Enhancing the Security of Local Danger Warnings in VANETs - A Simulative Analysis of Voting Schemes

Benedikt Ostermaier*
Inst. for Pervasive Computing
ETH Zurich
8092 Zurich, Switzerland
ostermaier@inf.ethz.ch

Florian Dötzer*
ASKON ConsultingGroup
60313 Frankfurt, Germany
florian.doetzer@askon.de

Markus Strassberger
BMW Group
Forschung und Technik
80992 Munich, Germany
markus.strassberger@bmw.de

Abstract

The upcoming deployment of vehicular ad-hoc networks does not only facilitate novel telematics applications, but also poses strong requirements on security. Especially the adoption of active safety applications may raise new threats to road safety if security issues are not properly handled, thus thwarting their initial purpose. In this paper, a special active safety application is considered that enables cooperative foresighted driving through the exchange of local danger warnings, which are based on individual observations and refer to the current road condition. From a security point of view, the decision whether or not such an application should rely on a reported hazard, is a crucial issue, which cannot be completely protected by conventional security measures. We propose an additional security mechanism based on an information centric evaluation of the plausibility of received hazard messages. We developed four decision methods, which are based on voting schemes, and evaluated them by simulation using two attacks trying to manipulate the decision process by distributing false information. Our results indicate that the proposed information centric evaluation of remote observations is a reasonable means to increase the stability and security of a cooperative local danger warning service.

1 Introduction

VANETs - Vehicular Ad hoc NETworks - are a subset of the class of mobile, self-organizing and decentralized networks, called mobile ad hoc networks (MANETs), that consist of cars acting as mobile routers. A couple of research projects have addressed technologies and applications dedicated to VANETs (e.g. [1, 2, 3, 4]). One of the major stimuli for VANETs is the desire to further increase road safety and traffic efficiency by using communication. In this way, vehicles will be able to utilize sensor readings of other vehicles and may thus extend their own sensing capabilities. The exchange of local danger warnings (LDW), which are based on local sensor readings, is considered one of the most promising active safety applications for inter-vehicle communication.

The security of such a system is of utmost concern, since wrong or manipulated information could lead to a decrease of road safety. Forged messages could ultimately provoke accidents, a threat denoted as *intelligent collisions* in [7]. However, in such a highly dynamic system, with a potentially large number of nodes, conventional security measures such as digital certificates, tamper-proof hardware and network security schemes are not sufficient.

Therefore we propose a novel information centric approach - lightweight plausibility checks on application level - that takes advantage of the large number of nodes that can be expected in most traffic scenarios and complements conventional security measures. We developed and analyzed four decision methods and evaluated them by simulation using two attacks trying to manipulate the decision process by distributing false information. We could show that the proposed information centric evaluation of remote observations is a reasonable means to increase the stability and security of a cooperative local danger warning service.

This paper will focus on improving the security of LDW systems by following an information centric approach and evaluate the concepts by simulation. While it will relate to specifics of message distribution, sensor / detection approaches and privacy where applicable, an in-depth discussion of these topics is not within the scope of this paper.

The article is structured as follows: First, the basic paradigms and characteristics of a local danger warning

*Major parts of this work have been carried out at BMW Group Forschung und Technik within the scope of the PReVENT / WILLWARN research project, funded by the European Commission in the context of the 6th framework.

service are described. Section 3 discusses the deficiency of conventional security, specific security requirements, related work and attack scenarios. In section 4 we describe our concepts of plausibility evaluation. Next, section 5 presents the simulation environment and the scenarios used for the evaluation of the proposed approach. Section 6 contains a detailed discussion of the results. Finally, section 7 sums up the results and points out to future work.

2 The Local Danger Warning Application

Predictive driver assistance systems are a key issue in the visionary field of accident-free driving. Vehicles should be enabled to foresee critical driving conditions and therefore inform their drivers timely. In this context, direct communication among vehicles that form a spontaneous ad-hoc network will complement on-board sensor systems, enabling a new paradigm in driving assistance: collaborative and predictive situation-awareness. The exchange of local danger warnings (also known as *regional alerts*), where vehicles directly exchange information about dangerous traffic situations based on local sensor readings, strives to realize this paradigm. In the following, we introduce the basic concepts of such a service, as outlined in a variety of project descriptions and articles (e.g. [19, 26, 5, 4]).

Cars[1] equipped with an LDW Application try to automatically detect hazards whilst driving, using their on-board sensors. Possible hazards for example include low friction, reduced visibility and obstacles on the road. Whenever a potential critical road condition is detected, a new warning message is generated and subsequently disseminated within a certain area. By forwarding received messages, vehicles act as relay nodes, thus enabling the distribution of messages beyond the immediate transmission range of the detecting vehicle. Cars receiving such a message evaluate its content. If there is sufficient evidence for a critical road condition on the route ahead, the system notifies the driver accordingly and may take appropriate actions. When a driver is warned, he can react timely and thus may avoid critical conditions and accidents. Note that the local detection of hazards, the message dissemination and the notification of drivers must be conducted autonomously, i.e. without any interaction of the driver.

In short, cooperative local danger warning basically comprises the following three main steps:

Detection Process The integral requirement of an LDW application is the autonomous detection of hazards on the road. This is not a trivial task and poses some open research questions (see e.g. [5, 12]). For the scope of this paper, we assume that vehicles are able to individually detect certain hazards without the need of cooperation with other vehicles. In the following, the process of detecting the presence or absence of a hazard is referred to as *experience* of the vehicle.

Message Dissemination After an experience has been made, a corresponding warning message is created by the detecting vehicle and broadcasted accordingly. Receiving vehicles store the remote information and relay the message to other vehicles as long as they travel within a dedicated area. By combining all available information from remote experiences, each vehicle is able to conclude a picture of the road situation ahead of its estimated route. Again, an efficient and scalable message dissemination is beyond the scope of this paper (e.g. [15, 26]). Two kinds of messages are considered: *warning messages* and *revocation messages*. Whenever a vehicle detects a hazard, a corresponding warning message is created. On the other hand, when a vehicle passes the location of a hazard which has been previously reported by other cars, and no hazard is detected, a revocation message is created, informing other cars that the potential hazard has possibly disappeared.

Decision Process The remote experiences received from other cars need not to be evaluated continuously. Instead, this is only necessary if a vehicle has been approaching a potential dangerous area, which has been experienced and reported before, up to a critical distance. In this case, the LDW Application has to decide whether or not to take action or notify the driver. As mentioned before, leading a system into a wrong decision is one of the major threats.

The considerations and results presented in this paper are based on the following system model. Each hazard is assigned three surrounding geographic regions (Fig. 1). The innermost area is thereby denoted as *recognition area*, specifying the area where the hazard can be detected by the on-board sensors of vehicles. Only vehicles inside the recognition area will be actually able to detect the presence or absence of a hazard. This region is enclosed by the *decision area*, which in turn is enclosed by the *dissemination area*. Whenever a vehicle enters the dissemination area of a hazard, it will start to collect and distribute the messages concerning this hazard. As soon as it enters the decision area, the LDW Application decides whether or not to take action or notify the driver. The size of these areas may depend on several factors, like the type of the hazard, the current traffic density and the course of the road network. To simplify matters, for the scope of this article, these areas are assumed to be circular. However, they may be adapted to follow the road network more closely.

[1] Although it is the LDW Application itself, which carries out that task, the terms "car" and "vehicle" will occasionally be used to denote the LDW Application.

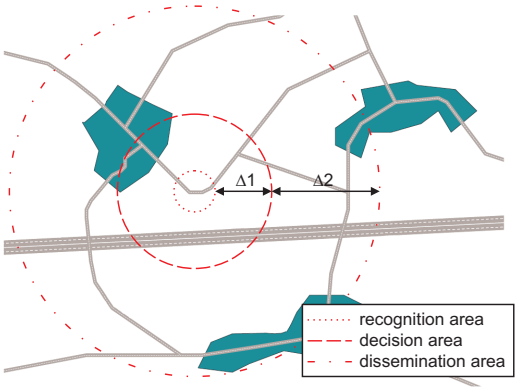

Figure 1. The three geographic areas

The distance between the dissemination and the decision area (marked in Fig. 3.2 as $\Delta 2$) should give vehicles approaching a hazard some time to collect relevant messages before they have to reach a decision. However, a large $\Delta 2$ thereby may lead to a waste of communication resources, because the benefit of the disseminated information decreases with increasing age and distance to its origination. On the other hand, a small $\Delta 2$ may not provide sufficient reception rates, in particular in sparsely connected networks. If the vehicle has approached the hazard up to the distance of $\Delta 1$, it enters the decision area and initiates the decision process. If this process is triggered too early, the quality of the conclusion decreases, because the hazard may change with time. On the other hand, if the experiences are evaluated too late, i.e. the hazard is already very close, timely driver intervention may not be possible anymore (for a more detailed discussion see e.g. [5, 14, 1]).

The process of collecting messages at first and reaching a decision when entering the decision area is called *delayed decision*. This procedure should compensate a possible lack of connectivity, since the application gets some time to collect all relevant messages for the hazard. It is worth mentioning that a vehicle may already be inside the decision area (but outside the recognition area) of a hazard when receiving the first respective message. If so, a decision has to be reached at once. This is denoted as *late decision*, since the decision is not reached when entering the decision area, but is made at a later point in time[2].

3 Security

As pointed out before, security plays an important role in traffic-related inter-vehicle communications. Conventional solutions focus either on securing the communication network or on restricting the access to vehicle components by utilizing well-known security measures, such as digital signatures and trusted hardware. In this way, attackers are prevented from manipulating the network or certain parts of the vehicle. Even when assuming that attackers can neither manipulate the VANET itself (for example, by injecting a fabricated local danger message) nor vehicles (for example, by exchanging certain parts of the hardware), there is still a possibility of manipulation: Since the detection of hazards is based on local sensor readings, an attacker may trick the detection process of his vehicle by "altering the physical environment" around his sensors [6]. In this way, a faked local danger message could be created and subsequently disseminated within the corresponding geographic region in the VANET.

Although required for protection against other forms of attacks, conventional security measures cannot protect against this threat, because cryptographic protection concepts cannot verify information itself. In other words, manipulating sensor readings to stimulate a false message may still result in a perfectly signed and certified message. Revocation of digital credentials due to detected false messages is difficult because of the decentralized nature of VANETs and the problem of automated on-time verification of information in a dynamically changing environment.

Therefore, an additional application-level approach is required. There are two possibilities to mitigate the effects of this threat:

1. Checking local sensor values for plausibility during the *detection process*.

2. Evaluating the plausibility of information received during the *decision process*.

We focus on the second issue, thus hardening the decision process against attacks.

3.1 Related Work

Manipulating sensors and thus affecting the *detection process* can be made much harder by securing car components against removal, alteration or replacement, so they cannot be replaced easily such as proposed by [27] using tamper-resistant technologies. If components have been protected in such a way, a next step is to use a *sensor reasoning* module that catches the most obvious sensor deceptions. But ultimately there is no way of preventing manipulation of sensor readings.

Manipulating the *decision process* (see sect. 4), the question arises whether or not the cars are mostly known to each other. On one hand, if the vast majority of cars is assumed to be known, reputation systems may be used. [11] suggests

[2]The term "late decision" refers to a global time reference. From a local point of view, one could also say that this is an early decision, since the vehicle has no time for collecting additional messages.

a distributed reputation system which establishes individual trust relationships between the participants of a VANET, applying the ideas of reputation in ad hoc network routing, such as in [21] in combination with reputation systems in social contexts as in [16] to the field of LDWs. In this system, the decision, whether or not to consider a reported hazard is mainly based on these individual trust relationships. A detailed analysis, which can be found in [23], shows several weaknesses of this approach.

On the other hand, even if most cars in the vicinity are unknown, information contained in received messages can be collected and evaluated. In [17], a theoretical framework to validate received data in VANETs is presented, based on the most plausible explanation. However, the authors do not provide methods to assess the plausibility of local danger messages and furthermore assume perfect communication, neglecting delays, transmission losses and collisions in practical networks. [6] mentions plausibility checks as a way of increasing security in traffic-related inter-vehicle communication.

Another insider attack as pointed out in [9] is related to resource restrictions of mobile devices such as battery life or computational power. This may lead to nodes that forward messages selectively only if they have an immediate advantage, thus seriously deteriorating the performance of the network. Further work on this topic has been done by [8] and [22]. Since these resource restrictions are less rigorous in cars, we will not deal with this problem in this paper.

3.2 LDW Threats and Security Objectives

We identified four threats on application level, which may arise from the deployment of an LDW Application:

1. **Road Traffic Interference:** Influencing road traffic, which may finally result in accidents (see also [7]).

2. **Subversion of Accountability:** Stimulating wrong accusations against a road user or enable attackers to stay unidentified in LDW application types requiring legal liability as pointed out in [25].

3. **Impairment of Privacy:** Generating movement patterns by exploiting information such as position and time in received messages (see also [18, 10]).

4. **Remote Compromise of Vehicles:** Manipulating vehicles remotely by exploiting existing vulnerabilities. This is possible theoretically, since the LDW Application will have to be connected both to the internal electronic bus systems[3] and to the interface of the wireless network of a vehicle (see also [27]).

[3]e.g. to access sensor readings and to carry out certain actions upon receipt of LDW messages

In the rest of this paper, we will focus on the first threat - Road Traffic Interference.

VANETs are expected to become very large (see also [7]), so messages received will typically originate from participants which are a priori unknown. Therefore, an initial trust relationship between the communication partners cannot be assumed. On the other hand, the high number of participants should allow for multiple experiences of the same hazard by different vehicles, and therefore enable the compensation of false messages, provided that the number of attackers is sufficiently smaller than the number of honest nodes. We thereby assume that a node cannot possess multiple identities, otherwise an attack known as the *Sybil Attack* [13] could render voting-based solutions useless.

The main objective of this work is to minimize false decisions conducted by the decision process. We thereby assume that an attacker may only change the type of a local danger message (i.e. whether it is a warning or a revocation message), by tricking the detection process of his vehicle. Hence, the correctness of the type of a local danger message cannot be guaranteed, so the system has to be robust against this kind of manipulation. To be more specific, *Byzantine Robustness* or at least *Byzantine Detection* [24] is aspired, exploiting the large average number of available experiences. Applying voting schemes to the set of received local danger messages thereby aims to enable the assessment of the plausibility of remote hazard information.

3.3 Attack Scenarios

It is assumed that a certain fraction of the simulated vehicles will misbehave, thus conducting some kind of attack. Since the only possibility to manipulate is the message item containing the type of experience, only malicious data attacks are considered. Two different attacks were identified, namely the *fake attack* and the *flip attack*, which will be explained subsequently. In both cases, the attackers do only cooperate implicitly by acting equally in identical situations, and do not explicitly collude, e.g. by driving in convoys.

During the *fake attack*, each attacker creates a warning message whenever he enters the recognition area of a fictitious hazard. All attackers thereby share the same fictitious hazard. Since that hazard cannot be detected by honest participants, these will create a revocation message once they enter the recognition area. The goal of this attack is to trigger a false positive decision of the LDW Application at the attacked vehicles.

When conducting a *flip attack*, attackers invert the type of experience included in their messages created, whenever entering the recognition area of an actual hazard. Thus, when attackers actually detect the presence of a hazard, they will send a revocation message, while detecting the absence

of a previously existing hazard will result in a warning message. The goal of this attack is to reach false negative decisions at the victims whenever a hazard is present and to reach false positive decisions when that hazard has disappeared.

4 Decision Process

The need for a decision process results from three reasons. First, wrong detections cannot be completely excluded, so there is always the possibility that incorrect information is distributed. Second, the state of the hazard may change with time, so that the received messages show an inconsistent picture. For example, this may be the case when the hazard has recently disappeared. Third, as was shown in the previous section, attackers may try to disturb the system by disseminating wrong local danger messages.

Three important requirements for the decision process were identified:

1. **Adaptivity.** Since VANETs are highly dynamic and the environment is changing continuously, the decision process should quickly adapt to those changes. For example, it should take only a small amount of time until all newly approaching cars detect the disappearance of a previously reported hazard, thus minimizing the amount of wrong decisions.

2. **Robustness.** The necessity of a decision process follows from the understanding that received local danger messages may not always reflect the current state of the environment. It is therefore crucial that a decision method is robust against wrong messages, thus providing correct decisions even under attacks.

3. **Scalability.** Due to the specifics of a certain hazard and the dynamic nature of VANETs, the number of experiences made with regard to a hazard may highly vary. This in turn directly corresponds to the number of local danger messages which can be utilized by the decision process. Hence, the process should perform equally well for most situations.

4.1 Decision Methods

We suggest lightweight plausibility checks, which estimate the plausibility of a reported hazard solely by performing voting schemes on the corresponding received local danger messages.

Four basic decision methods have been developed and analyzed:

1. **Freshest Message:** When a decision has to be reached, only the most recent message of a hazard is considered. If it is a warning message, then a positive decision is reached, if it is a revocation message, a negative decision is made. It is assumed that this decision method will not provide protection against adversaries, however, it should achieve a high adaptivity in attacker-free scenarios, resulting in only a few false decisions.

2. **Majority Wins:** This decision method performs a local voting over all received messages regarding a certain hazard. Duplicates are not considered, hence only distinct messages are counted. If the majority of the messages are warnings, then a positive decision is reached, otherwise a negative decision is made. It is assumed that this decision method provides a high robustness against attacks.

3. **Majority of Freshest X:** This decision method is a combination of the previous two methods. To reach a decision, a vehicle will perform a voting, considering only the recent x distinct messages, regarding the hazard in question.

4. **Majority of Freshest X with Threshold:** Finally, extending the previous decision method with a threshold check results in this decision method. Thereby, it is checked if the distinct messages received so far exceed a certain threshold. If this is not the case, a negative decision is reached for the hazard in question, otherwise the result of the decision process is determined by *Majority of Freshest X*.

5 Simulation

In this section, our simulation toolchain is outlined, followed by a specification of the simulated scenarios.

5.1 Simulation Setup

Our simulation setup consists of two pipelined simulators. The first one, *GenMobTrace*, simulates vehicle movements based on a selected mobility model and a given road map. A peculiarity of this simulator is the computation of the direct reachability between two arbitrary nodes. This is done by taking into account the transmission range and a line-of-sight model, which simulates the obstruction of the communications caused by buildings. It is thereby assumed that buildings exist on both sides of each street, an assumption which can be justified for inner-city scenarios. GenMobTrace is a time-discrete simulator and has a resolution of one second. The simulator generates a trace-file of its simulation run, which is read by the second simulator, *AppSim*.

Simulation area:	8 km^2
Number of vehicles:	250
Simulation time:	1200 sec
Lifetime of a message:	200 sec
Communication range:	400 m
Hazard appearance time:	100 sec
Hazard disappearance time for fake attack:	1100 sec
Hazard disappearance time for flip attack:	600 sec
Start of fake/flip attack:	100 sec
Stop of fake/flip attack:	1100 sec
Diameter of recognition area:	50 m
Diameter of decision area:	300 m
Diameter of dissemination area:	700 m

Table 1. Simulation Parameters

AppSim was developed in order to enable rapid prototyping of applications for VANETs. Besides the application logic, message forwarding algorithms can be simulated, based on the reachability information computed by GenMobTrace. AppSim is also a time-discrete simulator, its resolution is up to one millisecond. For a more detailed description of the simulation setup please refer to [14].

5.2 Simulation Scenarios

In all the simulation scenarios, 250 vehicles were simulated for 1200 seconds on a 8 km^2 section of a digital map of Munich. Vehicles were placed randomly across the street network and move according to a model after KRAUSS [20], while choosing their destinations according to the random waypoint model. The maximum communication range was set to 400 meters, and messages will expire 200 seconds after they were created. A simplified message forwarding algorithm is utilized, which transmits messages to every node which is directly reachable, at every second of simulation time.

A hazard was placed at an intersection in the center of the simulation area, appearing after 100 seconds of simulation time. For fake attacks, this hazard is fictitious and disappears after 1100 seconds of total simulation time, thus resulting in an attack time of 1000 seconds. For flip attacks, this is a real hazard, which is present for 500 seconds. In this case, attackers start their flip attack as soon as the hazard appears, and continue for 500 seconds of simulation time after the hazard has disappeared, hence also resulting in an attack time of 1000 seconds. A summary of the simulation parameters is shown in Tab. 5.2.

The number of messages being considered by *Majority of Freshest X* was set to 22, and the message threshold necessary for *Majority of Freshest X with Threshold* was set to 2. Hence, the latter decision method requires at least three received messages for a hazard in order to be able to reach a positive decision. These values were determined by the analysis of various traffic patterns, considering the above-mentioned parameters.

In order to evaluate the four decision methods introduced, they were simulated first without any attackers, in order to be able to investigate their basic performance. Subsequently, both the *fake attack* and the *flip attack* were simulated with increasing fractions of attackers, ranging from 5% to 40% in steps of 5%, to investigate the robustness of the four decision methods. Hence, 68 simulation runs[4] were conducted in total for the succeeding analysis, which are all based on the same trace-file generated by GenMobTrace. The decision, whether a simulated vehicle is a well-behaving node or an attacker is based on the internal node number generated by GenMobTrace. For a fraction of 5% of attackers, the first 5% of the nodes are considered to be attackers, thus making every set of attackers a superset of the smaller sets of attackers. As initial vehicle positions are assigned randomly, so is the initial distribution of attackers.

6 Results

The simulation results are visualized in Fig. 2 and 3. In order to measure the performance of the decision methods, the percentage of false decisions with respect to the total number of decisions reached within a simulation run was utilized. A decision is considered as false whenever its result does not match the status of the actual hazard at the time when the decision is reached. Both diagrams show the performance of each decision method with respect to an increasing number of attackers, and for the fake attack, also in an attacker-free scenario.

Detailed progression diagrams have been utilized, in order to analyze the simulation results. These diagrams show the development of important key figures during the course of simulation time, and are generated from trace-files of AppSim, for each simulation run. An exemplary progression diagram is shown in Fig. 4.

In the following, we will first discuss the simulation results of each proposed decision method. Subsequently, we will provide a summary and an evaluation of the results.

6.1 Freshest Message

Without any attackers, no false decisions at all were made. This can be explained as follows: When the hazard appears, there are already some vehicles inside its recognition area, which detect the danger at once and disseminate warning messages. Thus, subsequent vehicles are warned

[4]The total number of 68 simulation runs results from 4 decision methods $*$ (1 scenario without attackers + 8 fake attack scenarios + 8 flip attack scenarios).

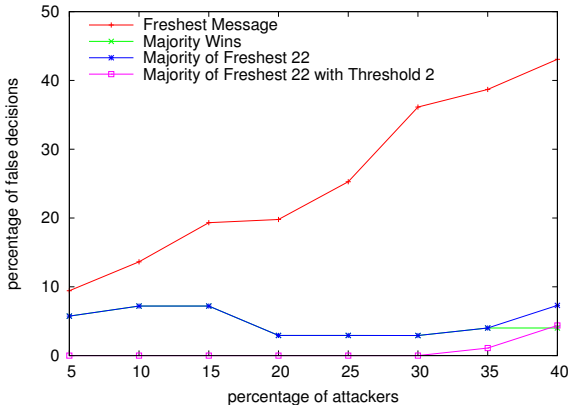

Figure 2. Comparison of the four decision methods regarding the Fake Attack

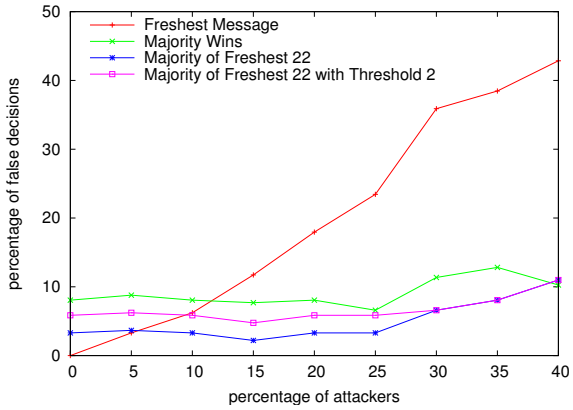

Figure 3. Comparison of the four decision methods regarding both an attacker-free scenario and the Flip Attack

and will make a correct decision. In case the hazard disappears, there has to be a vehicle which is inside the decision area but outside the recognition area, and which is informed about the hazard, thus having already reached a positive decision. This vehicle has to enter the recognition area of the hazard after it disappeared and before any other vehicle makes a decision. In this way, a revocation message is created and distributed, informing other vehicles about the disappearance of the hazard.

Looking at the fake attack, it reveals that the number of false decisions increases almost linearly with the number of attackers. The fraction of false decisions thereby is near the fraction of attackers. As expected, a protection against this type of attack cannot be identified.

Regarding the protection against the flip attack, the situation seems similar. However, as long as there are less than 20% of attackers, the results are better than those of the fake attack. The reason is that at the beginning of the fake attack, there are a lot of false decisions, which can be explained as follows: As soon as the first warning message is disseminated (which is faked), all vehicles inside the decision area have to reach a late decision, solely resulting in false decisions. In contrast thereto, at the beginning of the flip attack, most likely there are already some messages around, if the hazard was recognized by an honest vehicle at first. So, for this attack, attackers do not necessarily possess the initial advantage they have for the fake attack. With an increasing number of attackers, the initial advantage of the fake attack becomes negligible and the number of false decisions for the fake attack and the flip attack converge.

6.2 Majority Wins

It reveals that Majority Wins produces a non-negligible number of false decisions even when there are no attackers around. Analyzing the corresponding progression diagram, it shows that all of these false decisions take place after the hazard has disappeared, and persist as long as there are more warning messages than revocation messages around. We denote this period of time as the *second adaption phase*, which is influenced by the lifetime of the messages. It lasts from the point in time the hazard actually disappears until about $(disappearance\ time) + (message\ lifetime)/2$. Unlike the fake and the flip attack, there is no first adaption phase when there are no attackers nearby.

Looking at the fake attack, the number of false decisions is limited and does not significantly rise with an increasing number of attackers. This is because all false decisions are reached at a short period of time, after the faked hazard is announced. We denote this period of time as the *first adaption phase*, it is usually much shorter than the second adaption phase. After this period, the number of revocation messages continuously outnumbers the number of warning messages, thus resulting in exclusively correct decisions within the remaining simulation time. The decline of false decisions between 15% and 20% of attackers (Fig. 2) is an interesting aspect. In the former case, a lot of vehicles are within the decision area when the first attacker is sending his faked warning message, so all of these vehicles have to reach a late decision, resulting in false decisions. In the latter case, since there are more attackers around, the attack is launched at an earlier point in time. Because there are less vehicles in the decision area at that time, less false decisions are made.

Considering the flip attack, the number of false decisions remains almost constant and even decreases marginally until a fraction of 25% of attackers is reached. Looking at the corresponding progression diagrams, it reveals that up to this point, false decisions are still reached only within the

second adaption phase, thus providing robustness against false messages. The effect of a decrease of false decisions with an increasing fraction of attackers can be explained by an increasing number of (false) revocation messages before the hazard disappears. Therefore, the difference between warning and revocation messages is smaller, which shortens the second adaption phase. The reverse effect, namely an extension of the second adaption phase due to more false warning messages after the hazard has disappeared occurs with 30% of attackers, resulting in more false decisions. Additionally, there are now some false decisions for a short period of time after the hazard appears. Just like for the fake attack, this period of time is denoted as the *first adaption phase*, and is usually much shorter than the second adaption phase. The decrease of false decisions at 40% of attackers can again be explained by a shortening of the second adaption phase.

Comparing both attacks, it reveals that the fake attack is handled in a substantially better way. This is because there is no second adaption phase for this attack. Since for both attacks, all false decisions are reached within the adaption phases, a shortening of these two phases may help to decrease the number of false decisions. This is aspired by the next decision method.

6.3 Majority of Freshest X

By considering only the x most recent messages for the determination of the majority, a shortening of the adaption phases is aspired. As noted before, x was set to 22 for the simulation runs conducted.

The false decisions of Majority of Freshest X in a scenario without attackers are substantially lower than those for Majority Wins. This is because the duration of the second adaption phase is successfully reduced (as before, there is no first adaption phase in such a scenario).

The performance with regard to the fake attack is almost identical as compared to that of Majority Wins. This is because the first adaption phase is very short, and there is no second adaption phase for this attack, so this decision method does not improve the protection against the fake attack. Furthermore, the result for 40% of attackers is worse than that of Majority Wins. An analysis of the progression diagram reveals that in this case, there are false decisions after the first adaption phase, which are caused by considering only the recent 22 messages received.

Looking at the flip attack, in general there are less false decisions as compared with Majority Wins, except for a fraction of 40% of attackers. The reason for the latter is similar to the corresponding scenario with the fake attack, in such a way that there is a considerable number of false decisions between the two adaption phases. An analysis of the progression diagrams reveals that this kind of false decisions can be observed first with 30% of attackers.

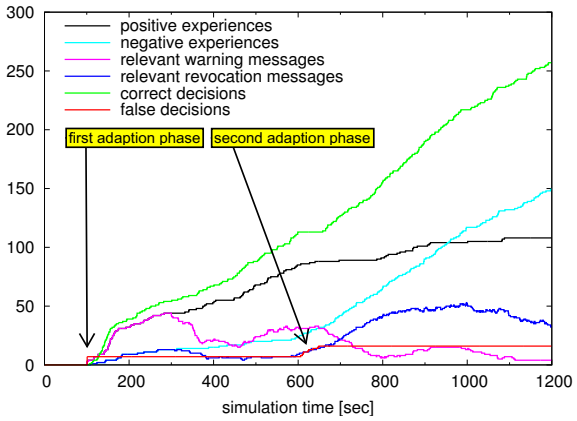

Figure 4. Progression diagram for Majority of Freshest 22 with Threshold 2, Flip Attack, 20% Attackers

In summary, this decision method significantly improves the protection against the flip attack. However, considering fake attacks, the performance is not enhanced. As almost all false decisions are reached at the beginning of this attack, a shortening of the first adaption phase would generally improve the performance of this decision method with respect to fake attacks.

6.4 Majority of Freshest X with Threshold

The last decision method tested is an enhancement of Majority of Freshest X. The goal is to reduce the number of false decisions regarding fake attacks, by utilizing a message threshold. In order to be able to reach a positive decision at all, a minimum number of messages has to be received at the time a decision is made. As noted before, the threshold was set to 2, thus requiring at least 3 messages in order to invoke Majority of Freshest X.

The performance in a scenario free of attackers is slightly worse than that of Majority of Freshest X. This is because the threshold adds a first adaption phase, in which all vehicles reach a false decision. As soon as there are at least 3 warning messages disseminated, correct decisions are reached until the second adaption phase.

Looking at the results of the fake attack, it reveals that this decision method provides a significant protection against this type of attack. Up to a fraction of 30% of attackers, no false decisions are reached at all. With 35% of attackers, a small first adaption phase emerges, and with 40% of attackers, false decisions are reached in the middle of the attack.

Considering the protection against flip attacks, it shows that it is comparable to the one provided by Majority of Freshest X. However, the number of false decisions is generally higher for fractions of attackers below 35%. This is a direct effect of the utilized threshold, which results in false decisions after the hazard appears, even when there are no attackers nearby at that time.

Altogether, this decision method almost completely avoids false decisions with regard to fake attacks, while its performance on flip attacks slightly decreases, when compared with Majority of Freshest X (Fig. 3).

6.5 Summary

The analysis of the simulation results revealed some interesting aspects. Since the results of Freshest Message differ completely from those of the other decision methods, this decision method is not considered in the following. For all remaining decision methods, nearly all false decisions are reached within two adaption phases. For real hazards, the second adaption phase is inevitable because of the change of the state of the hazard. In addition, there may be a first adaption phase, depending on the utilized decision method. Concerning the fake attack, there exists only the first adaption phase. Looking at the flip attack, the situation is based on that of the real hazard, whereas a first adaption phase is becoming more probable with an increasing number of attackers.

Usually, the first adaption phase is much shorter than the second adaption phase, so the number of false decisions reached in the first phase is usually lower than that reached in the second phase. The simulation results indicate that within the flip attack, the second adaption phase may be extended as well as shortened by attackers. Furthermore, late decisions may lead to a lot of false decisions, which are reached at once. In the simulated scenarios, late decisions exclusively occur for a short time after a real or faked hazard was announced. This is because in our simulation scenarios the VANET is of sufficient density, so vehicles approaching a hazard will be able to collect the respective messages before reaching a decision. Whenever an attacker is able to send the first message concerning a hazard, he will trigger late and false decisions at all vehicles currently within the decision area. This is especially problematic for fake attacks.

6.6 Evaluation

Of the four decision methods tested, Freshest Message achieves the highest adaptivity and thus produces the best results in the absence of attackers. However, it provides no protection against attacks, and is therefore considered as an inappropriate solution.

Majority Wins provides a significant robustness against false messages. However, the adaptivity is not very high, which explains the considerable percentage of false decisions in the absence of attackers. Interestingly, both attacks do not scale, so for most cases in the simulated scenarios, more attackers do not achieve more false decisions.

Evaluating the results of Majority of Freshest X, this decision method seems to offer a good tradeoff between the high adaptivity of Freshest Message and the high robustness of Majority Wins. However, simulation results indicate that this decision method may produce false decisions outside the adaption phases, which is the price for its higher adaptivity. The choice of the parameter x enables an adjustment between adaptivity and robustness, which can be considered as contrary goals.

Finally, Majority of Freshest X with Threshold provides the best result with regard to the fake attack (Fig. 2). However, for the flip attack, it performs slightly worse than Majority of Freshest X, which provides the best results for this attack (Fig. 3). However, as false positive decisions could possibly be more dangerous than false negative decisions, Majority of Freshest X with Threshold is considered to be the best solution of the decision methods introduced.

Regarding scalability, it is assumed that all decision methods can be utilized with an increasing number of vehicles, since the number of relevant messages can be limited by an adaptive message lifetime. However, without further simulation, no statements can be made about scenarios with fewer vehicles, since the effects of reduced connectivity may affect the dissemination of messages, which in turn affects the decision process.

Concerning the percentages of false decisions noted in Fig. 2 and 3, these values have to be treated with care. As the simulation results have shown, for most decision methods, false decisions only occur within the adaption phases, whose length is independent from the overall attack time. Therefore, increasing the attack time would probably lead to more correct decisions, while keeping the number of false decisions constant. Hence, attacks which last longer would produce a lower percentage of false decisions, while shorter attacks would increase the percentage of false decisions.

7 Conclusions

Our simulation results indicate that even simple decision methods, which are solely based on the number of received distinct messages, might provide sufficient protection against attackers. Especially the results of the fourth decision method, *Majority of Freshest X with Threshold*, look promising. However, a lot of additional parameters have to be investigated with the help of simulations, such as different traffic scenarios (e.g. highways and rural roads), a more specific radio and message propagation model or a

more detailed traffic simulation, in order to further analyze this decision method. Furthermore, an important issue is the determination of the two parameters of *Majority of Freshest X with Threshold*. This has to be based on the current traffic situation, which therefore has to be analyzed automatically by the LDW Application.

In order to further improve the performance of the decision process, it is suggested that additional decision methods are considered. For example, taking into account local sensor readings of the vehicle which has to reach a decision and relate them to the reported hazard may help to reduce the remaining number of false decisions.

References

[1] The Network on Wheels Project. Website, 2004. `http://www.network-on-wheels.de`.

[2] The PATH Project. Website, 2005. `http://www-path.eecs.berkeley.edu/`.

[3] The Vehicle Safety Communications Project. Website, 2005. `http://www-nrd.nhtsa.dot.gov/pdf/nrd-12/CAMP3/pages/VSCC.htm`.

[4] The WILLWARN Project. Website, 2005. `http://www.prevent-ip.org/en/prevent_subprojects/safe_speed_and_safe_following/willwarn/`.

[5] C. Adler and M. Strassberger. Putting Together the Pieces - A Comprehensive View on Cooperative Local Danger Warning. In *Proceedings the 13th ITS World Congress and Exhibition on Intelligent Transport Systems and Services (ITS'06)*, London, UK 2006.

[6] A. Aijaz, B. Bochow, F. Dötzer, A. Festag, M. Gerlach, R. Kroh, and T. Leinmüller. Attacks on Intervehicle Communication Systems - an Analysis. In *Proceedings of the 3nd International Workshop on Intelligent Transportation*, Hamburg, Germany, March 2006.

[7] J. Blum and A. Eskandarian. The Threat of Intelligent Collisions. *IT Professional*, 6(1):24 – 29, January/February 2004.

[8] S. Buchegger and J.-Y. Le Boudec. The Selfish Node: Increasing Routing Security for Mobile Ad Hoc Networks. Technical report, 2001.

[9] L. Buttyán and J.-P. Hubaux. Enforcing Service Availability in Mobile Adhoc WANs. In *1st IEEE/ACM Workshop on Mobile Ad Hoc Networking and Computing (MobiHOC 2000)*, Boston, Massachusetts, 2000.

[10] F. Doetzer. Privacy Issues in Vehicular Ad Hoc Networks. In *Workshop on Privacy Enhancing Technologies*, Cavtat, Croatia, May 2005.

[11] F. Doetzer, L. Fischer, and P. Magiera. VARS: A Vehicle Ad-Hoc Network Reputation System. In *Proceedings of the IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks*, Taormina, Italy, 2005.

[12] F. Doetzer, T. Kosch, and M. Strassberger. Classification for traffic related inter-vehicle messaging. In *Proceedings of the 5th IEEE International Conference on ITS Telecommunications*, Brest, France, 2005.

[13] J. R. Douceur. The Sybil Attack. In *IPTPS '01: Revised Papers from the First International Workshop on Peer-to-Peer Systems*, pages 251–260, London, UK, 2002. Springer-Verlag.

[14] S. Eichler, B. Ostermaier, C. Schroth, and T. Kosch. Simulation of Car-to-Car Messaging: Analyzing the Impact on Road Traffic. In *Proceedings of MASCOTS '05*, Washington, DC, USA, 2005.

[15] S. Eichler, C. Schroth, T. Kosch, and M. Strassberger. Strategies for Context-Adaptive Message Dissemination in Vehicular Ad Hoc Networks. In *Proceedings of the Second International Workshop on Vehicle-to-Vehicle Communications (V2VCOM 2006)*, 2006.

[16] S. Fähnrich and P. Obreiter. The Buddy System - A Distributed Reputation System Based On Social Structure. Technical Report 2004-1, Universität Karlsruhe, Faculty of Informatics, February 2004.

[17] P. Golle, D. Greene, and J. Staddon. Detecting and Correcting Malicious Data in VANETs. In *VANET '04: Proceedings of the 1st ACM international workshop on Vehicular ad hoc networks*, pages 29–37, New York, NY, USA, 2004.

[18] J.-P. Hubaux, S. Capkun, and J. Luo. The Security and Privacy of Smart Vehicles. *IEEE Security and Privacy*, 2(3):49–55, 2004.

[19] T. Kosch. Local Danger Warning based on Vehicle Ad-hoc Networks: Prototype and Simulation. In *Proceedings of the 1st International Workshop on Intelligent Transportation (WIT 2004)*, pages 43–47, Hamburg, Germany, 2004.

[20] S. Krauss. *Microscopic Modeling of Traffic Flow: Investigation of Collision Free Vehicle Dynamics*. PhD thesis, Universität zu Köln, 1998.

[21] S. Marti, T. J. Giuli, K. Lai, and M. Baker. Mitigating Routing Misbehavior in Mobile Ad Hoc Networks. In *Proceedings of MOBICOM 2000*, 2000.

[22] P. Michiardi. *Cooperation enforcement and network security mechanisms for mobile ad hoc networks*. PhD thesis, Eurecom, 2004.

[23] B. Ostermaier. Analysis and Improvement of Inter-Vehicle Communication Security by Simulation of Attacks. Master's thesis, Technische Universität München, October 2005.

[24] R. Perlman. *Network Layer Protocols with Byzantine Robustness*. PhD thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 1988.

[25] M. Raya and J.-P. Hubaux. Security Aspects of Inter-Vehicle Communications. In *Swiss Transport Research Conference (STRC)*, 2005.

[26] Q. Sun and H. Garcia-Molina. Using Ad-hoc Inter-vehicle Networks for Regional Alerts. Technical report, Stanford University, 2004.

[27] A. Weimerskirch, C. Paar, and M. Wolf. Cryptographic Component Identification: Enabler for Secure Vehicles. In *Proceedings of Vehicular Technology Conference 2005*, Dallas, USA, September 2005.