Dependable and Secure Geocast in Vehicular Networks

Elmar Schoch, Boto Bako
Ulm University
Institute of Media Informatics
Ulm, Germany
{elmar.schoch, boto.bako}@uni-ulm.de

Stefan Dietzel, Frank Kargl
University of Twente
Distributed and Embedded Security Group
Enschede, The Netherlands
{s.dietzel, f.kargl}@utwente.nl

ABSTRACT

Inter-vehicular communication envisions many applications to enhance traffic safety. One fundamental communication paradigm used to realize a wide range of such applications is called Geocast, that is, multi-hop broadcast dissemination of messages within a geographically restricted destination region. Because of the safety-related nature of many VANET applications, it is crucial that Geocast protocols ensure dependable dissemination of information. Here, dependability has two aspects. First, a Geocast protocol needs to scale to varying node densities – reliable delivery should be provided both in sparsely connected networks and also in scenarios with high channel load due to high node density. In addition, Geocast needs to be protected against attacks to achieve dependable dissemination of information even in presence of malicious nodes trying to suppress information delivery. In this work, we focus on the goal of Geocast security. We evaluate the impact of several attacks, and, based on these results, we introduce specific countermeasures against the discovered threats. Particularly, we highlight the intrinsic security properties already present in scalability-enhancing mechanisms. Thus, we show how security and scalability complement each other in Geocast protocols. In summary, our focus are lightweight and efficient measures to secure Geocast for usage in VANETs.

Categories and Subject Descriptors

C.2.0 [**Computer-Communication Networks**]: General

General Terms

Security, Reliability, Performance

Keywords

Vehicular networks (VANETs), Security, Geocast

VANET'10, September 24, 2010, Chicago, Illinois, USA.
Copyright 2010 ACM 978-1-4503-0145-9/10/09 ...$10.00.

1. INTRODUCTION

Today, advanced driver assistance systems in vehicles use, for example, radar to detect obstacles in a vehicle's way or in the blind spot of the driver. In the future, the capabilities of such active safety systems will be improved notably by wireless ad hoc communication between vehicles, termed VANETs. While single-hop communication is sufficient for numerous applications, multi-hop routing and information dissemination is important as well, for example, to inform drivers at a larger distance about an accident.

One of the main requirements for applications like accident warning is that approaching vehicles must be warned quickly and in a larger area so that drivers can react carefully and keep the situation under control. For that class of applications, Geocast is generally considered a suitable dissemination mechanism. Through so-called Geocast or Geobroadcast, vehicles disseminate a message via multi-hop relaying to all receivers within a designated geographic region. The simplest way to implement Geocast is the use of geographically restricted flooding. In this scheme, every node located in the addressed geographic region performs a link layer broadcast exactly once to disseminate a message over the whole region.[1] A more detailed overview on communication patterns occurring in VANETs and a description of Geocast can be found in [6].

To achieve the full potential of car-to-car communication, one core requirement is the dependability of the communication system. In the accident warning example, all vehicles approaching the accident site should be informed reliably about the hazard. This is particularly important from a security perspective: malicious intruders should not be able to disrupt communication or inject falsified or modified messages, for instance. Hence, providing secure communication protocols is crucial, because otherwise no dependable service can be maintained. For Geocast, considering security is particularly relevant for several reasons. One aspect is the potentially large impact region of Geocast. Due to multi-hop broadcast, messages are received by many nodes. Hence, malicious information can reach a broader distribution using Geocast. Moreover, the flooding principle of Geocast consumes considerable channel resources and, therefore, may be misused for very effective denial-of-service attacks.

In this work, we focus on the security of Geocast. Section 2 reviews earlier work in this field. In Section 3, we identify and analyze potential security issues and quantitatively evaluate the impact of serious problems. Based on this

[1] Note that we do not consider the potential transport from the sender towards a remote destination region.

security analysis, Section 4 proposes security mechanisms to thwart major issues. The evaluation of their effectiveness follows in Section 5. Section 6 concludes the paper with a summary and an outlook.

2. RELATED WORK

2.1 Efficient Geocast

Although the implementation of flooding as the basis for Geocast message dissemination is very simple, it has a huge drawback: with increasing node density in the network, channel congestion becomes a major issue. In extreme cases, this can lead to so-called broadcast storms [12]. From an attacker's perspective, the redundancy can be exploited for amplification denial-of-service attacks, where few packets get replicated and saturate the channel in a large area.

For both scalability and security reasons, more efficient and scalable broadcast mechanisms are desirable. Such efficient broadcast mechanisms are based on the idea that not all nodes need to forward a message, but a smaller subset of forwarding nodes is sufficient to still reach all nodes in the addressed region. More efficient dissemination protocols, therefore, try to minimize this subset of forwarding nodes, relieving the broadcast storm problem. To achieve such a reduction of forwarders, there are two basic approaches: deterministic and probabilistic schemes.

Deterministic approaches make use of topology information to determine a small set of forwarding nodes. Because the determination of an optimal subset (i.e., with minimal size) is NP-hard, deterministic broadcast approaches use heuristics, which are often based on two-hop neighborhood knowledge to determine a not necessarily optimal but still small and sufficient set of relays. Examples of such broadcast mechanisms are Multipoint Relaying [14] and Dominant Pruning [9].

Although deterministic approaches achieve a very high forwarding efficiency, they usually require significant signaling and overhead to determine the optimal forwarding set, which is not desirable in VANETs.

The main difference between probabilistic and deterministic protocols is that probabilistic approaches do not try to approximate this optimal forwarding set in advance to sending a message. Instead, receiver nodes decide on their own, whether to relay a message. In reference to the probabilistic spreading of rumors, probabilistic broadcasting schemes are often referred to as gossiping schemes.

A naïve implementation of such a probabilistic broadcasting scheme is the use of a static forwarding probability at each node (e.g., [15, 16, 11]). Obviously, determining an optimal, static forwarding probability is impossible in dynamic environments like VANETs. An optimal forwarding probability highly depends on the local network density and topology. If the network characteristics are not static, homogeneous, and known in advance, any fixed value results either in a low delivery ratio or a high number of redundant messages.

Although many simple heuristics were suggested to improve the performance of static gossiping (e.g., [16, 12]), they still cannot determine an efficient forwarding probability as shown in [8]. Therefore, more advanced gossiping protocols were proposed, which try to adapt the forwarding probability independently at each node based on localized topology information. In [1], the authors introduced the so called Advanced Adaptive Gossiping protocol (AAG) and showed that this probabilistic dissemination scheme achieves a reliable and efficient information dissemination in a wide range of different VANET scenarios.

The key advantage of such adaptive protocols is that they achieve high efficiency by locally determining an optimal forwarding probability, but require less overhead compared to deterministic schemes due to the probabilistic nature of the protocol. For the same reason, they are more robust to singular node failures.

In the following, we specifically address the role of efficient Geocast protocols from the security perspective.

2.2 Secure Geocast

The security of Geocast in vehicular networks has not been a frequent subject in previous research. One of the major aspects considered so far is message integrity. In [3], the authors propose that both the original sender and intermediate forwarders sign Geocast messages. The latter is intended to secure mutable fields in the message, which are modified by each forwarder. The authors also consider message injection as a threat, but only propose static frequency thresholds for different source node types as limitation of the problem. However, such static limits are presumably well-known to attackers, and can, therefore, be exploited. Other work, like the DRG scheme by Joshi et al. [7], also addresses the robustness of Geocast but not from the perspective of malicious activities. However, their proposed mechanisms can be transferred to the security domain, because they can also detect intentional message losses created by attackers.

3. SECURITY ANALYSIS

Before providing countermeasures, it is important to know which security goals are important and which attacks are feasible. Therefore, we first revisit traditional security goals to decide on their relevance for Geocast in VANETs. In a second step, we investigate two major attacks on Geocast. In particular, we give quantitative figures that show the severity of these attacks.

Regarding security goals, the following are well established in classic network security [10]:

1. *Confidentiality* denotes the concealment of information.

2. *Integrity* refers to the trustworthiness of data, i.e., the prevention of unauthorized changes.

3. *Availability* denotes the ability to use a system at all times.

Confidentiality.

Keeping transferred information private is the number one goal in today's wireless networks. From the perspective of VANET-based safety applications, however, the opposite is the case: information disseminated by Geocast is intentionally public and should reach all vehicles in the destination region. Because of constant topology changes and broadcast communication, potential receivers are not known to the sender, thus rendering data encryption impractical. Therefore, we conclude that confidentiality is not a relevant goal for Geocast in VANETs.

Integrity.

Several attacks on Geocast become feasible because messages can be easily manipulated by intermediate nodes during the multi-hop dissemination. Moreover, due to wireless communication, any device with suitable radio equipment can interfere with the communication in general, not only with Geocast. Thus, even more than in traditional networks, integrity is very important for VANETs. Without preventing unauthorized message changes or message spoofing, an attacker can easily manipulate the system, in the worst case resulting in accidents provoked by false warnings. To ensure the integrity of VANET communication, a number of solutions were proposed.

Authentication and integrity verification of messages is commonly assumed to be a basic protection primitive in VANETs [13]. This can be achieved by using digital signatures: the sender computes a hash of the message, signs it using his signing key, and all receivers will be able to verify the signature using the corresponding verification key. The verification key is encapsulated in a certificate that is attached to messages. Certificates are assumed to be issued by a trusted third party that verifies the validity of vehicles. In summary, this scheme ensures the integrity of messages as well as the authenticity of senders, so that only insider attackers are able to forge new messages.

Availability.

The third security goal, availability, gets more and more important in traditional networks. Similarly, availability is also of key importance for the deployment of VANET applications. Without availability protection, attackers can disrupt and congest message dissemination with low effort. Even considering integrity protected messages, denial-of-service attacks are still feasible. Yet, availability is only rarely addressed by previous research so far.

Summarizing the security goals, we find that confidentiality is not relevant for Geocast in VANETs, integrity can be addressed by asymmetric cryptographic schemes, but availability is under-investigated. Therefore, we specifically focus on availability of Geocast in VANETs in this work. The goals are to give figures on the danger of certain attacks and to propose lightweight and efficient countermeasures to achieve dependable and secure Geocast.

3.1 Attacker model

Before we are able to precisely analyze the consequences of attacks on availability, we have to define the capabilities of potential attackers. We assume the following attacker characteristics in our attacker model:

- *Single attacker* with only a local view on the network. Multiple concurrent but non-colluding attackers may be present in the network.

- *Insider and outsider attacker.* As insider attacker, we consider nodes which are able to authenticate themselves as proper participants of the network using valid cryptographic keys. Outsider attacks work without the ability of authentication.

- *Active attacker.* An attacker may listen to the communication and may also interfere actively with communication protocols.

- *Short-term attacker.* We consider attacks lasting only over a relatively short period of time. Long lasting attacks (days or months) are assumed to be prosecuted by a supervising entity that is able to detect and remove attackers from the network.

- *Attacker knowledge.* We assume that the attacker is aware of all communication protocols and algorithms and is able to use and manipulate them. Moreover, attackers may completely control captured nodes.

- *Stationary attacker.* Attackers do not move around in our analysis, the stay at a fixed position. Moving attackers could affect single nodes more seriously, but from a network perspective, the effect of mobile attackers is expected to be similar.

3.2 Impact of attacks on availability

Considering this attacker model, several attacks on availability are possible. For example, authentic messages may be replayed by an outsider after a certain time or at another location. Similarly, attackers may violate forwarding rules or interfere with lower layers. Moreover, insider attackers may forge messages, which carries a large potential to cause damage. All these attacks are possible even in presence of common integrity protection primitives. In the following we present two representative, generic attack approaches: saturating the medium by overloading it with messages and selective jamming of the channel to hinder message propagation. For both attacks, we focus on quantifying their severity, an aspect that is missing in many attack and risk analyses.

3.2.1 Denial-of-service by overloading

The intention of Geocast is to disseminate information in a potentially large destination region. In a basic implementation, messages are flooded within the region, that is, every node inside the region forwards each message once. This simple mechanism opens opportunities for denial-of-service attacks, because the dissemination consumes considerable communication channel resources. If an attacker forges messages with large destination regions and at high frequency, the channel gets congested.

In order to get a more precise picture of the potential damage, we simulated Geocast dissemination in city scenarios with fixed field size of 4×4 km, but a varying number of nodes in order to test the effect with different node densities. In addition to Geocast messages sent by normal nodes, a single adversary creates new Geocast messages with a certain frequency and addresses them to the whole simulated area. As we use the same simulation to test our proposed countermeasures, more details on the simulation environment can be found in Section 5.

Figure 1 depicts the success of information delivery under normal conditions and under attack. A total of 500 messages are sent by various nodes during simulated 60 s. We find that delivery success without attack improves with increasing vehicle density, as network partitioning becomes less influencing with higher density. Delivery success converges to 100 % when vehicle density reaches a certain level (e.g., approximately 97% with 1000 vehicles in the city).

These conditions change notably when the attacker is active and sends forged messages. The delivery success metric in this case still considers regular Geocast messages only.

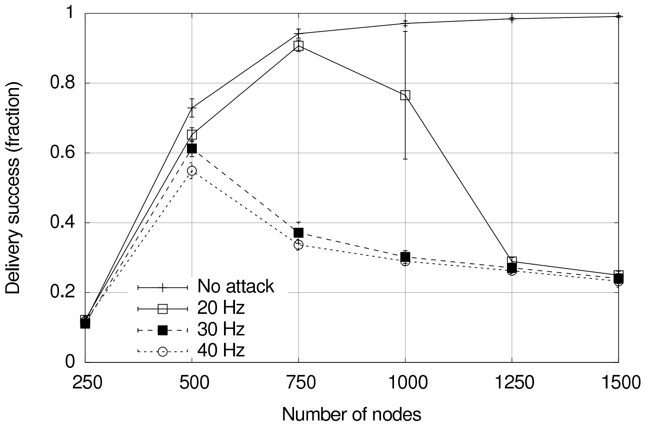

Figure 1: Effect of massive message injection on flooding-based Geocast.

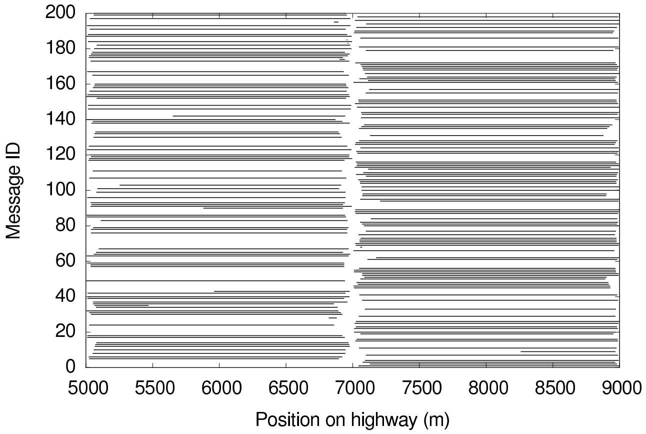

Figure 2: Effect of selective jamming on Geocast message dissemination along a highway.

We observe two effects of the attack. In scenarios with lower node density, the attack is automatically inhibited because of frequent network partitions, and, thus, can influence delivery success only marginally. On the other hand, we can see that the redundancy of flooding-based Geocast makes the dissemination collapse at a certain vehicle density. The large variance with 20 Hz message frequency and 1000 nodes on the field clearly delimits this point. When the attacker increases the injection frequency to 30 Hz and 40 Hz, the dissemination of regular Geocasts starts to decline even with low node density and converges to a saturation level of about 20 % with higher density. Hence, we can conclude that frequent injection of Geocast messages can easily congest the network and massively harm delivery of regular messages in a very large area.

3.2.2 Denial-of-service by selective jamming

Another class of attacks uses radio interference to harm successful delivery of Geocast. In this case, attackers send a jam signal immediately after detecting a transmission, that is, after the wireless medium is sensed to be busy. This way, messages in transmission cannot be received by other nodes in the vicinity of the attacker due to adverse noise conditions. The physical destruction of transfers particularly affects broadcast packets, because these messages are not acknowledged and retransmitted by the 802.11-based MAC protocol in case of failures. Since Geocast uses MAC-layer broadcasts for single hops, the overall dissemination can be seriously affected.

Figure 2 shows the result of a selective jamming attack simulation on the dissemination along a highway. Geocast messages are created by nodes on the highway and address the whole road segment between 5 and 9 km. Each line symbolizes the geographical dissemination of a single message, that is, when a line starts at 5000 m and ends at 6000 m, the corresponding message has been received by nodes between these two locations on the highway. The attacker uses a manipulated MAC layer protocol, which starts to send as soon as the MAC enters the busy state. The break in the center, where the attacker resides, is immediately visible. A closer look reveals that none of the Geocast messages has been delivered beyond the attacker's location, neither from left to right nor vice versa. This shows that an attacker can successfully interrupt dissemination on long stretches of road segments such as highways.

4. SECURE AND EFFICIENT GEOCAST

As we have seen in the previous Section, there are different possibilities for attackers to achieve their goals. The analysis showed that, even in the presence of basic security measures such as authentication and integrity protection, two major attacks on availability are still possible and may have serious impact on the availability of the Geocast dissemination. These are denial-of-service by overloading and denial-of-service by selective jamming. In this section, we now discuss three mechanisms that directly enhance the security of a Geocast protocol with respect to these two attacks.

4.1 Efficient Geocast

An overloading attack has a disastrous effect on the delivery success rate. However, we acknowledge that the redundancy of flooding facilitates the attack. Therefore, our first claim to thwart overloading is to use more efficient multi-hop broadcast protocols instead of flooding. The assumption is that using an efficient and scalable Geocast protocol, i.e., a protocol that minimizes message redundancy, mitigates the extent of such attacks. Beyond that, efficiency is also of eminent importance to overcome the broadcast storm problem. As robustness of communication protocols is a fundamental requirement for safety applications in VANETs, using efficient Geocast serves both scalability and security.

As the discussion of the related work reveals, probabilistic broadcast protocols offer a good tradeoff between these properties. We chose the advanced adaptive gossiping protocol (AAG) from [1] as a basis to evaluate the gain of security of efficient Geocast protocols. In the following, we briefly describe the AAG protocol.

The basic idea of the AAG protocol is to compute an accurate forwarding probability at each node independently of other nodes, based on two-hop neighborhood information. Having these two-hop neighbor information, on reception of a message from sender S, a potential forwarder node F can determine all its neighbors that have not yet received the message from S, which we call *child nodes* of F. These nodes are direct neighbors of F, but are not in communication range of the sender S. Thus, node F is responsible for forwarding the message to these nodes. However, due to redundant dissemination paths, there might be more nodes that potentially forward the same message to F's child nodes. Thus, F checks for each child node C whether other nodes can forward the message to it. These nodes are called *siblings* of F or *parents* of C and are the common neighbors of nodes S and C.

Now, knowing the number of parents, an efficient forwarding probability can be determined to meet a required application-determined reception rate while, at the same time, minimizing the forwarding rate. In the following, a more formal definition of the protocol according to [1] is given.

Let N be the set of all nodes in the network and $F \in N$ be a forwarding node which receives a message M from a sender $S \in N$. S is a one-hop neighbor of F ($S \in \mathit{neighbor}(F)$). The set of nodes that received the message M is called M_r and is equal to the neighbors of S. F determines its neighbors that have not received the message using two-hop neighborhood information:

$$\mathit{child}(F) = \{C \in N \mid C \in \mathit{neighbor}(F) \wedge C \notin M_r\}$$

Having the child set of node F, for each child $C_i \in \mathit{child}(F)$ all nodes are determined which are possible forwarders of message M:

$$\mathit{parent}(C_i) = \{P \in N \mid P \in \mathit{neighbor}(C_i) \wedge P \in M_r\}$$

Knowing the number of parents ($K = \#\mathit{parent}(C_i)$), the forwarding probability can be determined with the following equation:

$$(1 - p_{forward})^K < (1 - \tau_{rel})$$

where τ_{rel} is a so called per hop reception probability which also considers the network diameter δ to ensure the required application reception rate (τ_{arp}) is met even in large networks. τ_{rel} can be determined by the following equation:

$$(\tau_{rel})^\delta = \tau_{arp}$$

4.2 Adaptive load control

Efficient message dissemination protocols reduce the redundancy of communication, and, thus, reduce the vulnerability of Geocast against message injection attacks. However, increasing the frequency of bogus messages suffices to cause congestion again. Therefore, a dedicated security mechanism is required to make Geocast robust against such DoS attacks.

4.2.1 Flexible load metric

Previous work has proposed static rate limitations in order to thwart message injection. Yet, we argue that more parameters have to be taken into account that determine the damage caused by a massive message injection. For example, the size of the Geocast dissemination region clearly is a factor. The available network bandwidth could cope with comparatively high frequencies of new messages if the destination regions are small. Vice versa, if a large destination region is required by an application, the sending frequency should be low. Therefore, we propose a load metric that determines the load caused by one node using multiple influence factors. We define the load l_i created by node i over a period p as

$$l_i(p) = \prod_k w_k(x_k(p))$$

where k denotes several influence factors, w_k describes an individual weighting function of each factor, and $x_k(p)$ is the measured value of factor k within p. Here, we discuss three influence factors k:

- k_1: frequency of messages produced by a node i during p,
- k_2: accumulated destination region sizes, and
- k_3: accumulated payload sizes.

The weighting functions are useful to assign a non-linear influence to certain factors, for example, twice the payload size does not directly relate to twice the effective load.

As a result, we get a more precise estimation of the load created by a node and can determine a threshold for acceptable load. Involving multiple factors also offers the advantage of improved flexibility for applications. For example, in case of an accident warning, a large destination region can be used, because the frequency of messages is low. In contrast, other applications may broadcast with higher frequency, but will have to apply a smaller destination region.

4.2.2 Traffic limitation

In order to be effective, every node constantly monitors incoming Geocast messages from various nodes and updates the load caused per source node. Based on this data, a node can decide whether to forward individual messages. This decision could be made using a static load threshold, that is, when a source node exceeds the acceptable load threshold, its messages are no longer forwarded. However, we propose to adapt this threshold dynamically to the current measured overall load. Since we have defined a load metric, we can locally monitor the amount of data traffic currently being disseminated by all nodes. Hence, the threshold of the acceptable load θ created by one node will be low when the general load level is currently high. Adapting the acceptable threshold has two key advantages: first, an attacker cannot predetermine the amount of load that is safe to create without being detected. Second, an attacker cannot easily overload the channel when it is already saturated from regular traffic. Moreover, if multiple attackers are present, their created load adds up and lowers the threshold accordingly. This also implicitly addresses the Sybil attack: even if an attacker manages to act under multiple identities, this is not a rewarding option because the threshold of acceptable load simply adapts itself.

We propose the following adaption model for the dynamic threshold. First, we define θ_{min} and θ_{max} as a minimum and maximum threshold. Defining a minimum is useful to keep the chance of successful dissemination for regular Geocast messages even in high load situations. In particular, high overall load may be caused artificially by an attacker, who then could inhibit forwarding of other messages if θ was reduced without lower boundary. The maximum threshold θ_{max} sets a fixed upper bound so that instantly starting heavy attacks are quickly detected as well. Otherwise the adaption to a reasonable level could take too long. Between θ_{min} and θ_{max}, the current θ is adapted in j steps, or in other words, in j intervals of width $\delta = \frac{\theta_{max} - \theta_{min}}{j}$. The periodic adaption then works as follows:

$$\theta_{new} = \begin{cases} \theta - \delta, & l > \theta + \epsilon, \theta > \theta_{min} \\ \theta + \delta, & l < \theta + \epsilon, \theta < \theta_{max} \\ \theta, & \text{otherwise} \end{cases}$$

The effect of this scheme is that the threshold θ decreases when the load increases and vice versa. When creating load, an attacker implicitly also lowers the thresholds at other nodes. Effectively, the attacker inhibits the effectiveness of the performed attack.

4.3 Message loss avoidance

Besides channel overload by attackers, we have seen that attackers may cut off Geocast dissemination by interrupting broadcast transmissions on the radio layer. Especially topologies like long road segments are affected by the attack. Effectively, the selective jamming attack creates artificial gaps that any Geocast protocol cannot jump across. In order to bridge such gaps, we propose an enhancement that can be applied broadcast dissemination schemes like flooding or AAG. The rationale of this message loss avoidance scheme is to store messages, if necessary, transport them physically with the moving vehicles over a certain distance, and retransmit them afterwards. This way, Geocasts are not lost at fortuitous network partitions or maliciously created gaps.

Our proposed scheme consists of two parts: the detection that a message might not have been received and a strategy for retransmission. For the detection of forwarding by other nodes, we employ implicit acknowledgements, an idea which has been used before. New is that nodes store messages to exploit physical movement of nodes for transporting messages. Therefore, every node stores new messages first, forwards them, and then starts counting the number of subsequent receptions of the same message. We assume that later receptions of a message indicate a successful continuation of the forwarding process. However, this assumption may not hold if only few retransmissions are received. Hence, our scheme uses a parameter C that denotes the number of subsequent receptions, after which a message is assumed to have continued forwarding. In AAG, the counting can even be restricted to the reception from child neighbor nodes. In case that less than C receptions are counted, we assume that a retransmission is required. Retransmissions of such messages are initiated after a defined delay d. After being sent, counting receptions continues and eventually reaches C. If not, replays are repeated no more than k times, and messages with expired lifetime L are purged from the memory.

The set of parameters C, d, k, and L determines both the success of message loss avoidance and the created overhead. If C and k are large, many messages may be retransmitted multiple times, leading to more channel load but also leading to better delivery success. On the other hand, if C and d are small, the successful forwarding may be assumed too early. We propose to let applications decide to which extent they require successful delivery and thereby define suitable parameters.

5. EVALUATION

5.1 Simulation environment

In order to evaluate both the impact of attacks and the effectiveness of our proposed countermeasures, we developed corresponding simulations. As a base, we use the JiST/SWANS network simulator [2] with various extended models to run sufficiently realistic simulations of vehicular networks. We use the STRAW model [4] to simulate vehicle movements in cities and apply the highway scenarios of the FleetNet project. To vary node density, we use different numbers of nodes on a fixed field size. Throughout the simulation, we create a fixed number of Geocast messages, sent by randomly selected nodes distributed over the whole simulated time period of 60 s. The radio and MAC layers operate according to the IEEE 802.11a standard, which is close to the IEEE 802.11p variant foreseen for vehicular communication. We do not expect that the difference has major impact on the fundamental validity of our results. For the physical layer, we use the two-ray ground model and receivers capable of noise cumulation, which is important to see effects of colliding packets. Again, we consider noise accumulation as more relevant than a very accurate physical signal propagation model. Message sizes include required space for signatures and certificates of 200 bytes. For statistical reasons, all simulation configurations were run ten times.

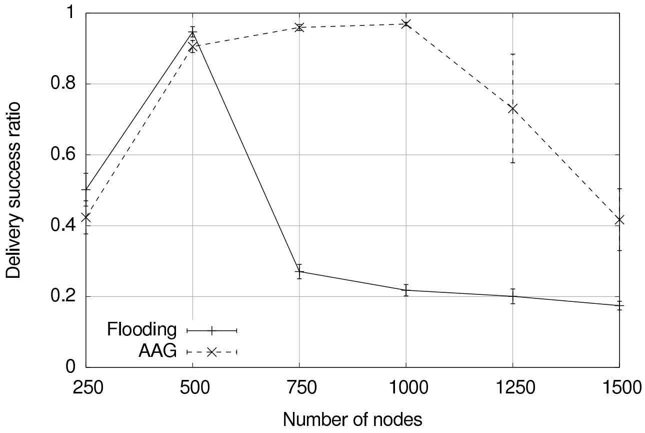

Figure 3: Efficiency helps security: the AAG protocol reduces redundancy of flooding and thus alleviates the severity of massive message injection.

5.2 Effect of efficient Geocast

As shown in Section 3.2.1, the delivery success of simple flooding drops drastically in presence of an overloading attack. Thus, an attacker can easily achieve a denial-of-service attack by massively injecting messages. The high message redundancy in simple flooding facilitates such attacks, whereas efficient Geocast protocols are supposed to mitigate this effect implicitly. Therefore, we evaluate the behavior of an efficient protocol like AAG in the presence of such attacks.

In order to compare the performance of flooding and AAG, we simulated a city scenario with a field size of 3×3 km with varying node densities. We reduced the field size compared to the previous analysis in order to further increase node density, so that an overloading effect on AAG becomes visible. Besides the Geocast messages sent by normal nodes, a single attacker in the center of the scene creates forged messages with a Geocast region that spans the whole scene at a frequency of 20 Hz. As before, Figure 3 shows, simple flooding is heavily affected by this injection attack. Already with 750 nodes, the delivery success ratio drops drastically and converges to 20% with higher densities. On the other hand, AAG mitigates the extent of the attack significantly. It maintains a high delivery ratio up to 1000 nodes. Never-

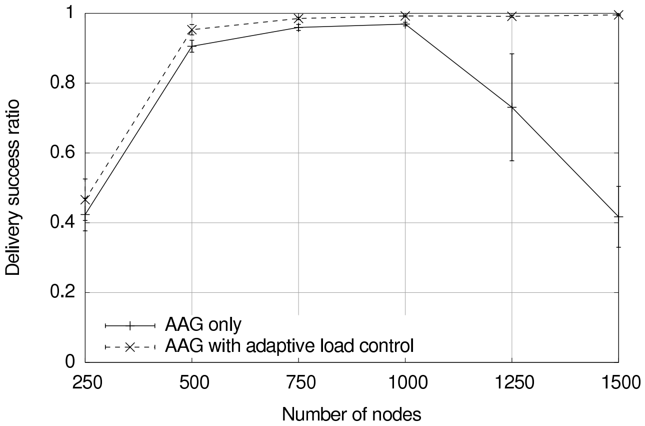

Figure 4: With activated load control, AAG manages to deliver the utmost amount of messages in all considered node densities despite of an ongoing attack.

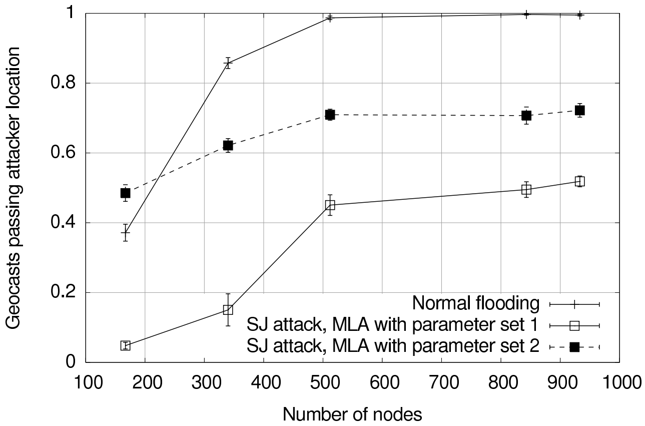

Figure 5: Geocast messages passing a selective jamming attacker when the message loss avoidance is active. Parameter set 2 uses more strict parameters than parameter set 1.

theless, the performance of AAG also declines notably with even higher node density.

Thus, in the case of overloading attacks, the efficiency of a Geocast protocol clearly helps security. However, even considering an efficient protocol, an attacker can still successfully achieve denial-of-service. But the effort to achieve results equivalent to those possible with a Geocast based on flooding is much higher. That is, an efficient Geocast protocol necessitates a much higher injection frequency or destination region in order to achieve a successful attack. This implicates yet another advantage of efficient protocols: attackers can be detected much easier because of the abnormal behavior needed to achieve their goals.

5.3 Effect of load control

We will now evaluate the performance of the adaptive load control mechanism described in Section 4.2. For the evaluation, we use the same simulation parameters as in the last scenario and compare the performance of the standard AAG protocol and AAG with the adaptive load control extension.

Figure 4 shows the simulation results. Like before, we see that the standard AAG protocol is also affected by the overloading attack when a certain node density is present. Thus, the delivery success rate drops in dense scenarios. When we add the adaptive load control scheme to AAG, however, we find the protocol operating optimally, delivering almost 100% of all messages.

These results prove the effectiveness of the proposed overload control. It manages to control the load adaptively and thus to detect and isolate misbehaving nodes that produce high amounts of traffic. This way, a normal operation mode of the communication protocol is enabled even in the presence of such attacks. Moreover, the scheme is lightweight and operates locally, so that not much overhead except the packet monitoring is necessary. Moreover, by taking into account all relevant factors into the load metric and through its threshold adaption mechanism, it can cope with most kinds of overload attacks, even if attackers conduct a Sybil attack using multiple identities.

5.4 Effect of message loss avoidance

The simulation results depicted in Figure 5 show the effect of our message loss avoidance mechanism. In this case, we consider 200 Geocast messages being disseminated along a highway stretch of four kilometers length, with a selective jamming attacker at the center. Recall that, without protection, none of the Geocasts passed the attacker from one side to the other. In other words, Geocast dissemination was entirely disrupted by the attacker. With the message loss avoidance, the situation improves notably. In Figure 5, we compare two different parameter sets of the message loss avoidance with normal dissemination without attack. We find that up to 50% of the Geocast messages pass the attacker in case of parameter set 1 and up to about 70% with parameter set 2.

The two sets differ in important parameters for the message loss avoidance. Set 1 is not very strict – only one implicit acknowledge is required to stop replay, and if this is not received, one replay after three seconds is done. However, we can see that this relaxed setting already copes well with the attack. If an application needs better delivery to the expense of higher overhead, it may work with set 2. In this case, we require two acknowledgements, allow two replays, and delay the replays ten seconds. This scheme copes even better with the attack, and even outperforms normal dissemination in case of low node density, because regular network partitions can be bridged. The drawback of the message loss avoidance is clearly that it increases the overhead because of message replays. However, our simulations show that the additional load is still considerably lower compared to flooding in dense networks. In sparse networks the rebroadcasts are more frequent due to network partitions, resulting in a similar overhead as flooding. But in such environments the overhead is not significant because the absolute number of messages is considerably low.

6. CONCLUSION

In this work, we investigated Geocast for vehicular networks from the security perspective. We showed that, despite usage of basic security primitives (like digital signatures), there still exist significant attack opportunities to disrupt message dissemination: denial-of-service by overloading the communication system and denial-of-service by selective jamming. As our simulation-based evaluation demonstrated, both attacks can significantly affect the correct operation of a Geocast protocol.

As a first step towards improving the robustness of Geocast against overloading attacks, we showed that efficient broadcast protocols like AAG are suitable to stand such DoS attacks for a while. But even the most efficient protocol cannot withstand massive message forging. For additional protection, we therefore introduced adaptive load control. The adaptive load control dynamically monitors both the load created by single nodes and the overall load in a network. No fixed threshold value is used, but controlling parameters are constantly adapted to the current situation. As our simulations reveal, this mechanism efficiently overcomes overloading attacks. Moreover, the scheme is lightweight, does not involve any computationally costly cryptographic operations or cause additional transmission overhead.

To counter the selective jamming attack, we introduced the message loss avoidance mechanism. The idea is to transport messages physically over an artificial or natural network gap in a vehicle; thus, stationary attackers can be easily jumped over and more dependable message dissemination is enabled. This protects a Geocast protocol in the case of a selective jamming attack, but also in the case of partitioned networks.

As an overall result, we get very efficient security measures in order to achieve more dependable Geocast. Moreover, we want to conclude that the two aspects of efficiency and security should not be seen in isolation but have a deep interrelation: *efficiency enhances security* and *security enhances efficiency.*

7. REFERENCES

[1] B. Bako, F. Kargl, E. Schoch, and M. Weber. Advanced Adaptive Gossiping Using 2-Hop Neighborhood Information. In *Proc. IEEE Global Telecommunications Conference IEEE GLOBECOM 2008*, pages 1–6, Nov. 30 2008–Dec. 4 2008.

[2] R. Barr, Z. Haas, and R. van Renesse. JiST: An efficient approach to simulation using virtual machines. *Software Practice & Experience*, 35(6):539–576, 2005.

[3] Charles Harsch, Andreas Festag, and Panos Papadimitratos. Secure Position-Based Routing for VANETs. In *IEEE 66th Vehicular Technology Conference (VTC2007-Fall)*, Oct. 2007.

[4] David R. Choffnes and FabiÂůn E. Bustamante. An Integrated Mobility and Traffic Model for Vehicular Wireless Networks. In *Proc. of the 2nd ACM International Workshop on Vehicular Ad Hoc Networks (VANET)*, Sept. 2005.

[5] Elmar Schoch, Frank Kargl, Tim Leinmüller, and Michael Weber. Vulnerabilities of Geocast Message Distribution. In *2nd IEEE Workshop on Automotive Networking and Applications (AutoNet 2007, in conj. with GlobeCom 2007)*, Washington, DC, USA, Nov. 2007.

[6] Elmar Schoch, Frank Kargl, Tim Leinmüller, and Michael Weber. Communication Patterns in VANETs. *IEEE Communications Magazine*, 46(11):2–8, Nov. 2008.

[7] Harshvardhan P. Joshi, Mihail L. Sichitiu, and Maria Kihl. Distributed Robust Geocast Multicast Routing for Inter-Vehicle Communication. In *Proc. of the First Workshop on WiMAX, Wireless and Mobility*, May 2007.

[8] P. Kyasanur, R. R. Choudhury, and I. Gupta. Smart Gossip: An Adaptive Gossip-based Broadcasting Service for Sensor Networks. In *Proc. IEEE International Conference on Mobile Adhoc and Sensor Systems (MASS)*, pages 91–100, Oct. 2006.

[9] H. Lim and C. Kim. Multicast tree construction and flooding in wireless ad hoc networks. In *MSWIM '00: Proceedings of the 3rd ACM international workshop on Modeling, analysis and simulation of wireless and mobile systems*, pages 61–68, New York, NY, USA, 2000. ACM.

[10] Matt Bishop. *Computer security: art and science.* Addison Wesley, 2004. ISBN: 0-201-44099-7.

[11] Matthew J. Miller, Cigdem Sengul, and Indranil Gupta. Exploring the Energy-Latency Trade-Off for Broadcasts in Energy-Saving Sensor Networks. In *ICDCS '05: Proceedings of the 25th IEEE International Conference on Distributed Computing Systems*, pages 17–26, Washington, DC, USA, 2005. IEEE Computer Society.

[12] S.-Y. Ni, Y.-C. Tseng, Y.-S. Chen, and J.-P. Sheu. The Broadcast Storm Problem in a Mobile Ad Hoc Network. In *MobiCom '99: Proceedings of the 5th annual ACM/IEEE international conference on Mobile computing and networking*, pages 151–162, New York, NY, USA, 1999. ACM.

[13] P. Papadimitratos, L. Buttyán, T. Holczer, E. Schoch, J. Freudiger, M. Raya, Z. Ma, F. Kargl, A. Kung, and J.-P. Hubaux. Secure Vehicular Communication Systems: Design and Architecture. *IEEE Communication Magazine*, 46(11), Nov. 2008.

[14] A. Qayyum, L. Viennot, and A. Laouiti. Multipoint Relaying for Flooding Broadcast Messages in Mobile Wireless Networks. In *Proc. 35th Annual Hawaii International Conference on HICSS System Sciences*, pages 3866–3875, 7–10 Jan 2002.

[15] Ranveer Chandra, Venugopalan Ramasubramanian, and Kenneth Birman. Anonymous Gossip: Improving Multicast Reliability in Mobile Ad-Hoc Networks. Technical report, Cornell University, Ithaca, NY, USA, 2001.

[16] Zygmunt J. Haas, Joseph Y. Halpern, and Li Li. Gossip-based ad hoc routing. *IEEE/ACM Trans. Netw.*, 14(3):479–491, 2006.