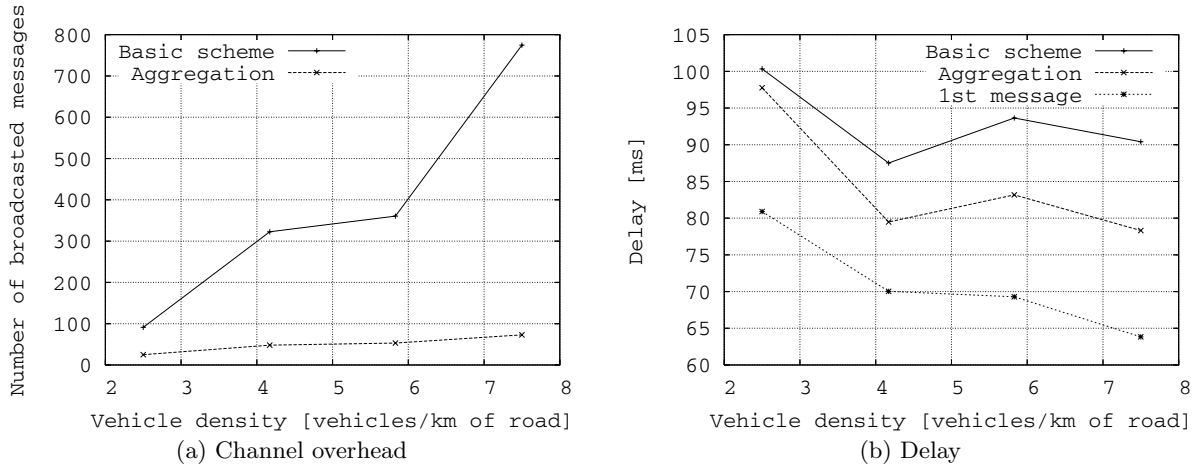


Figure 6: Effects of vehicle density on the performance of message aggregation.

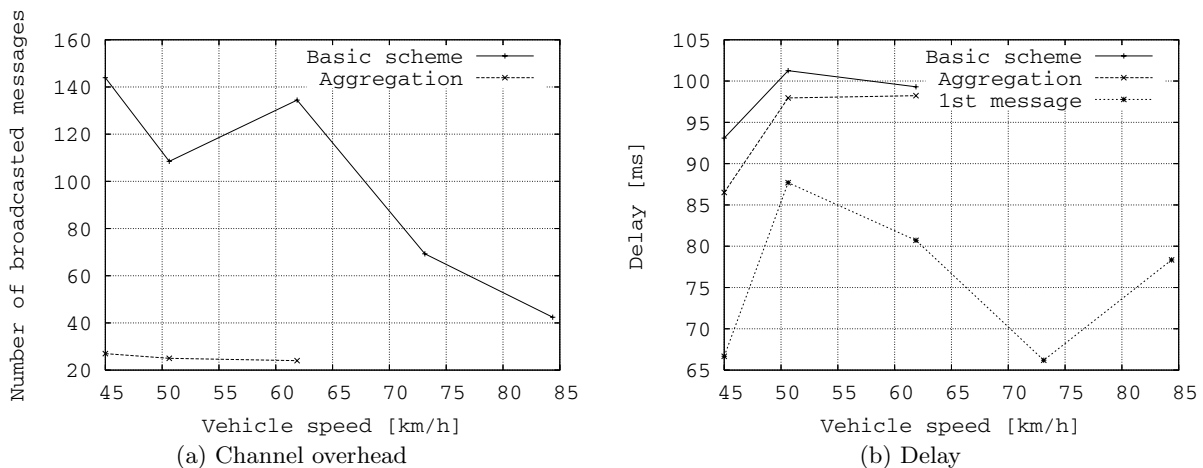


Figure 7: Effects of vehicle speed on the performance of message aggregation.

For message delivery delay, it is interesting to observe in Figure 6(b) that aggregation also results in a smaller delay than the basic scheme. The explanation of this effect is twofold:

1. Although vehicles in the destination area receive the first message earlier under the basic scheme, they still have to accumulate the necessary number of messages to reach the proper correctness level. Hence, the aggregated message is received before the last required message in the basic scheme; the reason is explained in the next point.
2. Because the number of broadcasted messages is considerably smaller in the aggregation scenario, there is much less congestion on the wireless channel. This in turn favors faster delivery of messages.

In addition to the basic and aggregation schemes, Figure 6(b) illustrates the delay achieved when the correctness level is set to 1, i.e., when the first received message is considered valid. We can see that boosting the correctness level (by

4 times) results in a tolerable additional delay (around 10 ms). This suggests that data verification can be done at affordable costs.

5.3 Effects of Vehicle Speed

We have also explored the effects of vehicle speed (Figures 7(a) and 7(b)) with a density of 2.5 vehicles/km of road. The reason for choosing this value is that it represents the worst case scenario according to the results of the previous section; the performance gap between the aggregation and basic schemes increases at higher densities in favor of the first.

According to Figure 7(a), channel overhead decreases with increasing speed. Although this seems to be surprising, it can be explained by the fact that at higher speeds a smaller fraction of messages is delivered, resulting in smaller effective channel usage (lost messages are not represented on the graphs).

As for the delay, Figure 7(b) shows high variability with speed. These fluctuations can be explained as above by the varying number of successfully received messages. In fact,

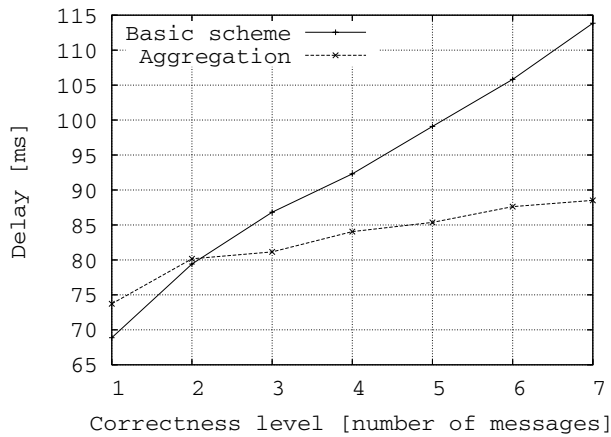


Figure 8: Efficiency vs. Security

vehicle mobility can speed up message delivery but can also result in higher message losses. We can see that at high speeds, the correctness level could not be achieved in both the basic and the aggregation schemes.

These results suggest that further investigations should be done on the effects of mobility in VANETs. But in the context of our study, both graphs show, as with vehicle density, that aggregation performs better than the basic scheme in terms of channel overhead and message delivery delay.

5.4 Effects of Correctness Level

Last but not least, we examined how changing the correctness level affects the delay of message delivery. As information correctness affects the security of the network, the resulting graph represents, in a sense, the *Efficiency vs. Security tradeoff*. And not surprisingly, Figure 8 shows that the higher the correctness level, the higher the delivery delay. Moreover, aggregation performs far better than the basic scheme when the correctness level increases. It is also interesting to observe that the increase in the delay is rather slow when using aggregation. This leads to the conclusion that better data assurance can be efficiently achieved through aggregation.

6. CONCLUSION

In this paper, we have addressed the tradeoff between efficiency and security in VANETs. We have proposed several mechanisms, including combined signatures, overlapping groups, and dynamic group key creation. We have evaluated the performance of one type of combined signatures, namely concatenated signatures. The results confirm that secure aggregation is a promising approach for increasing the channel efficiency and decreasing message delivery delay in VANETs. Moreover, we have come to the interesting conclusion that aggregation contributes to better data correctness and, in some sense, a higher level of security.

Our plans for future work include exploring other aggregation techniques, as well as different types of efficient mechanisms. We also intend to study the effect of mobility on VANET performance because this seems to be an intriguing topic.

7. REFERENCES

- [1] 5.9 GHz DSRC. <http://grouper.ieee.org/groups/scc32/dsrc/>.
- [2] IEEE P1609.2 Version 1 - Standard for Wireless Access in Vehicular Environments - Security Services for Applications and Management Messages.
- [3] T. Aura. Cryptographically Generated Addresses (CGA). In *Proceedings of ISC'03*, 2003.
- [4] D. Goldschlag, M. Reed, and P. Syverson. Onion routing. *Communications of the ACM*, 42(2):39–41, 1999.
- [5] P. Golle, D. Greene, and J. Staddon. Detecting and correcting malicious data in VANETs. In *Proceedings of VANET'04*, 2004.
- [6] H. Harney and C. Muckenhirn. Group Key Management Protocol (GKMP) architecture. RFC 2094, 1997.
- [7] L. Hu and D. Evans. Secure aggregation for wireless networks. In *Proceedings of SAINT'03*, 2003.
- [8] J.-P. Hubaux, S. Capkun, and J. Luo. The security and privacy of smart vehicles. *IEEE Security and Privacy Magazine*, 2(3):49–55, May–June 2004.
- [9] B. Parno and A. Perrig. Challenges in securing vehicular networks. In *Proceedings of HotNets-IV*, 2005.
- [10] B. Przydatek, D. Song, and A. Perrig. SIA: secure information aggregation in sensor networks. In *Proceedings of SenSys'03*, 2003.
- [11] M. Raya and J.-P. Hubaux. The security of vehicular ad hoc networks. In *Proceedings of SASN'05*, 2005.
- [12] H.-J. Reumerman, M. Roggero, and M. Ruffini. The application-based clustering concept and requirements for intervehicle networks. *IEEE Communications Magazine*, 43(4):108–113, April 2005.
- [13] A. K. Saha and D. B. Johnson. Modeling mobility for vehicular ad-hoc networks. In *Proceedings of VANET'04*, 2004.
- [14] K. Sampigethaya, L. Huang, M. Li, R. Poovendran, K. Matsuura, and K. Sezaki. CARAVAN: providing location privacy for VANET. In *Proceedings of escar'05*, 2005.
- [15] D. Wagner. Resilient aggregation in sensor networks. In *Proceedings of SASN'04*, 2004.
- [16] Y. Yang, X. Wang, S. Zhu, and G. Cao. SDAP: a secure hop-by-hop data aggregation protocol for sensor networks. In *Proceedings of Mobihoc'06*, 2006.
- [17] M. El Zarki, S. Mehrotra, G. Tsudik, and N. Venkatasubramanian. Security issues in a future vehicular network. In *Proceedings of European Wireless*, 2002.