Resilient Secure Aggregation for Vehicular Networks

Stefan Dietzel, Elmar Schoch, Bastian Könings, and Michael Weber, Ulm University
Frank Kargl, University of Twente

Abstract

Innovative ways to use ad hoc networking between vehicles are an active research topic and numerous proposals have been made for applications that make use of it. Due to the bandwidth-limited wireless communication medium, scalability is one crucial factor for the success of these future protocols. Data aggregation is one solution to accomplish such scalability. The goal of aggregation is to semantically combine information and only disseminate this combined information in larger regions. However, the integrity of aggregated information cannot be easily verified anymore. Thus, attacks are possible resulting in lower user acceptance of applications using aggregation or, even worse, in accidents due to false information crafted by a malicious user. Therefore, it is necessary to design novel mechanisms to protect aggregation techniques. However, high vehicle mobility, as well as tight bandwidth constraints, pose strong requirements on the efficiency of such mechanisms. We present new security mechanisms for semantic data aggregation that are suitable for use in vehicular ad hoc networks. Resilience against both malicious users of the system and wrong information due to faulty sensors are taken into consideration. The presented mechanisms are evaluated with respect to their bandwidth overhead and their effectiveness against possible attacks.

Wireless communication between vehicles — often termed vehicular ad hoc networking (VANET) — is assumed to contribute to a safer, more efficient, and more comfortable driving experience in the foreseeable future. One use case that appears in many research efforts is the dissemination of status information in a certain region, e.g., 5 km, surrounding the vehicles. An example is the co-operative detection of traffic jams achieved by vehicles exchanging information about their current position and speed. In a very simple implementation, each vehicle could periodically broadcast such messages using multi-hop flooding.

Scalability Through Aggregation

The major drawback of this periodic flooding approach is that it does not scale well. Consider a simplified example of a traffic jam on a three-lane highway with 20 m distance between vehicles. That means a total of 750 vehicles on a 5 km highway segment. Assuming a theoretic transmission range of 200 meters and a message generation rate of 1/sec, every vehicle on these 5 km is expected to receive $3 \cdot 2 \cdot 200\text{m}/20\text{m} \cdot 750 = 45{,}000$ packets per second. If each packet is 100 bytes in length (including all headers), this consumes a total of 36 Mb/s, not considering MAC delays, bandwidth demands of other applications, and broadcast collisions due to the multi-hop dissemination. Overall, the application will over-saturate the communication medium and cannot be implemented that way.

Several approaches have been proposed to mitigate this problem, including transmit power control, dynamic reduction of transmission frequency, and efficient Geocast protocols. For applications like traffic jam detection, another suitable solution is in-network data aggregation.

The idea of aggregation is simple: instead of disseminating individual messages, vehicles compare messages that they receive with their own information base, semantically combining messages where possible. Thus, after collaborative detection of a traffic jam in this way, only one message suffices to convey the information about an arbitrarily long traffic jam. Compared to the dissemination of several hundreds of messages from all vehicles in the traffic jam, this is a huge bandwidth saving.

Security is an Unresolved Problem

Several examples [1, 2] discuss the benefits and scalability of such aggregation schemes and propose solutions specifically tailored to the needs of VANETs. However, introducing aggregation also opens new opportunities for attackers to disrupt normal operation. In the above-mentioned example, an attacker could misuse the aggregation scheme to report a congestion on the next 10 km of the road, claiming that he received reports from hundreds of vehicles that contributed to the aggregate value.

As we will discuss later, vehicular networks pose strict requirements on security and existing approaches cannot be easily applied. We then introduce a generic aggregation model that serves as a basis for further considerations. We present two security mechanisms that can detect attacks and mitigate their effects. We evaluate the protection level and the performance that these mechanisms can achieve. We then conclude with a summary and outlook on our ongoing work.

0890-8044/10/$25.00 © 2010 IEEE

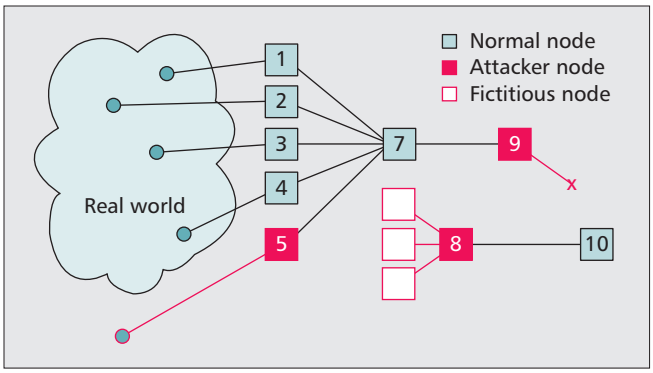

Figure 1. *Schematic overview of different attacks influencing aggregation. Node 5 forges atomic reports, node 8 assembles entirely forged aggregates, and node 9 drops aggregates selectively.*

Security Issues

Difficulties for the security of aggregation arise directly from the advantage of aggregation: messages carry information with higher semantic value and less redundancy than reports from a single vehicle. While atomic reports only claim a single fact (e.g., one vehicle's speed), aggregates contain the combined view of many vehicles in a larger area. Therefore, the process of aggregation and dissemination should be protected against malicious attacks, which we discuss in the following.

Adversary Model and Goals

We consider an adversary to be a single entity which may control several stations in a certain area of the network. The attacker stations can actively participate in the communication, that is, they can send and receive any message conforming to the aggregation protocol. Focusing only on the aggregation process, we can identify the attacks depicted in Fig. 1:

- *Forging of atomic reports*. An attacker station may forge its own speed and thus influence further aggregation (node 5). As a result, aggregates may be biased.
- *Forging of aggregates*. Instead of just influencing aggregates, an attacker may directly create aggregates with arbitrary data and inject them into the network (node 8).
- *Suppression of aggregates*. Because of the larger information value of aggregates, attacker stations may suppress aggregates, resulting in biased information dissemination (node 9).

While suppression of aggregates and forging of atomic reports certainly influence aggregation, the most effective attack on aggregation is the creation of entirely fictitious aggregates, because such aggregates can pretend to carry information about arbitrary dimensions and values. For example, one single aggregate may claim a congestion in a very large area, causing all other vehicles to avoid that area. Therefore, the major goal for secure aggregation is to safeguard the aggregation process in order to prevent attackers from forging aggregates.

Constraints

Due to the specific characteristics of vehicular networks, protecting aggregation is not trivial. High mobility causes frequent topology changes, that is, nodes constantly encounter new neighbors and leave others behind. Therefore, security mechanisms should not assume stable structures such as a static node topology. These high dynamics together with the bandwidth limited wireless channel constrain the usage of interactive solutions. Multi-round protocols that involve the exchange of multiple messages to convince neighbors that an aggregate was correctly merged may fail, because mutual reachability may be very short. Therefore, re-use of existing protocols for wireless sensor networks [3] is not possible.

Secure aggregation without multi-round communication has, for example, been considered by Picconi et al. [4]. The proposed system uses random checks similar to interactive attestation in WSNs. However, a *tamper-proof service* in each vehicle acts as a proxy for the receiver. Thus, multi-round communication over the wireless channel is avoided. An aggregator passes each aggregate to be disseminated to this service, which requests a randomly chosen original record to prove its integrity. Only if the original record is valid, it will be added to the newly created aggregate record as a *proof item* for other vehicles. This method allows to achieve a probabilistic verification of aggregates without a large communication overhead. The main disadvantage of this approach is the dependency on the tamper-proof service. An attacker can easily bypass this service and compose malicious aggregates including a valid proof item.

Node reputation management is also not a viable solution to safeguard self-organized networks [5]. To build node reputation for semantic aggregation in particular, each aggregate would need to be accompanied by a list of all participating nodes' identities. However, this list would lead to exponentially increasing aggregate sizes, thereby destroying the bandwidth savings of aggregation. Even if such a list was available, derivation of reputation for individual nodes would be difficult because it is hard to determine which one of the contributors maliciously modified the aggregate at an earlier stage of the aggregation process.

Another approach for secure aggregation is presented by Raya et al. [6]. Here, streets are divided in segments of static size corresponding to the coverage of wifi signals. However, this grouping approach implies an underlying aggregation scheme employing a fixed segmentation of the road. It has been shown [7] that this type of aggregation cannot scale well with a large number of vehicles and larger areas, for example, long traffic jams spanning over several kilometers.

Further, it is assumed that all group members share the same view of their environment, which is calculated in a group agreement process. Securing data aggregation is then achieved by hybrid signatures to achieve a trade-off between computation efficiency and bandwidth efficiency. The assumption that group members agree on one common view of their environment means that either a multi-round scheme is employed, which is infeasible due to high node mobility. Or, a quantization is applied, for example, speed might be reported only with 10 km/h granularity. However, this would unnecessarily reduce the information value of the aggregates. Further, the assumption that all vehicles in a segment drive at a roughly common speed could be wrong considering different vehicle types, e.g., fast cars and trucks.

We will therefore introduce a security mechanism that does neither rely on agreeing to common views nor node reputation and also works with dynamic road partitioning. For this, we will further develop the idea of probabilistic attestation, removing the need for a tamper-proof device in each car.

Generic Aggregation Model

Before discussing our proposed security mechanisms, we will now introduce a generic model for the underlying aggregation that is applicable to a wide range of applications like traffic information systems, icy road warning, counting of parking spots, fog warnings, and so forth. For this, we consider the following generic aggregate structure:

$$A = [\underbrace{(a_1,b_1),\ldots,(a_n,b_n)}_{\text{index dimensions}} | \underbrace{(v_1,\ldots,v_p)}_{\text{values}} | \underbrace{(m_1,\ldots,m_p)}_{\text{meta-information}}].$$

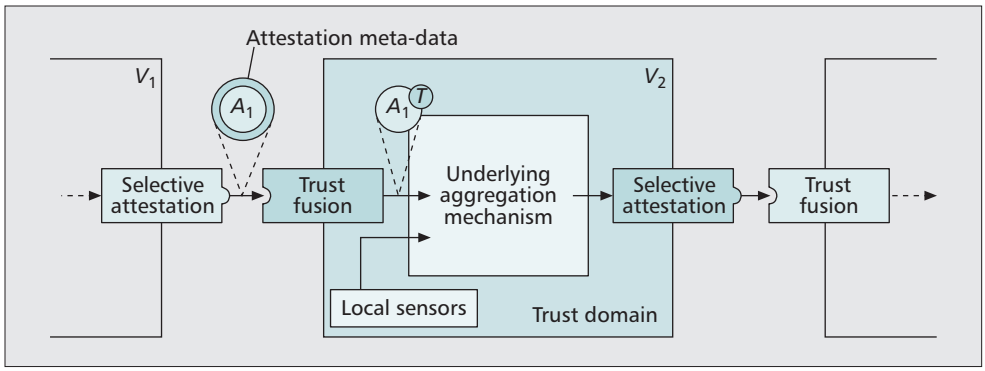

Figure 2. *Overview of the system model showing the flow of aggregates and atomic reports through our system.*

The index dimensions indicate the area, and possibly time, about which an aggregate contains information. The values are the actual information contained, e.g., average speed, minimum temperature, number of parking lots, or range of vision.

The meta-information items contain all additional information used to verify the aggregate's correctness according to our security mechanisms. The same notation can be used to express signed atomic reports by setting $a_i = b_i$ and m_1 to be the signature over the aggregate's content. Whenever atomic reports are combined to a new aggregate — or when existing aggregates are further combined due to hierarchical aggregation mechanisms — an aggregation function calculates the new values of the resulting aggregate, e.g., the new speed average based on the speed values contained in the existing aggregates and atomic reports. Moreover, the index dimension intervals are adjusted to reflect the new area in time and space that the resulting aggregate holds information about. That is, if two aggregates A' and A'' are combined to the new aggregate A, then $a_1 = \min(a'_1, a''_1)$, $b_1 = \max(b'_1, b''_1)$, and so forth.

In the rest of this article, we will consider the case of two index dimensions. One defining the interval of a road, the other defining the time interval that an aggregate contains information about. As an exemplary value we will use the average speed driven by vehicles in that area. However, the mechanisms presented can be applied to many other aggregation applications.

Security Mechanisms

On top of the generic aggregation mechanism, we design a security mechanism adhering to the constraints discussed earlier. The basic system model is shown in Fig. 2. We assume that information originating from both local sensors as well as other vehicles enters a vehicle's local information base where it is stored, possibly further processed, and finally disseminated to other vehicles in the vicinity. Information originating from local sensors is considered to be trustworthy whereas remote information needs to be further assessed to assign a certain trust in its validity. For this assessment, we employ a combination of selective proofs using cryptographically signed atomic reports and probabilistic verification of the information contained in a given aggregate. This methodology results in a data-centric trust [8] in contrast to a node-centric trust as it would result from a node reputation system.

As soon as a trust value has been assigned to an aggregate, it is further processed locally using only this trust value as basis for decisions by the underlying aggregation mechanism. We call this scope the local *trust domain* of a vehicle. However, as we do not employ a node reputation system, the trust domain does not extend beyond vehicle borders. Therefore, as soon as aggregates are selected for dissemination to nearby vehicles, the assigned trust value looses its significance and will not be communicated. Receiving vehicles will again judge the aggregates according to the security mechanisms.

Selective Attestation

A common technique to secure vehicular communication, especially beaconing applications, is to cryptographically sign each outgoing message, thereby giving a proof of the sending vehicle's identity, given that the accompanying public key is signed by a central, mutually trusted authority. Such a PKI system is commonly agreed to authenticate messages in vehicular networks [9]. If an adversary then sends out beacons with a high frequency to give his information a higher weight, other vehicles can easily detect the high frequency due to the attached signature and discard attacker messages. However, when using aggregation, atomic reports of several vehicles will be combined to aggregates. Even if all atomic reports were cryptographically signed initially, the signature cannot be verified after aggregation because the information that was signed has been altered and one cannot reproduce the single speed reports given only the resulting speed average. Thus, deterministic verifiability is lost.

To achieve a certain verifiability of the aggregates nonetheless, we need to attach a certain amount of meta-information to aggregates, which serves as a witness for the correctness of the aggregation. This meta-information will be called *attestation meta-data*. A simple approach to select such witnesses would be to add all cryptographically signed atomic reports that served as input to the aggregation as attestation meta-data. In that case, any receiver could deterministically verify the aggregation by first checking all signatures of the atomic reports and then re-calculating the speed average and verifying that the result is the same as the one contained in the

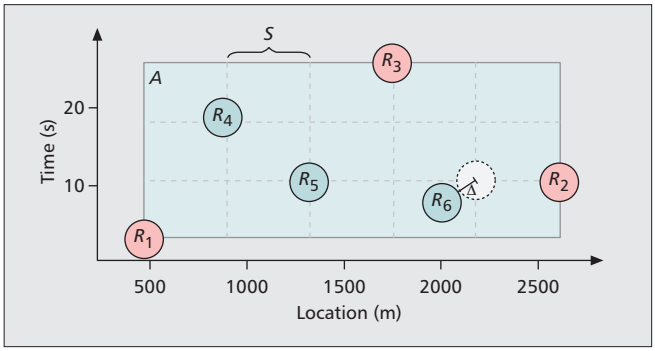

Figure 3. *An example aggregate* A *with attestation meta-data attached. The attestation meta-data is comprised of three border atomic reports* $\{R_1, \ldots, R_3\}$ *and the additional reports* $\{R_4, \ldots, R_6\}$ *selected according to the granularity defined by the security parameter* S.

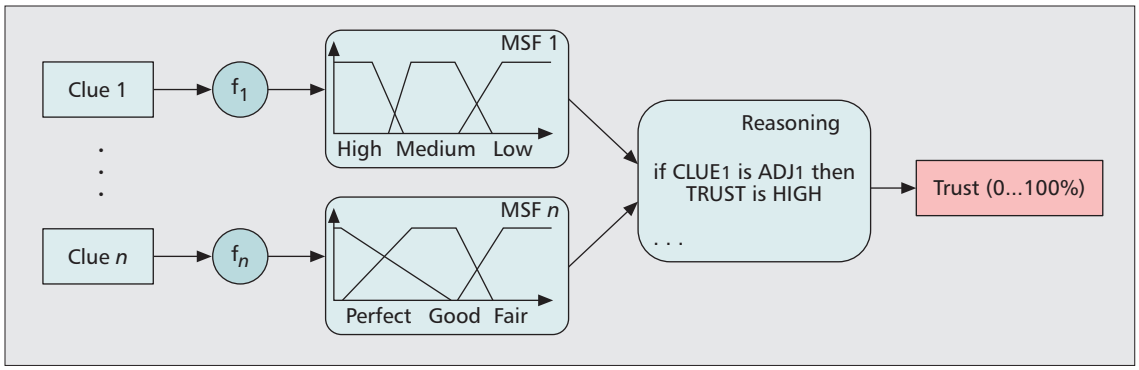

Figure 4. *Methodology of the fuzzy reasoning employed for trust fusion. The membership functions (MSFs) determine how input values map to adjectives. The same value can be mapped to multiple adjectives with different degrees.*

aggregate values. However, if all atomic reports are attached to aggregates as attestation meta-data, there will be no reduction of bandwidth usage.

Therefore, we select a *strategically chosen subset* of all atomic reports that served as input to an aggregate to allow for a probabilistic verification. To stipulate the detection of malicious information with as few meta-data as possible, we employ a list of criteria to select the meta-data during the aggregation process.

First, we can uniquely identify those atomic reports that led to an aggregate's current maximum and minimum in time and space (e.g., $\{R_1, \ldots R_3\}$ for A in Fig. 3). Considering that aggregates will commonly represent an area whose vehicles share common values, e.g., common speed, all such atomic reports defining the borders of an aggregate's area will be added as attestation meta-data. Only if a vehicle is able to present valid signed information about all borders of an aggregate, it can be believed to be valid.

Especially if aggregates cover larger areas (as it is the case for A in Fig. 3), adding values from the borders will only lead to a first indication that an aggregate is valid. An adversary can still select arbitrary atomic reports and craft an aggregate where the claimed values are only present at the borders and not throughout the area. This could, for example, lead to the concealment of a traffic jam. Therefore, additional meta-data is needed. The goal is to select each additional signed atomic report such that they are evenly distributed throughout the aggregate ($\{R_4, \ldots, R_6\}$ in Fig. 3). The finer the granularity of the distribution, the higher the achieved security. To express this granularity, we introduce a security parameter S that defines the required granularity of the additional reports, marked by the dashed lines in Fig. 3. However, reports are not available at arbitrary positions. They can differ from the ideal positions defined by S, e.g., R_6 differs by an amount of Δ_6. Thus, checking the distribution of additional reports is not a binary process. Instead, the conformance to the ideal distribution can be characterized by considering all deviations Δ_i.

To select additional reports during the aggregation process, we proceed inductively. Assume that two aggregates, A_1 and A_2, which already contain additional signed reports according to S, are selected to be combined. The area covered by A_1 and A_2 in time and space can either already overlap, or they can be disjunct. If they overlap, their contained additional reports suffice to achieve a good distribution throughout the aggregate. In the overlap region, only those additional reports are kept that are nearest to the optimal grid. If A_1 and A_2 do not overlap, their distance in time or space can be either smaller than S or larger. If it is smaller, all additional reports of A1 and A_2 together will still approximate the ideal grid well. However, if the distance is larger than S, then there are regions for which no supporting reports can be found. The same considerations apply for the aggregation bootstrapping. When two signed atomic reports are first selected for aggregation and their distance is smaller than S, then the resulting aggregate will adhere to the criteria defined above. Otherwise, it will not. In all cases, we can find the necessary additional reports if aggregates or atomic reports selected for combination are not further apart than S.

Thus, S needs to be selected such that it matches the underlying aggregation mechanism. If, for example, two reports will be aggregated when their distance is less than 1 km, then S needs to be at least 1 km. Smaller values for S will result in honest nodes not being able to adhere to the ideal grid for additional report selection. Larger values of S are possible, resulting in a lower probability of detecting attacks but also in a lower bandwidth overhead. Considering these criteria, an adversary can craft malicious aggregates in two ways. First, he can omit additional attestation meta-data that would otherwise expose the attack. This would lead to an uneven distribution of attestation meta-data which can be detected. Second, an attacker can adhere to the distribution rules. Then, the values contained in the signed reports would make the forgery attempt obvious.

In summary, we have defined two strategies for the selection of signed atomic reports serving as witnesses attesting the correctness of aggregates:

- **Atomic reports from the aggregate borders** are always attached due to their exposed position. As the goal of aggregation is to combine reports from regions with similar characteristics, those values serve as clues for trust in aggregates.
- Further, additional signed atomic reports that are **evenly distributed throughout the area of the aggregate** are selected to underline the correctness of the values. The bandwidth/security trade-off can be adjusted according to application requirements due to a configurable security parameter.

Trust Fusion

As argued earlier, deterministic and cryptographically verifiable proofs of an aggregate's integrity are hard to achieve in the context of vehicular ad hoc networks. However, the presented selective attestation mechanism results in clues leading to trust in the correctness of an aggregate, namely:

- **BORDERS.** All borders of an aggregate are supported by a signed atomic report.
- **DISTRIBUTION.** Additional signed reports are distributed evenly throughout the aggregate, adhering to the security parameter S.
- **VALUE_APPROXIMATION.** All values contained in the signed atomic reports presented as attestation meta-data support the claimed values of the aggregate.

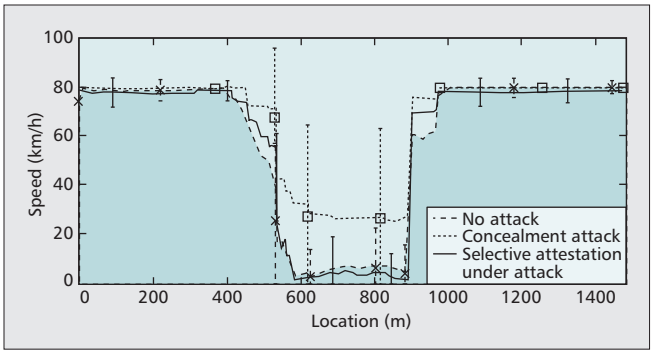

Figure 5. *Averaged traffic flow in three different situations: the actual situation with no attack, a congestion concealment attack, and the same attack with selective attestation security.*

In addition to the clues given by the selective attestation process, there can be a number of further clues that are defined by specific aggregation applications. One example is **VALUE_RANGE**, i.e., the presented values are within their natural range.

Due to the number of possible clues and their correlations with each other, their interpretation can be complex. Therefore, we employ a reasoning mechanism, namely fuzzy reasoning [10], to combine all collected clues to a single value that expresses the trust in the correctness of an aggregate in percent. The basic methodology is shown in Fig. 4. First, each of the clues needs to be expressed as a floating point number. For example, the above mentioned **VALUE_APPROXIMATION** can be expressed by the root mean square error of all atomic reports' values $v_1, \ldots, v_n$ and the claimed value v of the aggregate:

$$\left(\frac{1}{n}\sum_i (v_i - v)^2\right)^{1/2}.$$

Similarly, the distribution of additional signed reports can be expressed as a floating point value by merging all deviations Δ_i from the ideal grid defined by S. Then, a set of adjectives is assigned to each of the clues. Those adjectives characterize the floating point value of their corresponding clue using natural language. An exemplary set of adjectives for the clue **VALUE_APPROXIMATION** is {perfect, good, fair, poor}. Note, that the adjectives do not need to correspond to crisp intervals of the floating point value of a clue but can gradually fade in and out (Fig. 4). Next, if-then-statements can be used to reason about the correlation of several clues, e.g.:

if **DISTRIBUTION** is good and
VALUE_APPROXIMATION is (perfect or good) then
TRUST is high

It is possible to formulate as many of these rules as necessary for a given application. As shown in the example rule, the **TRUST** property is assigned adjectives in the same way as all the clues. Also, the resulting trust value (in the range of 0–100 percent) is mapped to the adjectives low, medium, and high. Thus, only a mapping between the input values' adjectives and the corresponding trust adjectives needs to be expressed by the rules without the need to explicitly express exact correlations between the underlying floating point values. The resulting trust percentage is then calculated by evaluating the rules, considering, for example, which trust adjective has been assigned to the most. Then, all further components of the underlying aggregation mechanism can use this single trust value to judge the correctness of the aggregate, e.g.:

- If the trust value is very low, the aggregate can be discarded.
- Fusion of aggregates with highly differing trust values can be prevented.
- Aggregates with high trust value can be preferred for further dissemination.

However, as soon as an aggregate is selected for further dissemination, the assigned trust value looses its significance. Any vehicle receiving the aggregate will re-evaluate the trust using the attestation meta-data.

The combination of the presented selective attestation mechanism with the trust fusion allows to combine cryptographic signatures used as trust anchors with probabilistic integrity criteria. Although arbitrarily complex correlations of those criteria can be expressed, due to the fusion to only a single trust value, only minor modifications to the underlying aggregation scheme are necessary to make use of the added security mechanisms.

Evaluation

For evaluation, we simulate an aggregation scheme to detect traffic jams [1], corresponding attacks, and our proposed security mechanisms using the JiST/SWANS simulation framework. The simulation model uses two-ray ground path loss and an IEEE 802.11-based MAC. A total of 30 nodes, among which 10 are attackers, move on a 1,500 m-long highway segment with three lanes. After 900 m, the highway is blocked by an obstacle, resulting in a developing traffic jam. The size of dissemination messages is fixed. Aggregated information and attestation atomic reports are added according to the rules specified earlier. For this example application, we chose the security parameter $S = 333$ m, i.e., 3 attestations per kilometer are required in addition to the atomic reports at the borders of an aggregate, corresponding to the settings of the underlying aggregation mechanism. For the temporal dimension, we do not apply S in the simulation, because outdated reports are already ignored by the aggregation mechanism. To assess the effectiveness of the proposed mechanisms, both the achieved security and the induced bandwidth consumption are evaluated.

Achieved Security

The mean aggregated view of the situation, without an active adversary, is shown in Fig. 5 (*no attack*). The simulated adversaries now try to conceal the congestion by crafting aggregates that pretend normal traffic flow on the whole road. Without any security countermeasures, the attackers can alter the aggregated view of the situation on average by 25 km/h (*concealment attack*). Moreover, under attack, the aggregated

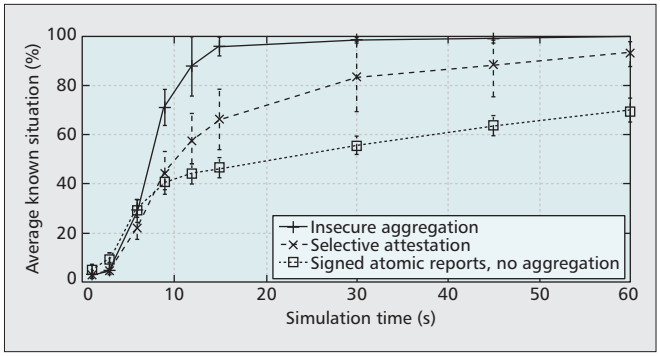

Figure 6. *Comparison of dissemination speed between the underlying aggregation mechanism without security mechanisms, the selective attestation mechanism, and the dissemination of signed atomic reports without aggregation.*

views of the vehicles differ notably, shown by a high standard deviation in the graph. With activated security mechanisms, the resulting mean aggregated view is close to the reference view without attack. An attacker trying to conceal the congestion is not able to present enough supporting information. This results in a lower trust in attacker aggregates leading to a correct representation of the real situation.

Performance

Despite the security gain, one important factor of secure aggregation mechanisms is their scalability. To gauge this performance, we use the dissemination speed as a metric, that is, the amount of time needed until all vehicles have information about a high percentage of the upcoming road. This metric implicitly includes the bandwidth overhead due to security mechanisms, because the dissemination packet size is fixed. Thus, information dissemination is slower if fewer aggregates fit in one packet. The simulation parameters are selected equal to the security evaluation setting. We compare the selective attestation with the performance of the underlying aggregation mechanism without any security considerations and against the periodic dissemination of cryptographically signed atomic reports without any aggregation. Figure 6 shows that the aggregation without security countermeasures outperforms both security aware protocols, as expected. The beaconing of atomic reports without aggregation achieves a certain awareness of the vehicles' neighborhood at first but then only increases linearly as the vehicles explore more of the simulated road by driving along. The selective attestation approach significantly outperforms the dissemination of atomic reports. Although the dissemination speed is slower than the insecure aggregation due to the size of added signatures, at the end of the simulation, over 90 percent of the road is known, whereas the atomic report beaconing only reaches about 70 percent.

It should be noted, that the bandwidth overhead induced by the selective attestation correlates only with the area that is covered by aggregates and is independent of the amount of vehicles in a certain area. Especially in scenarios with very high vehicle density, e.g., traffic jams, the selective attestation therefore clearly outperforms the dissemination of atomic signed reports. Moreover, the computational overhead for the verification of signatures is decreased notably, since only a small amount of atomic reports are transferred, compared to the approach that does not aggregate, i.e., that sends all information as signed atomic reports.

Conclusion

In this article, we have motivated the need for using aggregation in VANETs and highlighted arising security issues that are clearly unique and different compared to those that arise in other VANET scenarios.

We present selective attestation and trust fusion, two mechanisms that add and use a certain amount of additional atomic (i.e., non-aggregated) data attached to a message to evaluate the trustworthiness of the aggregate. In our evaluation, we show that these mechanisms are both efficient and effective. As we explain earlier, compared to earlier work on secure aggregation, our mechanisms fit better to the specific requirements of VANETs that mostly arise from the high network dynamics.

As our mechanisms are based on a generic aggregation model, they can easily be integrated in most existing aggregation schemes and will add significant additional robustness against attacks. In our future work, we plan to implement more complex scenarios and attacker models that will allow us a more detailed analysis of selective attestation and trust fusion.

Another interesting aspect is the relation of aggregation and privacy. As aggregation is removing individual person-related data from the aggregates, it can be per-se considered as a privacy enhancing technology. It is our goal to develop the secure attestation to a point where it can be performed without revealing details about the individuals or vehicles involved to bridge the gap between security and privacy.

References

[1] S. Dietzel *et al.*, "A Fuzzy Logic based Approach for Structure-free Aggregation in Vehicular Ad-Hoc Networks," *Proc. 6th ACM Int'l. Wksp. Vehic. Inter-Networking*, New York, NY, USA, 2009. ACM Press, pp. 79–88.

[2] T. Nadeem *et al.*, "TrafficView: Traffic Data Dissemination using Car-to-Car Communication," *SIGMOBILE Mob. Comp. Commun. Rev.*, vol. 8, no. 3, 2004, pp. 6–19.

[3] Y. Sang *et al.*, "Secure Data Aggregation in Wireless Sensor Networks: A Survey," *Proc. 7th Int'l. Conf. Parallel and Distributed Computing, Applications and Technologies*, Washington, DC, USA, 2006, IEEE Computer Society, pp. 315–20.

[4] F. Picconi *et al.*, "Probabilistic Validation of Aggregated Data in Vehicular Ad-Hoc Networks," *Proc. 3rd Int'l. Wksp. Vehicular Ad Hoc Networks*, New York, NY, USA, 2006, ACM Press, pp. 76–85.

[5] S. Buchegger, J. Mundinger, and J. Le Boudec, "Reputation Systems for Self-Organized Networks: Lessons Learned," *IEEE Technology and Society Mag.*, vol. 27, no. 1, 2008, pp. 41–47.

[6] M. Raya, A. Aziz, and J. Hubaux, "Efficient Secure Aggregation in VANETs," *Proc. 3rd Int'l. Wksp. Vehicular Ad Hoc Networks*, New York, NY, USA, 2006, ACM Press, pp. 67–75.

[7] B. Scheuermann *et al.*, "A Fundamental Scalability Criterion for Data Aggregation in VANETs," *Proc. 15th Annual Int'l. Conf. Mobile Computing and Networking*, New York, NY, USA, 2009, ACM Press, pp. 285–96.

[8] M. Raya *et al.*, "On Data-Centric Trust Establishment in Ephemeral Ad Hoc Networks," *IEEE Infocom 2008*, 2008.

[9] P. Papadimitratos*et al.*, "Secure Vehicular Communication Systems: Design and Architecture," *IEEE Commun. Mag.*, vol. 46, no. 11, Nov. 2008.

[10] L. Zadeh, G. Klir, and B. Yuan, "Fuzzy Sets, Fuzzy Fogic, and Fuzzy Systems," *World Scientific*, 1996.

Biographies

STEFAN DIETZEL (stefan.dietzel@uni-ulm.de) received his Diplom degree in computer science from Ulm University, Germany in 2008. He is now pursuing his Doctorate degree at the Institute of Media Informatics at Ulm University. His research interests include message dissemination mechanisms in general and in-network data aggregation in particular, as well as security and privacy aspects of vehicular communication.

ELMAR SCHOCH (elmar.schoch@uni-ulm.de) received his Doctorate degree in computer science from Ulm University in 2009, working on robust and efficient security mechanisms for communication in inter-vehicle networks. Both for Ulm University and DaimlerChrysler Telematics Research, he was involved in several European IVC projects, such as SEVECOM, PRECIOSA, and NOW. His current research interests center around advanced techniques for vehicular networks, trust, security, and privacy for mobile and ubiquitous systems.

FRANK KARGL (f.kargl@utwente.nl) is an associate professor at the University of Twente in the Distributed and Embedded Security Group. Until 2009 he was a senior researcher at Ulm University, leading a research team focusing on various aspects of VANET communications including information dissemination, applications, and security and privacy. He has co-authored over 80 peer-reviewed publications and is actively involved in research projects like SeVeCom and PRECIOSA. He also contributes to standardization activities, is a regular member of program committees, reviewer for selected journals, panelist and keynote speaker on ITS security and privacy.

BASTIAN KÖNINGS (bastian.koenings@uni-ulm.de) is a Ph.D. student at Ulm University. He is working in a research team with focus on car to car communications. His main research interests are security and privacy in the context of mobile and ubiquitous computing. He is actively involved in the EU-funded research project ATRACO.

MICHAEL WEBER (michael.weber@uni-ulm.de) holds a Ph.D. in computer science from the University of Kaiserslautern. After a number of years in industry, working on parallel and multimedia systems, he joined the University of Ulm as a professor for computer science in 1994 and was appointed director of the Institute of Media Informatics in 2000. He has authored and co-authored more than 150 peer reviewed contributions, edited five books and written a textbook. He has led projects funded by the German Science Foundation (DFG), the state of Baden-Württemberg, by the German Ministry for Education and Research (BMBF), by the European Commission and by industrial partners. His current research interests include mobile and ubiquitous computing systems and human-computer interaction.