Open issues in differentiating misbehavior and anomalies for VANETs

Rens W. van der Heijden and Frank Kargl
Institute of Distributed Systems
University of Ulm
Albert-Einstein-Allee 11, 89081, Ulm, Germany
Email: {rens.vanderheijden,frank.kargl}@uni-ulm.de

***Abstract*—This position paper proposes new challenges in data-centric misbehavior detection for vehicular ad-hoc networks (VANETs). In VANETs, which aim to improve safety and efficiency of road transportation by enabling communication between vehicles, an important challenge is how vehicles can be certain that messages they receive are correct. Incorrectness of messages may be caused by malicious participants, damaged sensors, delayed messages or they may be triggered by software bugs. An essential point is that due to the wide deployment in these networks, we cannot assume that all vehicles will behave correctly. This effect is stronger due to the privacy requirements, as those requirements include multiple certificates per vehicle to hide its identity. To detect these incorrect messages, the research community has developed misbehavior data-centric detection mechanisms, which attempt to recognize the messages by semantically analyzing the content. The detection of anomalous messages can be used to detect and eventually revoke the certificate of the sender, if the message was malicious. However, this approach is made difficult by rare events –such as accidents–, which are essentially anomalous messages that may trigger the detection mechanisms. The idea we wish to explore in this paper is how attack detection may be improved by also considering the detection of specific types of anomalous events, such as accidents.**

I. Introduction

The ultimate goal of the development of vehicular ad-hoc networks (VANETs) is to improve both safety and efficiency of road transportation. Although applications in the near future are designed with drivers in mind, it is conceivable that VANETs will be combined with recent developments on automated driving in the future. For this reason, it is especially important that we develop a secure communication platform from the ground up. Another important aspect is user-acceptance, which will deteriorate if the system reports incorrect warnings. Developments of these security mechanisms have required several innovations from the security community, especially in order to deal with the new challenges that VANETs pose. Notably, these challenges include the strict privacy requirements, bandwidth constraints, the ephemeral nature of the network, lack of permanent access to infrastructure and the public nature of the messages that are exchanged. We refer the interested reader to [1] for a more detailed discussion of these challenges. In summary, it has become clear that pro-active security mechanisms like digital signatures are not sufficient to provide security in VANETs; the research community has proposed complementary reactive security mechanisms, which detect malicious messages even when they are authentic. This process is also known as misbehavior detection.

Over the past decade, many security mechanisms for detecting misbehavior have been developed. These can be classified as either data- or node-centric, representing a focus the correctness of the content or the trust in a network participant respectively. A deeper classification and its applications to other types of networks can be found in [2].

During our study of the literature, we have observed that the evaluation of these mechanisms is typically performed against normal/baseline behavior and that same scenario containing one or more attacker. In this paper, we will focus on data-centric mechanisms that look at the data alone; good examples of these mechanisms include those proposed by Leinmüller et al. [3]; among other mechanisms, they describe the detection of unrealistic speeds, e.g. a claimed speed of 500km/h. One might imagine a similar detection mechanism could be used for sudden speed drops, but this could lead to problems: a crashing vehicle may also portray an unusual pattern (e.g., a very sharp drop in speed). Clearly we cannot classify a crashing vehicle as a malicious one. This illustrates that for data-centric misbehavior detection mechanisms, an important open issue is that we cannot automatically classify anomalous data as being malicious.

Partially in response to this issue, several authors have proposed a new class of data-centric detection mechanisms, which use the driver's response as a model for correctness [4], [5]. The idea is that a driver will correctly react to scenarios such as an accident, even when all detection mechanisms fail. These mechanisms can be used to eventually expel the malicious senders from the network. This can significantly reduce the impact of attackers and damaged sensors, but it does not prevent the spread of malicious or incorrect messages throughout the network until the driver should already have responded to the event. This is undesirable from a user-perspective: if the driver receives potentially false warnings all the time, the user acceptance of the applications will go down, or the users may simply turn the system off. Therefore, we identify a need to combine both approaches; we need detection mechanisms that use the driver's behavior as a baseline, as well as detection mechanisms that prevent malicious messages before they arrive.

In the remainder of this paper, we discuss two open issues regarding data-centric misbehavior detection: the similarity with the detection of events, such as a crashing vehicle, and the evaluation of misbehavior detection. Before discussing these open issues, we discuss how multiple misbehavior detection mechanisms can be combined into a framework, and we

elaborate on the state of our current research. Our research set out to improve detection accuracy, and as we discuss the open issues we will elaborate how our framework may help solve them.

II. Frameworks for Misbehavior Detection

In the literature, several authors have already observed this challenge and proposed the combination of several different mechanisms [3], as well as frameworks that allow more complex operations [6], [7]. Specifically, Golle et al. [7] discussed techniques to decide which conclusion is most likely based on a set of received messages. On the other hand, Raya et al. [6] and other authors have discussed the more abstract idea of trust between participants in the network. The idea of their work is to use node- and data-centric information to provide a trust value for each participant. This trust value is then used to predict the likelihood that a message is incorrect.

In our ongoing research, we are developing a framework to unify the detection of misbehavior in VANETs. Using subjective logic [8], we build a modular system that can incorporate arbitrary amounts of detection mechanisms and usefully combine their results. The goal is to provide filtering of malicious messages, the exchange of evidence between nodes and the tools for local or global revocation. Although the specific advantages of our framework are beyond the scope of this paper, a core focus for us is the idea that different mechanisms may perform poorly or very well depending on the specific scenario. In order to cope with this, we keep the results from multiple mechanisms, and allow the expression of uncertainty about a result (for which subjective logic is our chosen mathematical representation). However, we realized that because a mechanisms' (un)certainty about a result may depend on the context in which it is running; is it designed for highway or urban scenarios? Is there a connection available to a back-end system or certificate authority?

III. Combining Security and Functionality

We have previously noted that there is a significant parallel between data-centric misbehavior detection and the recognition of legitimate events that vehicles should notify their drivers about. In particular, we note that pure data-centric misbehavior detection is very hard especially for this reason: a sufficiently powerful attacker will always attempt to imitate a legitimate event as precise as possible. On the other hand, the detection of legitimate events is challenging because sensor data may not be reliable, or have a significant margin of error. We propose that these mechanisms can complement each other.

Specifically, this position paper proposes the idea that our framework can facilitate the mechanisms by providing the appropriate context based on the history. This history, or the trustworthy subset there-of, can be provided to a mechanism that is designed to recognize a particular scenario – for example, a crashing vehicle. Moreover, this process could be triggered by detection mechanisms, which typically detect such anomalous events. An approach for this process is illustrated in Figure 1. This figure shows an arriving message (1), which triggers a misbehavior detection mechanism to detect an anomaly (2). This allows the situation recognition to create a context (3). This context can be used to update

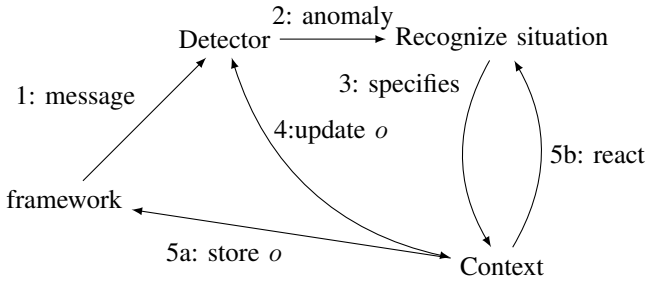

Fig. 1. This figure shows how our framework might detect and adapt to a legitimate anomalous event.

the detection mechanism (4), or it can be used to modify the opinion directly. The opinion can then be stored (5a), and if necessary, a reaction can be processed (5b). In practice, this means that a single misbehavior detection mechanism may detect a sudden drop in speed, which it recognizes as a potential attack. Instead of discarding the message, or marking it as untrustworthy, the message and associated history is first passed through a mechanism that searches for a specific pattern –in this example, it could be a crash on a highway (the highway is the context in this case). Because this mechanism identifies the message as part of a crash, the framework updates the trust values towards higher uncertainty.

We propose that our framework extensions may be applied to decouple the development of event detection mechanisms for the different settings (e.g., highway or urban), as well as the specific events that are to be detected (e.g., traffic jams or crashing vehicles). By describing the individual events for specific settings, a much clearer and more accurate event detection mechanism can be developed, which allows for more certainty regarding the correctness of the misbehavior detection mechanisms. In addition, it allows us to avoid significant pitfalls during detection (of both attackers and events) caused by the attacker producing messages that relate to settings that are distinct from the actual situation. We can decouple the correct recognition of a scenario from validating the misbehavior detection mechanism, which reduces the required simulation time and the effort required to design the simulations. Ideally, this decoupling could even allow a formal proof for individual components, which makes the analysis stronger.

IV. Examples

We now present two brief examples to motivate the decoupling of event and attack detection. Consider an urban setting, with three vehicles driving on a road at about 50km/h. Just before a small side-street, the second vehicle breaks hard to be able to turn into the street. This sudden drop in speed will be accompanied by a break warning DEMN, and may be considered anomalous (i.e., suspicious) by several misbehavior detection mechanisms, due to the absence of an accident. Nevertheless, it is intuitively clear that this is not misbehavior, but rather poor driving, because the driver braked too late.

Similarly, consider a highway in a dense forest, where three vehicles driving behind each other, with an average speed of around 120km/h. In this situation, the first vehicle performs a hard break because it detected a stray animal on the road. Again, the sudden break warning may trigger a

misbehavior detection mechanism in the following vehicles to detect misbehavior, until they detect the animal.

In both cases, the context (the side street and the dense forest, respectively) provides an additional explanation for the detected event that might otherwise have identified the sending vehicle as malicious. Similarly, attack detection may be improved by this decoupling, as it allows the detection of attacks that imitate the incorrect context. For example, when an attacker attempts to create a high-speed crash scenario based on a highway setting in an urban one, we can detect that the sequence of messages does not match the twists and turns of the road. In addition, the approach simplifies the development of misbehavior detection mechanisms, because they no longer need to provide a generic mechanism capturing all possible scenarios, but rather can be designed to deal with specific scenarios.

V. Conclusion

In this position paper, we have proposed several ideas on how to improve data-centric misbehavior detection in VANETs. We have briefly discussed existing work, including approaches that attempt to provide a framework for general misbehavior detection, and we have pointed out several open issues. We have then proposed several ideas that can be used to improve data-centric misbehavior detection and its evaluation, which we hope provide the material for an exciting discussion of the topic during the Fachgespräch.

Acknowledgment

The authors would like to thank the participants of the previous edition of the Fachgespräch Inter-vehicle Communication for their insightful discussion and ideas, parts of which have inspired this work.

References

[1] M. Raya and J.-P. Hubaux, "Securing vehicular ad hoc networks," *Journal of Computer Security*, vol. 15, no. 1, pp. 39–68, Jan. 2007. [Online]. Available: http://iospress.metapress.com/content/CH4D4DG8YL2QHR0W

[2] R. van der Heijden, S. Dietzel, and F. Kargl, "Misbehavior Detection in Vehicular Ad-hoc Networks," in *1st GI/ITG KuVS Fachgespräch Inter-Vehicle Communication (FG-IVC 2013)*, Innsbruck, Austria, February 2013.

[3] T. Leinmüller, E. Schoch, F. Kargl, and C. Maihöfer, "Decentralized position verification in geographic ad hoc routing," *Security and Communication Networks*, vol. 3, no. 4, pp. 289–302, Jul. 2010. [Online]. Available: http://doi.wiley.com/10.1002/sec.56

[4] W. Bamberger, J. Schlittenlacher, and K. Diepold, "A trust model for intervehicular communication based on belief theory," in *2010 IEEE Second International Conference on Social Computing (SocialCom)*, 2010, pp. 73–80.

[5] M. Ghosh, A. Varghese, A. Gupta, A. A. Kherani, and S. N. Muthaiah, "Detecting misbehaviors in VANET with integrated root-cause analysis," *Ad Hoc Networks*, vol. 8, no. 7, pp. 778–790, Sep. 2010. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S157087051000034X

[6] M. Raya, P. Papadimitratos, V. D. Gligor, and J.-P. Hubaux, "On data-centric trust establishment in ephemeral ad hoc networks," in *The 27th Conference on Computer Communications IEEE INFOCOM 2008*. IEEE, Apr. 2008, pp. 1238–1246.

[7] P. Golle, D. Greene, and J. Staddon, "Detecting and correcting malicious data in VANETs," in *Proceedings of the 1st ACM international workshop on Vehicular ad hoc networks*, ser. VANET '04. ACM, 2004, p. 2937. [Online]. Available: http://dl.acm.org/citation.cfm?id=1023881

[8] A. Jøsang, R. Hayward, and S. Pope, "Trust network analysis with subjective logic," in *Proceedings of the 29th Australasian Computer Science Conference - Volume 48*, ser. ACSC '06. Darlinghurst, Australia, Australia: Australian Computer Society, Inc., 2006, pp. 85–94. [Online]. Available: http://dl.acm.org/citation.cfm?id=1151699.1151710