Proceedings of the 2nd GI/ITG KuVS Fachgespräch Inter-Vehicle Communication (FG-IVC 2014)

Luxembourg, February 2014

Raphael Frank, University of Luxembourg

Markus Forster, University of Luxembourg

Christoph Sommer, University of Innsbruck

Frank Kargl, Ulm University

Thomas Engel, University of Luxembourg

www.securityandtrust.lu

University of Luxembourg • Interdisciplinary Centre for Security, Reliability and Trust • 4, rue Alphonse Weicker • L-2721 Luxembourg-Kirchberg

Proceedings of the 2nd GI/ITG KuVS Fachgespräch Inter-Vehicle Communication (FG-IVC 2014)

February 20-21, 2014, Luxembourg City, Luxembourg

Preface

GI/ITG KuVS Fachgespräch was organized on the topic of inter-vehicular communication. Discussions revolved around both the state of the art and around future directions of inter-vehicle communication research, from physical layer optimizations to novel applications of vehicular networks and from microscopic evaluation metrics to problems of scale, crime, and privacy. The goal of this new Fachgespräch was again to bring together talented young researchers to follow-up on the discussions.

Over the last few years, significant efforts are being carried out by industry, academia and government agencies to improve driving safety, increase vehicle traffic efficiency and decrease fuel consumption by exploiting vehicular communications and networking technologies. These technologies, which are generally referred to as VANET (Vehicular Ad Hoc Networks), include Vehicle-to-Infrastructure (V2I), Vehicle-to-Vehicle (V2V) communications and can be based on short- and medium-range communications as well as on cellular systems.

In February 2013 the first edition of the GI/ITG KuVS Fachgespräch was held in Innsbruck, Austria. Due to its success, the Fachgespräch was reorganized in February 2014, this time in Luxembourg City, Luxembourg. With more than 25 participants, the event was well attended. Young researchers have presented twelve papers on different topics around IVC. Every presentation was followed by a discussion leading to a dynamic interaction between the participants. In addition to an academic keynote held by Prof. Claudio Casetti, a VEINS tutorial has been organized.

February 2014

Raphael Frank
Markus Forster
Christoph Sommer
Frank Kargl
Thomas Engel

Contents

DRIVE: An evaluation in a Closed Border Environment

Markus Forster, Raphael Frank, Thomas Engel

Interdisciplinary Centre for Security, Reliability and Trust

University of Luxembourg, 1359, Luxembourg

{markus.forster|raphael.frank|thomas.engel}@uni.lu

Abstract—In this short paper we address the problem of abrupt breakdowns in traffic flow as a consequence of high traffic demand in combination with small driver inaccuracies. First we provide a brief overview on a *Cooperative Advanced Driver Assistance System* (CADAS) called *Density Redistribution through Intelligent Velocity Estimation* (DRIVE) that has been proposed to mitigate jam formations. Next, we perform an analysis on the traffic dynamics in a close border environment that enables us to run simulations with well defined vehicular densities.

Keywords— *Traffic Modeling, Vehicular Networks, Congested Flow, Shock Waves*

I. INTRODUCTION

Traffic demand has significantly increased over recent decades. The traditional solution to face this problem has been the extension of the road network by constructing new roads or adding lanes to already existing ones. This strategy however has come to a point where in most areas physical road extensions are no longer possible [1]. As a possible remedy, one should also consider the fact that, due to current completely uncoordinated vehicular traffic, the already available capacity is not fully exploited. Simulations and empirical analysis have shown that vehicular traffic is in *free flow* only for relatively low densities when interactions between vehicles are negligible [2].

To improve the current situation it is necessary to rely on *Advanced Driver Assistance Systems* (ADAS) systems that equip vehicles with several sensors, enabling them to recognize their surroundings and to notify motorists to take corrective actions or, in some cases, to implement such corrections automatically. An example of such a system is the *Adaptive Cruise Control* (ACC) that can maintain minimum safety headway to the vehicle ahead. Such systems can react much faster than a human driver to sudden downstream vehicle maneuvers and therefore allow to use smaller safety distances [3].

It has been until recently that researchers consider the flow dynamics to increase the efficiency of *Cooperative Adaptive Cruise Control* (CACC). In [4], the authors propose a strategy to stabilize traffic flow by periodically beaconing relevant traffic information to close-by vehicles. Vehicles that detect perturbations downstream try to keep a larger gap to their predecessor in order to compensate traffic inhomogeneity. They show that at an equipment rate of 30% free flow can be restored even at high traffic densities. An improvement of this work has been proposed by the same authors in [5], where the model includes the effect of the human reaction time, confirming the previous findings.

In this paper we perform an analysis of a previous spec-ified network protocol used within a *Cooperative Advanced Driver Assistance System* (CADAS) that connects vehicles via an *Vehicular Ad Hoc Network* (VANET) and lets them exchange recent and relevant traffic information.

The proposed DRIVE protocol [6] broadcasts with an event driven messaging scheme critical information such as position and velocity. This information can be used by following vehicles far behind to learn about the traffic situation downstream. This information is collected and evaluated in a way that provides a consistent picture of the traffic situation ahead. By using the *Lighthill-Whitham-Richard model* (LWR) [7], [8], it is possible to estimate a density gradient between two communicating vehicles and use this information for velocity estimations. The aim of this protocol is to redistribute oncoming vehicles in a way that prevents congestion shock waves from forming. We show that, by giving individual velocity recommendations, even with a low number of equipped vehicles one can achieve significant improvements in overall traffic flow and average travel times. Since this does not lead to a reduction in safety distance, full system coverage is not required to ensure collision-free driving.

II. PROTOCOL DESCRIPTION

The aim of the presented protocol is to improve vehicular flow in situations of high traffic demand without requiring excessive network resources. Furthermore, the protocol should improve the overall traffic flow even with low equipment rates. The flow improvement is achieved by the use of *Density Redistribution through Intelligent Velocity Redistribution* (DRIVE). For a full description of the protocol please refer to the original paper [6].

In the next subsection we provide an overview of the messaging mechanism of DRIVE. In the next but one subsection we provide a brief insight of the recommended velocity estimation. Finally we demonstrate the execution of DRIVE with an example.

A. Message propagation

The network protocol, used for communication of DRIVE messages is fully compliant with the IEEE 802.11p standard for *Wireless Access in Vehicular Environments* (WAVE) [9], [10].

For our analysis, we assume channel characteristics as listed in Table I. The protocol is divided into three main parts. First part is the *Message Dissemination*, second part is the *Message Reception* and the third part describes the *Message*

Forwarding. A complete description of those three phases is given in the next subsections.

TABLE I: Properties of radio channel

PHY 802.11p	
Frequency	5.9 Hz
Maximum Sending Power	20 mW
Thermal Noise	−110 dBm
Sensitivity	−89 dBm
Signal Attenuation Threshold	−89 dBm
MAC 1609_4	
Transmission Power	20 mW
Bitrate	18 Mbps
Service Channel	2

1) Message Dissemination: DRIVE is an event driven protocol, meaning that messages are only disseminated when a certain event occurs. This implies that the WAVE beaconing is not used in our analysis.

The event to trigger the dissemination of a DRIVE notification is a significant slowdown or a unusually low velocity. A slowdown is considered to be significant if the actual velocity is 10 km/h lower than the velocity one second ago. Due to the fact, that our protocol has been designed for highway scenarios velocities lower than 50 km/h are considered as low velocities.

When detecting a slowdown event, an equipped vehicle broadcasts a notification message containing the information:

$$m_{\mathrm{h}} = [id, x, y, t, v, TTL] \quad (1)$$

where $\mathrm{h} \in \mathbf{H}$ is a unique message identifier and id is a unique identification of the originating vehicle. x and y are the GPS-coordinates of the originator. The remaining values are the timestamp of the originator at message creation (t), the actual velocity of the originator at message creation (v) and a preselected Time To Life (TTL).

2) Message Reception: The protocol will only check every second if a relevant notification has been received. This implies that messages arriving within this time interval ΔT_m have to be precomputed and stored for further handling.

Vehicles receiving a message m_j as given by the message Specification 1 have to check first if this message is relevant for them. This means, if no message has been received before, within the current time interval the actual message will be stored in a temporary variable m_{act}. In case of the presence of an older message the algorithm first has to find out if the new message is more recent than the already stored one. This is done by comparing the spatial distances of the message originators to the receiver. This means that a vehicle i, receiving a message m_j will consider this message more recent than message a m_{act} if the spacial distance between vehicle i and j is less than that between i and the vehicle originating message m_{act}, both vehicles are driving in the same direction and vehicle j is ahead of vehicle i.

3) Message Forwarding: Each message received by a vehicle with opposite driving direction will be rebroadcast if the message has not been received before or the TTL is equal zero. This is done to increase the sending range by using opposite lane vehicles as relays.

In case of equal driving directions between sender j and receiver i the message will only be rebroadcast if it is considered recent (cf. *Message Reception* phase).


Fig. 1: DRIVE Example

The message forwarding ends if the maximum propagation radius is reached. Since patterns of a wave are recurrent, the maximum message propagation distance has been set to the maximum wavelength of a shock wave, in our case 2000 m [11].

B. Computation of Optimal Velocity

We assume that a driver always tries to drive as fast as allowed by the traffic rules. Furthermore, we assume that the only cause of a slow down is an increasing density the traffic flow ahead. Then, having the information of the next known
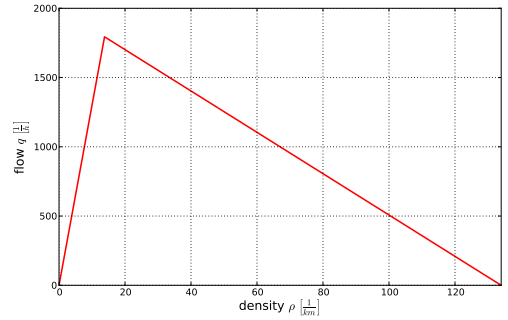

Fig. 2: Triangular fundamental diagram

vehicle in front, we estimate the unknown traffic situation between the receiving vehicle i and the sending vehicle j by applying the LWR model with the solution for the Transport Equation:

$$\tilde{c} = \frac{Q_j - Q_i}{\rho_j - \rho_i} \quad (2)$$

where Q_j and Q_i are the calculated traffic flows for vehicle j and i, respectively. ρ_j and ρ_i denote the vehicular densities corresponding to the actual velocities. Based on this information, it is possible to estimate the traffic conditions between vehicles i and j. Obviously, only situations where the density at the sender position ρ_j is greater than at receivers position ρ_i are of interest. Next, we can compute a velocity recommendation $v_{i,\mathrm{safe}}$:

$$v_{i,\mathrm{safe}} = v_j + \frac{d_{i,j}^* - v_j T}{\tau} \quad (3)$$

$$d_{i,j}^* = d_{i,j} + \tilde{c}\tau \quad (4)$$

$$\tau = \frac{d_{i,j}}{\Delta v_{i,j}} \quad (5)$$

with $d_{i,j}$ being the distance between vehicles j and i, $\Delta v_{i,j}$ being the velocity difference between vehicles i and j and τ being the time for vehicle i to reach vehicle j with the given

(a) Velocity density diagram   (b) Average hop-count   (c) Average number of sent messages
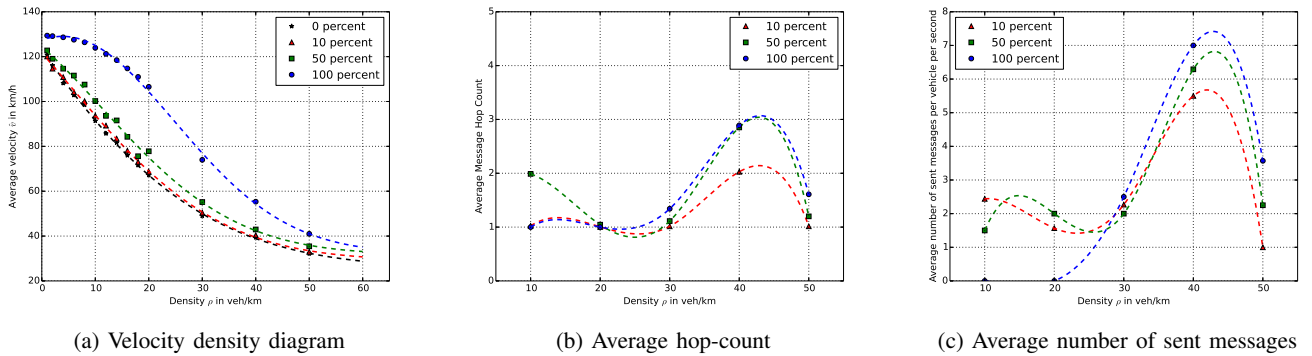
Fig. 3: Simulation results

velocity difference. $v_{i,\text{safe}}$ ensures that vehicle i will adapt its velocity in a way not to hit the tail end of the shock wave produced by vehicle j within the time τ. For more information we refer the reader to [6].

C. DRIVE Example

An example of the DRIVE protocol is given in Figure 1. Suppose, vehicle j experiences a slowdown from $v_j(0) = 35$ m/s to $v_j(1) = 25$ m/s. The onboard DRIVE equipment broadcasts a message $m_{\text{h}}(j) = [id, x_j, y_j, t, v_j, TTL]$ according to the message given in Definition 1. A vehicle i, receiving $m_{\text{h}}(j)$ and considering it as most recent event notification, first checks if $v_j(1) < v_i(1)$. If this holds true, the densities ρ_j and ρ_i at position $\mathbf{x}_j$ and $\mathbf{x}_i$ are estimated. Having an estimate of the densities and the measured values of the velocities of both vehicles we can easily compute the density gradient between them. In the given example and based on the triangular fundamental diagram as given in Figure 2 this leads to an approximated shock wave traveling with $\tilde{c} = -15$ km/h against traffic direction.

Knowing this, from Equations 3, 4 and 5 DRIVE computes a safe velocity for vehicle i given by $v_{i,\text{safe}}(1) = 29$ m/s in order to avoid reaching the tail end of the shock wave produced by vehicle j within the time τ as given by Equation 5, and thus reducing the risk of a traffic perturbation.

III. PROTOCOL EVALUATION

In this section we describe the simulation setup and discuss the obtained results.

Simulation Setup

All simulations presented in this paper have been performed using the Veins framework [12]. Veins offers a connection link between the two well-established simulators Omnet++ [13] and SUMO [14]. With this setup we have a bi-coupled network, meaning that the mobility simulation interacts with the network simulation and vice versa.

The general simulation setup is an octagonal one lane closed loop with an overall length of 10 km. A given number of vehicles are driving around this loop counter clockwise for 2 hours. The mobility scenario is crash-free and specifies that every driver tries to increase his speed up to the maximum

allowed velocity $v_{\text{max}} = 130$ km/h. Previous studies have shown that after exceeding a certain density, the inefficiency of human driving will lead to slow downs [7], [8]. To reproduce this behavior with the *Intelligent Driver Model* (IDM) [15] used for our simulations, it is necessary to introduce a small probabilistic factor. In our scenario this probability is set to $P_{\text{d}} = 1\%$ meaning that a driver will reduce its current speed by 1 m/s for 1 s with the given probability. Several simulation runs have been performed with different vehicular densities and different equipment rates. The aim is on the one hand to show how DRIVE optimizes traffic flow and on the other hand to provide an insight of the network characteristics of the protocol.

The simulated densities ρ are given by the set $\mathcal{R} = \{1, 2, 4, 6, 8, 10, 12, 14, 16, 18, 20, 30, 40, 50, 60\}$ vehicles per km per lane. To provide a detailed overview on the DRIVE protocol, four equipment rates, ranging from uncoordinated (0%) over 10% and 50% to full coverage (100%) have been simulated.

Simulation Results

In this section we discuss the simulation results. Four different analyses with the above mentioned equipment rates of the DRIVE system have been performed. The results show the impact of DRIVE on traffic flow with vehicles that might be equipped with a CACC but additionally takes into account the effects of human behavior.

Figure 3a shows the average velocity over the given set of densities $\mathcal{R}$ with the above mentioned equipment rates. As expected, the average velocities are decreasing with increasing density. However the DRIVE protocol alleviates this effect by redistributing the vehicles in the system in a way that enables them to drive with higher speeds. Between completely uncoordinated traffic and 100% equipment rate with the DRIVE system, the average velocity can be increased by up to 25% in the critical density region between 15 veh/km and 25 veh/km.

Figure 3b illustrates the average message hop-count. This metric indicates how far messages travel backward in the traffic chain or how often they have been forwarded. One can see that with increasing density, the message travel distance increases up to a certain point (40 veh/km). Beyond that density, most vehicles are within transmission range reducing the message

dissemination to one single hop.

These findings also agree with the results given in Figure 3c, where the average number of sent and forwarded messages per vehicle per second is shown. For low densities the ratio of send and forwarded messages is between 1 and 3 messages per vehicle per second. For higher equipment rates it increases to 8 messages per vehicle per second for a density of 40 vehicles per km and then suddenly decreases. In agreement with the trend shown in Figure 3b the communicating vehicles are close enough that no message forwarding is needed to reach all relevant vehicles. This means that the road is already fully congested and that the traffic situation can no more be improved.

IV. CONCLUSION

In this short paper we evaluate the performance of our protocol *Density Redistribution through Intelligent Velocity Estimation* (DRIVE) in a closed border environment. We show that DRIVE allows to increase the average velocity and therewith the vehicular throughput on a congested road. Moreover we show that those improvements can be achieved with only a few messages exchanged between the vehicles.

The redistribution of vehicles in terms of increasing the distances between them has a positive effect on road safety and driver convenience. Furthermore, the avoidance of perturbations by a better distribution of the traffic demand results in smoother traffic, reducing fuel consumption and emissions.

REFERENCES

[1] M. Shinkman, M. Buchanan, and E. I. U. G. Britain), *Driving change: how policymakers are using road charging to tackle congestion.* The Economist Intelligence Unit, 2006.

[2] B. Kerner, *Introduction to Modern Traffic Flow Theory and Control.* Springer Berlin Heidelberg, 2009.

[3] L. Davis, "Effect of adaptive cruise control systems on mixed traffic flow near an on-ramp," *Physica A: Statistical Mechanics and its Applications*, vol. 379, no. 1, pp. 274 – 290, 2007. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0378437106013690

[4] F. Knorr and M. Schreckenberg, "Influence of inter-vehicle communication on peak hour traffic flow," *Physica A Statistical Mechanics and its Applications*, vol. 391, pp. 2225–2231, Mar. 2012.

[5] F. Knorr, D. Baselt, M. Schreckenberg, and M. Mauve, "Reducing traffic jams via vanets," *Vehicular Technology, IEEE Transactions on*, vol. 61, no. 8, pp. 3490–3498, 2012.

[6] M. Forster, R. Frank, M. Gerla, and T. Engel, "A cooperative advanced driver assistance system to mitigate vehicular traffic shock waves," in *Proceedings of the 33rd Annual IEEE International Conference on Computer Communications (INFOCOM'14), April 2014, Toronto, CND*, 2014 (to appear).

[7] M. Lighthill and G. Whitham, "On kinematic waves. II. A theory of traffic flow on long crowded roads," *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, vol. 229, no. 1178, pp. 317–345, 1955.

[8] P. Richards, "Shock Waves on the Highway," *Operations research*, vol. 4, no. 1, pp. 42–51, 1956.

[9] IEEE, *802.11-2012 - IEEE Standard for Information technology - Telecommunications and information exchange between systems Local and metropolitan area networks - Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications.* IEEE Computer Society, 2012.

[10] "Ieee draft standard for wireless access in vehicular environments (wave) - networking services," *IEEE P1609.3/D8, August 2010*, pp. 1–159, 2010.

[11] M. Treiber, A. Kesting, and D. Helbing, "Three-phase traffic theory and two-phase models with a fundamental diagram in the light of empirical stylized facts," *ArXiv e-prints*, Apr. 2010.

[12] C. Sommer, R. German, and F. Dressler, "Bidirectionally Coupled Network and Road Traffic Simulation for Improved IVC Analysis," *IEEE Transactions on Mobile Computing*, vol. 10, no. 1, pp. 3–15, January 2011.

[13] A. Varga, "The omnet++ discrete event simulation system," *Proceedings of the European Simulation Multiconference (ESM'2001)*, June 2001.

[14] M. Behrisch, L. Bieker, J. Erdmann, and D. Krajzewicz, "Sumo - simulation of urban mobility: An overview," in *SIMUL 2011, The Third International Conference on Advances in System Simulation*, Barcelona, Spain, October 2011, pp. 63–68.

[15] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations," *Phys. Rev. E*, vol. 62, pp. 1805–1824, Aug 2000. [Online]. Available: http://link.aps.org/doi/10.1103/PhysRevE.62.1805

Anonymous Credentials and Attribute-Based Authorization Tickets in Car-to-X Communication

Carsten Büttner
Adam Opel AG
Advanced Technology
Rüsselsheim, Germany
carsten.buettner@de.opel.com

Sorin A. Huss
Technische Universität Darmstadt
Integrated Circuits and Systems Lab
Darmstadt, Germany
huss@iss.tu-darmstadt.de

Abstract—In this paper we propose an approach that allows an ITS vehicle station (IVS) to anonymously authenticate itself at applications with attribute-based access restrictions. For authentication purposes we use attribute-based authorization tickets, which are granted after the IVS proved, that it fulfills all access restrictions. For the proof of attributes we rely on anonymous credentials, which are sets of verified attributes. The resulting system supports billing of service usage and revocation of an IVS, tickets, anonymous credentials or single attributes, while protecting the privacy of the IVS.

I. Introduction

In safety relevant communication over ETSI ITS-G5A [1] an ITS Vehicle Station (IVS) authenticates itself by signing all messages with the help of an Authorization Ticket (AT). Besides these safety applications the IVS can also use non-safety applications over ETSI ITS-G5B or cellular communication. To authenticate this kind of communication it is not sufficient to sign all messages with the help of an AT for safety applications, as some applications will have access restrictions. A possible restriction could be for example, that the IVS is of a specific brand or certain sensors or features have to be present. Some Service Providers (SP) might also want to bill the IVS for usage of their application.

In this paper we propose a system that relies on anonymous credentials to issue attribute-based authorization tickets. Each ticket can only be used for specific applications. An IVS can take these tickets to anonymously authenticate itself at applications with access restriction. If the application requires billing, the IVS may pay with anonymous digital money. In addition our system supports the revocation of an IVS, tickets, anonymous credentials, or single attributes. We thus protect the privacy of the IVS to the SP as well as to the other entities.

The rest of this paper is organized as follows: In Section II we discuss anonymous credentials. The entities used in our system are described in Section III. Previous work in authenticating an IVS in Car-to-X communication is reviewed in Section IV. We introduce the new system with attribute-based authorization tickets in V. In Section VI we discuss the different kinds of possible revocation supported in our system. Finally, we conclude in Section VII.

II. Anonymous Credentials

In an anonymous credential system, introduced by Chaum [2], an issuer provides credentials to users, whereas a credential denotes a set of attributes certified by the issuer. An attribute consists of a key (e.g., brand) and a value (e.g., Opel). With this credential the user can prove to a verifier that he or she possesses certain attributes certified in the credential. To convince the verifier, she derives a token that only contains the subset of attributes she wants to prove. So, the verifier learns only the necessary attributes of the user. A user can also prove that the content of a ciphertext produced for a third party is the value of a certain attribute without revealing the attribute to the verifier. Several proofs of the same attributes cannot be linked, because each time another token is being derived. It is possible to prove predicates like *or*, *and*, *greater than*, *smaller than* and *equals* over the attributes.

Possible attributes for an IVS are *brand* or *production date* and features like *traffic sign recognition*, *electric engine* or *Bluetooth*.

As an example, Equation 1 shows how it is possible to check if an IVS is manufactured by GM, without leaking the brand of the IVS. This is realized by checking for all GM brands, if the attribute *brand* equals to the name of the brand. The derived token then indicates, if the equation is true or false and so reveals only the manufacturer and not the brand to the verifier.

$$brand = \text{"}Opel\text{"} \lor brand = \text{"}Chevrolet\text{"} \lor brand = ... \quad (1)$$

Equation 2 can be used to prove that an IVS is not older than x days. From the result a verifier learns, that the IVS is manufactured less than x days ago but not the exact date.

$$production\_date > today - x \quad (2)$$

III. Entities

We use the following entities in our system to obtain attribute-based certificates:

Root CA: The Root Certification Authority (CA) is the trust anchor of the system. It certifies the Enrolment and Authorization Authorities. Every IVS in the system has to trust the Root CA in order to trust the other entities.

IVS Enrolment Authority: The IVS Enrolment Authority (IVS EA) issues the Enrolment Certificate (EC) and credential to the IVS. In order to issue credentials, the IVS EA has a database with all attributes of all ITS Vehicle Stations it is responsible for.
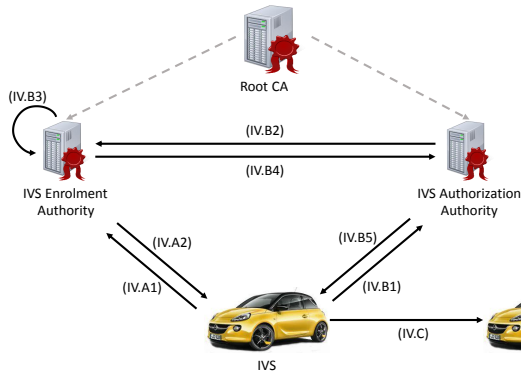
Fig. 1. Request and usage of the Enrolment Certificate and authorization tickets

IVS Authorization Authority: The IVS Authorization Authority (IVS AA) issues authorization tickets to an IVS, after it has checked the eligibility. The AT allows an IVS to use the application specified in the ticket.

SP Enrolment Authority: The Service Provider Enrolment Authority (SP EA) issues Enrolment Certificates to SPs.

Bank: The bank is a central entity which issues digital money. This money can be used to purchase an AT at the IVS AA in order to get access to an application.

Service Provider: Each Service Provider has an EC issued by the Service Provider Enrolment Authority to prove it is an authorized SP in the system. An SP can provide one or several applications to an IVS. Applications can have the condition that the IVS using it must have certain attributes like of a specific brand or presence of a certain sensor. An SP can also have its own IVS Authorization Authority, which is then used instead of the global one.

ITS Vehicle Station: An IVS requests an EC and credential, containing its attributes from the IVS EA. Later on, it uses them to obtain authorization tickets from the IVS AA. The ITS Vehicle Station proves its eligibility for using an application by showing the AT to the Service Provider.

IV. PREVIOUS WORK

The process to request Enrolment Credentials and authentication tickets for safety applications is standardized by ETSI [3]. It gives an IVS the possibility to request ATs from an IVS AA in order to send safety relevant messages over ETSI ITS-G5A. To request ATs, an IVS has to obtain an EC first. Figure 1 shows the single steps which an IVS must complete prior to sending safety relevant messages. In the following these steps are presented in more detail.

A. Enrolment Certificate Request

First, the ITS Vehicle Station has to request an Enrolment Certificate from the IVS Enrolment Authority. This is done in Step IV.A1, where the IVS sends an encrypted and signed message with its canonical certificate together with a certificate request to the IVS EA.

$$Enc_{IVS\_EA}(Sig_{IVS}(canonical\ cert., cert.\ request))$$
$$(IV.A1)$$

Then, the IVS EA validates the request and responses with the encrypted EC (Step IV.A2). After reception, the IVS encrypts, validates, and stores the EC.

$$Enc_{EC}(EC) \qquad (IV.A2)$$

B. Authorization Tickets Request

There are two different ways an IVS can obtain authorization tickets from the IVS AA. One was developed by the Car-to-Car Communication Consortium [4] and the other has been standardized by ETSI [3]. In this paper we only discuss the former one, as our concept can be seen as an extension.

After an IVS received its Enrolment Certificate, the IVS then requests authorization tickets at the IVS Authorization Authority as illustrated in Figure 1. To request ATs, the IVS creates key-pairs to validate, signs the Public Keys (PKs), encrypts the signature and the EC for the IVS EA, and sends the PKs together with both the encrypted signature and the EC to the IVS AA (Step IV.B1).

$$PKs, Enc_{IVS\_EA}(Sig_{EC}(PKs), EC) \qquad (IV.B1)$$

Upon reception the IVS AA calculates the hash of the public keys and sends the hash value together with the encrypted EC and signature, to the IVS EA (Step IV.B2).

$$H(PKs), Enc_{IVS\_EA}(Sig_{EC}(PKs), EC) \qquad (IV.B2)$$

Next, the IVS EA decrypts the EC and the signature and validates them (Step IV.B3) and then sends the result to the IVS AA (Step IV.B4).

$$OK\ or\ FAIL \qquad (IV.B4)$$

If the result is positive, then the IVS AA finally issues the ATs to the IVS (Step IV.B5). Otherwise, an error is returned.

$$ATs\ or\ error \qquad (IV.B5)$$

C. Authorization Tickets Usage

After the ITS Vehicle Station received its authorization tickets, it can start broadcasting signed messages for safety applications over ETSI ITS G5A. This is shown by Step IV.C1 in Figure 1. When the validation period of the ATs is over, the IVS has to request new tickets from the IVS Authorization Authority.

V. AUTHENTICATION

In the following, we detail the steps which are necessary for an ITS Vehicle Station to authenticate itself in order to contact a Service Provider featuring access restrictions.

A. Enrolment Certificate and Credential Request

First, the IVS must obtain an Enrolment Certificate and credential from the IVS Enrolment Authority, which is illustrated in Figure 2. The EC is a certificate to prove, that the IVS is an authorized ITS station. It is the same Enrolment Certificate as in the previous work [3]. The credential is a certification of the attributes of the ITS Vehicle Station. To get an EC and credential, the IVS sends in Step V.A1 an encrypted and signed
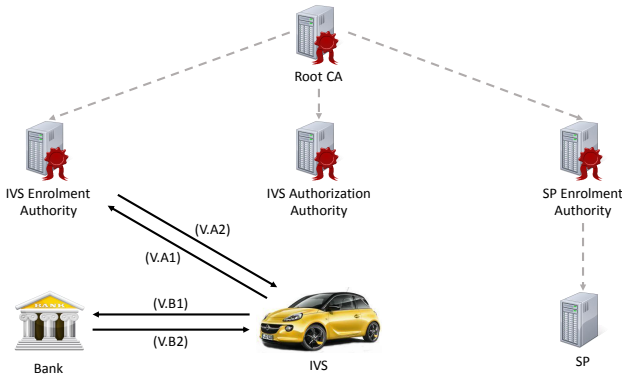
Fig. 2. Request of the Enrolment Certificate, credential and money

Fig. 3. Request of attribute-based authorization tickets

message with its canonical certificate and a certificate request to the IVS EA.

$$Enc_{IVS\_EA}(Sig_{IVS}(canonical\ cert., cert.\ request)) \tag{V.A1}$$

Then, the IVS EA validates the request, issues the EC and credential, and sends them encrypted to the IVS (Step V.A2). After reception, the IVS encrypts, validates, and stores them.

$$Enc_{EC}(EC, Credential) \tag{V.A2}$$

B. Money Request

After the IVS received its EC and credential, it can request digital money from the bank as illustrated in Figure 2. This can be done multiple times. Blind signatures can be used for digital money [5]. If blind signatures are used, the IVS generates random values, blinds them, and sends the blinded values together with its authorization to the bank in order to request digital money (Step V.B1). The bank then blindly signs the values, debits the bank account of the driver, and returns the blind signed values back to the IVS (Step V.B2).

$$blinded\_values, auth\_info \tag{V.B1}$$

$$digital\_money \tag{V.B2}$$

The IVS now has digital money from the bank, which can be used to pay for applications.

C. Attribute-Based Authorization Ticket Request

The ITS Vehicle Station can now request attribute-based authorization tickets to use specific applications, with the help of its Enrolment Certificate, credential and, if necessary, the digital money from the IVS Authorization Authority. When requesting attribute-based ATs, the IVS first checks if the attributes necessary for the application are present as illustrated in Figure 3 (Step V.C1).

Now, the IVS encrypts the hash of its EC for the IVS EA. Then it derives a token containing the necessary attributes and a proof for the ciphertext. Afterwards it creates the key pairs to sign. Now it signs the name of the application and the public keys to certify. This signature is then, together with

the EC, encrypted for the IVS EA. Finally, the IVS sends the application name, the encrypted parts, the token, the PKs, and the necessary money to the IVS AA (Step V.C2).

$$app\_name, Enc_{IVS\_EA}(Sig_{EC}(app\_name, PKs), EC),$$
$$Enc_{IVS\_EA}(hash), token_{attr+hash}, PKs, digital\_money \tag{V.C2}$$

On reception, the IVS AA validates the token and sends the encrypted parts together with a hash over the application name and the public keys to the IVS EA (Step V.C3).

$$H(app\_name, PKs), Enc_{IVS\_EA}(hash),$$
$$Enc_{IVS\_EA}(Sig_{EC}(app\_name, PKs), EC) \tag{V.C3}$$

The IVS EA encrypts everything and checks in Step V.C4 if the EC and the signature are valid and if the hash fits this EC. Afterwards, the IVS EA returns the result of this validation to the IVS AA. (Step V.C5)

$$OK\ or\ FAIL \tag{V.C5}$$

If the validation was successful and the IVS has to pay for the application usage, the IVS AA forwards the digital money from the IVS to the bank (Step V.C6). There the money is transfered to the the account of the IVS Authorization Authority. Subsequently, the bank returns the validity information (Step V.C7).

$$digital\_money \tag{V.C6}$$

$$OK\ or\ FAIL \tag{V.C7}$$

Finally, the IVS AA issues the attribute-based authorization tickets for the application to the IVS (Step V.C8), if all previous steps were executed successfully. Otherwise an error is transmitted.

$$ATs_{app\_name}\ or\ error \tag{V.C8}$$

The outlined system supports two different kinds of billing - per request or per time period. If the IVS has to pay per request, each AT can be used for one request only. When the IVS uses an AT to access an application, the AT is revoked immediately at the SP to prevent further usage. The IVS gets as many ATs as it has payed for. When per time period billing

Fig. 4. Application usage

is used, the IVS gets authorization tickets, which are only valid in the time period it payed for.

When the SP registers its applications at the IVS AA, it has to provide the necessary attributes, billing information, etc. to the IVS AA. It is possible that each SP has its own IVS AA. However, this affects the privacy of the SP, because the SP can now link the ATs to the IVS.

D. Application Usage

When the IVS successfully requested ATs from the IVS AA, it can start using the application with the issued tickets as illustrated in Figure 4.

VI. REVOCATION

There are different kinds of revocation the system should support. It needs to revoke a whole IVS from participating in the system as well as single ATs, the whole credential or a single attribute of an IVS. In the following these cases are discussed.

A. ITS Vehicle Station

To exclude an ITS Vehicle Station, the IVS Enrolment Authority has to revoke the Enrolment Certificate of the IVS. Then the IVS can no longer get new authorization tickets from the IVS AA, since the EC is checked each time ATs are requested. There is no difference in comparison to the revocation of an IVS in previous work.

B. Authorization Tickets

If a Service Provider wants to revoke a single authorization ticket for its application, the SP can revoke it by himself. The SP can locally mark it as invalid, since each application has its own ATs, which are only valid for this application. If the ATs are also used in application specific Car-to-Car communication, the SP can distribute Certificate Revocation Lists (CRLs) to its clients. If an AT was issued by accident, the IVS AA can report this AT to the SP so it can be revoked.

C. Anonymous Credential

The revocation of the credential of an IVS can be done by the IVS EA like the revocation of the EC, since upon each request the IVS EA is asked if the credential is still valid.

D. Attribute

To revoke a single attribute of an IVS, it is necessary to have some kind of misbehavior detection, since it is necessary to detect which attributes are no longer valid. There may be a central entity which collects misbehavior reports, containing the AT and in which manner the IVS misbehaved from the SPs and triggers the revocation process.

Like in the ETSI ITS-G5A communication, the identity of the holder of an AT can be revealed by a cooperation of the IVS EA and IVS AA. Then the IVS EA can mark the current credential of the IVS as invalid. The next time the IVS tries to request a new AT for an application, the IVS EA can detect that the credential is no longer valid and triggers the IVS to request a new credential. The new credential can have less or other attributes, according to the misbehavior of the IVS.

VII. CONCLUSION

In this paper we proposed an approach aimed to allow Service Providers to specify attributes an ITS Vehicle Station must have in order to use its applications. An IVS can prove to the IVS Authorization Authority, which is issuing attribute-based authorization tickets for applications, that it has the necessary attributes to exploit the application, without revealing the attributes not necessary for this purpose. Then the IVS can use its ATs to prove the SP that it is eligible to use them, but without revealing the real identity. We assume a central billing entity, which supports request and time based billing. Since all applications can use this billing service, applications do not need to implement their own billing.

The presented approach supports a revokation of the IVS, single ATs, the credential of an IVS or even single attributes of the credential. Also, the system preserves the privacy of the IVS against the SPs and IVS AA, since they know only that the IVS possesses the necessary attributes and not necessarily the identity of the IVS. The IVS Enrolment Authority is aware of the identity and all attributes of the IVS, but not which application is being exercised. The billing service knows real data about the IVS like its bank account, but not its EC, credential or which application it is using.

REFERENCES

[1] ETSI ES 202 663, *Intelligent Transport Systems (ITS); European profile standard for the physical and medium access control layer of Intelligent Transport Systems operating in the 5 GHz frequency band*, ETSI Final draft ETSI ES 202 663, Rev. V1.1.0, 2009.

[2] D. Chaum, "Security without identification: Transaction systems to make big brother obsolete," *Commun. ACM*, vol. 28, no. 10, pp. 1030–1044, 1985.

[3] ETSI TS 102 941, *Intelligent Transport Systems (ITS); Security; Trust and Privacy Management*, ETSI Technical Specification ETSI TS 102 941, Rev. V1.1.1, 2012.

[4] N. Bißmeyer, H. Stuebing, E. Schoch, S. Götz, J. P. Stotz, and B. Lonc, "A generic Public Key Infrastructure for securing Car-to-X Communication," in *18th ITS World Congress*, 2011.

[5] D. Chaum, "Blind signatures for untraceable payments," in *CRYPTO 82*, 1982, pp. 199–203.

Towards the City-scale Simulation and Performance Assessment of Electric Vehicles

Ekaterina Masalkina, David Eckhoff, Rüdiger Berndt and Reinhard German
Computer Networks and Communication Systems
Dept. of Computer Science, University of Erlangen, Germany
{ekaterina.masalkina,david.eckhoff,ruediger.berndt,reinhard.german}@fau.de

Abstract—Electric vehicles are believed to soon play an important role in urban traffic. In order to fully understand their performance in terms of range and battery management large-scale simulations can offer valuable insights. In this work we present a simulation framework to study these and other aspects of future mobility.

We develop a computationally inexpensive battery and energy consumption model that accounts for rolling and air resistances and only requires preset constant parameters and speed updates of a vehicle to compute the battery's State of Charge. We furthermore show that—based on real map data—it is possible to conduct city-scale simulations with 100 000 vehicles to investigate a wide range of effects considering an increasing proportion of electric vehicles. Finally, the framework also allows to investigate battery management strategies which are based upon inter-vehicle communication.

Index Terms—Electromobility; IVC; lithium-ion battery model

I. Introduction

The absence of direct CO_2 emission makes electric vehicles a promising approach to lower pollution and reduce the CO_2 footprint. Therefore, the German Government released plans to increase the number of electric vehicles on German roads to 1 million by the year 2020 [1]. To reach this goal, the main challenge for electric vehicles has to be overcome: their limited range. One approach is to use rechargeable lithium-ion batteries; given a high energy density and low self-discharge rate, they can boost the effectiveness and thereby the popularity of electric cars. Unfortunately, the performance of lithium-ion batteries deteriorates over time [2], caused by age and the number of charge cycles. One possible solution for this problem is advanced battery management [3].

Fully understanding the effects of different management strategies from Field Operational Test (FOT) is a cost-intensive and time-consuming approach. Large-scale simulations offer a promising solution to this problem, however, they require realistic traffic and battery models in order to produce meaningful results. Especially when battery management includes the utilization of inter-vehicle communication [3], a comprehensive simulation framework can help predict the batteries' behavior under numerous conditions and to find approaches for extending the batteries' life-times.

In this paper, we present our work towards such a framework. Our contribution can be summarized as follows:

- We present a computationally efficient battery and energy consumption model that only requires preset engine parameters and data obtained by a microscopic traffic simulator
- We show how to conduct city-scale traffic simulation using SUMO and the OMNeT++ simulator to examine (potentially aided by wireless communication) battery management strategies.
- We discuss possible directions where inter-vehicular communication can help extend battery life-time.

II. Related Work

Lithium-ion battery modeling has been subject to numerous studies.

The authors of [4] present a behavioral lithium-ion battery model in a general sense without going into kinematic details. They discuss the dependency of the battery discharge capacity on the discharge rate and temperature. In [3] Tielert et al. analyse how ambient temperature, road gradient, auxiliary consumers and driver behavior impact the energy consumption of electric vehicles [3]. They study how vehicular networks, in particular traffic-light-to-vehicle communication, can decrease battery energy consumption and find that electric cars could profit from longer information distances. By doubling the information distance from 300m to 600m up to tripled battery energy savings were observed.

Using the presented battery model, we aim to further investigate the possibilities of wireless communication in city-scale simulations. In this work we perform the simulation of electric vehicles on city roads, implement a battery module and build it in our simulation setup. This enables us to evaluate the traffic flow, vehicular networks and the impact of different factors on the battery State of Charge (SOC) simultaneously.

III. Battery Model

This section describes our computationally inexpensive battery and energy consumption model that only requires preset engine parameters and that can operate solely on vehicles speeds, for example obtained by a microscopic traffic simulator. At this stage the model of lithium-ion battery is temperature independent and does not account for battery recuperation or road gradients. These extensions are subject of future work.

In each time step of the traffic simulation, the SOC is computed as a function of the vehicle speed $\vec{v}$, the acceleration $\vec{a}$ and the angular speed of the wheels $\vec{w_R}$ (see Figure 1). However, given only $\vec{v}$, the acceleration $\vec{a}$ is found as a rate

Figure 1. Battery model: State of Charge (SOC) in relation to the vehicle speed

of change of the vehicle speed compared to the last time step. Furthermore, the angular speed of the wheels $\vec{w_R}$ and consequently, the motor's angular speed $\vec{w_M}$ can be derived. The motor power P_M in relation to the number of revolutions per minute n and torque M is specific to the type of vehicle and was provided by AUDI AG for this research.

The power of the battery P_B and current I_B can be calculated based on the motor efficiency η_M according to [5]:

$$P_B = \frac{P_M + P_R + P_A}{\eta_M},\tag{1}$$

where P_R and P_A denote the rolling and air resistance, respectively.

Assuming that the battery voltage U_B is constant, I_B can be calculated as follows:

$$I_B = \frac{P_B}{U_B}.\tag{2}$$

Finally, using the battery capacity C_N, SOC is approximated by Equation (3):

$$SOC = \frac{\int_{t_0}^{t_1} I_B(t)dt}{C_N}.\tag{3}$$

This formula serves as rough approximation since it assumes that current and temperature correction factors are both equal to one [4].

IV. Simulation Model

In order to assess the performance and the impact on energy consumption of electric vehicles in an urban environment, we extend the Veins framework [6] accordingly. Veins couples the network simulator OMNeT++ [7] and the microscopic traffic simulator SUMO [8]. We chose OMNeT++ to be able to later introduce inter-vehicular communication as a possible battery management strategy.

With the help of MATLAB we generate the input for the Urban Driving Cycles as defined in the New European Driving Cycle (NEDC) [9] for a single vehicle and use this to validate our model implemented in OMNeT++.

A. Implementation

To allow for the independent computation and monitoring of the SOC for each vehicle, we extended them to maintain an instance of our rechargeable lithium-ion battery model (Figure 2a).

Our simulation further allows the deployment of charging stations, which are maintained by a manager module (Figure 2b). This module reads parameters such as position and capacity

from an XML file and creates charging station modules at run-time.

(a) Car battery module is built in the car NED file

(b) Highway NED file includes charging station manager module

Figure 2. OMNeT++ Implementation

In a proof of concept simulation, an electric vehicle will choose the closest free charging station when its SOC is below a certain threshold. The vehicle is then rerouted to the charging station and will recharge until the battery is full. The vehicle will then proceed to its original destination.

B. Simulation of Nuremberg Road Traffic

The meaningfulness of a simulation heavily relies on the correctness of the used models [10]. It is therefore of high importance to simulate realistic traffic in order to draw conclusions on the performance of electric vehicles.

According to the Nuremberg transport statistics for 2011 there are 227,482 registered passenger cars [11]. For simplicity we carry out the simulation of the city of Nuremberg (see Figure 3) with $100\,000$ vehicles on a PC with quad-core CPU, whose maximal speed and cache are 3.5 GHz and 10 MB, respectively.


Figure 3. Map of the city of Nuremberg in SUMO

The simulation time step is set to 1 s and vehicle parameters possess default values. The resulting simulation runtime is approximately 22 minutes, whereas in real life the duration exceeds 27 hours.

To also investigate the performance of electric vehicles, we started by simulating one electric vehicle in a smaller section of the map (see Figure 4). The distance between the start and destination points of the route is about 3.5 km. We placed two charging stations in close proximity to the chosen route.

One possible outcome of such simulations is an approximation of the optimal positioning of charging stations. In a first step we assume that regular gas stations are also equipped with plug charging connectors for electric cars. This allows us to draw conclusions whether (and how many) additional charging stations are required considering a desired maximum waiting time.


Figure 4. Individually simulated electric vehicle on the Nuremberg map

V. SIMULATION RESULTS

Figure 5 shows an example SOC curve of an individually simulated electric vehicle on the Nuremberg map. It can be seen that after 377 s with an average speed of 10.8 m/s and a top speed of 14 m/s the SOC is reduced by less than 7 % (corresponding to 0.84 kW h).

To prove the correctness of our OMNeT++ model we used the SUMO mobility trace as an input to the MATLAB model. The comparison of the SOC behavior in Veins and MATLAB simulation models shows no difference.


Figure 5. Speed and State of Charge of one individually simulated electric vehicle

In a second simulation, we use our battery model to simulate a typical urban driving cycle as defined by the NEDC [9]. Figure 6 can be seen as a benchmark for the comparison against other battery models and to investigate the effect of future features such as recuperation and temperature dependency. It


Figure 6. Speed, energy consumption and state of charge of a vehicle's run through the NEDC

also allows us to validate our model based on data from real experiments. We observe that after 780 s and a covered distance of 4067 m the battery is drained by less than 7 %.

VI. INTER-VEHICULAR COMMUNICATION

The communication among vehicles and between vehicles and infrastructure can be a promising tool to further reduce battery drain and thereby increase the range of electric vehicles. For example, Green Light Optimized Speed Advisory systems [4], [12] and vehicular communication based on IEEE 802.11p [13] and/or UMTS/LTE cannot only be used to avoid unnecessary stops but also to choose an optimal recuperation stage to obtain the best possible energy gain.

But also other vehicles can provide valuable information: for example, a car driving ahead can inform the electric vehicle about its deceleration rate or planned lane changes to avoid an energetic suboptimal driving maneuver and again aid in choosing the correct recuperation stage.

Vehicular networks can also be helpful in reducing the waiting times at charging stations for electric cars. Receiving information about the location and current occupancy of charging stations could enable drivers of electric vehicles to reserve a spot at the most suitable charging station in advance.

VII. CONCLUSION AND FUTURE WORK

In this work we introduce a computationally inexpensive lithium-ion battery and energy consumption model which considers rolling and air resistance and is only based on predefined engine parameters and speed of a vehicle. We presented a simulation model, which is used to simulate electric vehicles in the city of Nuremberg and investigate their performance together with their impact on energy consumption. The simulation results show that the developed model can be used to further study the performance of electric vehicles under certain conditions.

Future work will concentrate on extending model to also account for recuperation as well as the influence of ambient temperature. Furthermore we will focus on the simulation of realistic city-wide traffic based on empiric data to evaluate the impact of the substitution of common combustion engine

vehicles with their electric counterparts. Based on these steps, we will then investigate possible benefits of inter-vehicular communication on battery management.

REFERENCES

[1] The German Federal Government, "German Federal Government's National Electromobility Development Plan," press release, August 2009.

[2] J. Vetter, P. Novák, M. R. Wagner, C. Veit, K.-C. Möller, J. O. Besenhard, M. Winter, M. Wohlfahrt-Mehrens, C. Vogler, and A. Hammouche, "Ageing mechanisms in lithium-ion batteries," *Journal of Power Sources*, vol. 147, pp. 269–281, September 2005.

[3] T. Tielert, D. Rieger, H. Hartenstein, R. Luz, and H. Hausberger, "Can V2X Communication Help Electric Vehicles Save Energy?" in *12th International Conference on ITS Telecommunications (ITST 2012)*. Taipei, Taiwan: IEEE, November 2012, pp. 232–237.

[4] L. Gao, S. Liu, and R. A. Dougal, "Dynamic Lithium-ion Battery Model for System Simulation," *IEEE Transactions on Components and Packaging Technologies*, vol. 25, no. 3, pp. 495–505, September 2002.

[5] D. G. Fink and J. M. Carroll, *Handbook for Electrical Engineers*. McGraw-Hill Inc, October 1978.

[6] C. Sommer, R. German, and F. Dressler, "Bidirectionally Coupled Network and Road Traffic Simulation for Improved IVC Analysis," *IEEE Transactions on Mobile Computing*, vol. 10, no. 1, pp. 3–15, January 2011.

[7] A. Varga and R. Hornig, "An overview of the OMNeT++ simulation environment," in *1st ACM/ICST International Conference on Simulation Tools and Techniques for Communications, Networks and Systems (SIMUTools 2008)*. Marseille, France: ACM, March 2008.

[8] D. Krajzewicz, G. Hertkorn, C. Rössel, and P. Wagner, "SUMO (Simulation of Urban MObility); An Open-source Traffic Simulation," in *4th Middle East Symposium on Simulation and Modelling (MESM 2002)*, Sharjah, United Arab Emirates, September 2002, pp. 183–187.

[9] "Council Directive 91/441/EEC of 26 June 1991 amending Directive 70/220/EEC on the approximation of the laws of the Member States relating to measures to be taken against air pollution by emissions from motor vehicles," The Council of the European Communities, Directive 91/441/EEC, June 1991.

[10] F. Dressler, C. Sommer, D. Eckhoff, and O. K. Tonguz, "Towards Realistic Simulation of Inter-Vehicle Communication: Models, Techniques and Pitfalls," *IEEE Vehicular Technology Magazine*, vol. 6, no. 3, pp. 43–51, September 2011.

[11] Stadt Nürnberg, "Auszug der Bezirkstabellen und Karten aus den Innergebietlichen Strukturdaten Nürnberg 2012," Amt für Stadtforschung und Statistik, Innergebietliche Strukturdaten Nürnberg 2012, December 2012.

[12] D. Eckhoff, B. Halmos, and R. German, "Potentials and Limitations of Green Light Optimal Speed Advisory Systems," in *5th IEEE Vehicular Networking Conference (VNC 2013)*. Boston, MA: IEEE, December 2013, pp. 103–110.

[13] "Wireless Access in Vehicular Environments," IEEE, Std 802.11p-2010, July 2010.

In-Car Communication based on Power Line Networks

Thomas Gehrsitz
Institute for Communication Networks
Technische Universität München
München, Germany
thomas.gehrsitz@tum.de

Wolfgang Kellerer
Institute for Communication Networks
Technische Universität München
München, Germany
wolfgang.kellerer@tum.de

Helmut Kellermann
BMW Research and Technology
München, Germany
helmut.kellermann@bmw.de

Abstract—**The number of electronic control units in today's cars is permanently increasing. As the demand for information exchange between these units is growing, the complexity of the in-car communication infrastructure also increases. In the past more and more bus-segments have been added to cope with the growing demands. In order to reduce the complexity and the costs for the future in-car communication infrastructure, the approach of using power line communication for the in-car communication has been analyzed. By using power line communication, the cabling can be reduced to the minimum. In this paper the challenges of today's and the requirements for future in-car systems are summarized. The HomePlug Green PHY standard, which has been designed for Smart Grid applications, is analyzed according to its applicability to the in-car communication. Additionally a proposal and its evaluation for the reduction of the protocol overhead is presented.**

I. INTRODUCTION

In the past decades, the demand for an information exchange between ECUs (electronic control units) in vehicles rapidly raised. Prior to the digital data transmission, mainly discrete electrical signals have been used for the information exchange. With the increasing number of safety and comfort functions, the demand for a data exchange between a huge number of ECUs is growing from vehicle generation to vehicle generation. Over the years different communication systems emerged for different applications.

In the 1980s the CAN-bus (Controller Area Network) [1] has been designed and released. The CAN-bus was one of the first widely spread communication buses in modern vehicles. With an increase in body and comfort functions, the wish for a cheap and simple communication system came up. For this purpose the LIN-bus (Local Interconnect Network) [2] has been defined. It is mainly used for the body and comfort functions where safety requirements are typically low and no high data rates are needed. In the beginning of the 21 century the requirement for a new in-car communication system came up. FlexRay [3] has been designed to provide a communication system for time-critical functions like safety functions.

In a today's upper class vehicle more than 40-50 of these communication buses can be found. This introduces challenges in planning, deployment and the maintenance of the harness. Additionally the high number of wires lead to higher costs and also weight, which directly translates into fuel consumption.

The highest rate of change is typically observed in the body and comfort area, where the LIN-bus is dominant. Upper class vehicles already contain more than 20 of these buses, which are distributed over the whole vehicle. If information has to be transmitted to a different LIN-bus, gateways have to be used.

One possibility to reduce the cabling of the ECUs is the usage of wireless communication. But with wireless communication additional challenges arise. The interference to and from other devices in the surrounding of the vehicle have to be considered. Measurements [4] show that the effective area of influence is quite large. The positioning of the antenna has to be chosen carefully and the transmit power should be as low as possible.

Another possibility is the application of Power Line Communication (PLC). With PLC the need for dedicated communication lines besides the power lines disappears. Thereby the cabling is reduced to the minimum - the power supply. Measurements [5] show that the physical transmission of PLC signals on a car body is feasible. But the EMC still has to be investigated in detail in order determine if it meets automotive requirements.

Various PLC standards for the in-house application like G.hn or HomePlug AV are available. In this paper, the latter one is investigated according to its applicability to the in-car communication. Starting from a today's vehicle, the challenges and requirements for the future in-car communication are stated.

In section II a short overview on a today's in-car communication infrastructure is given. In Section III requirements for a future in-car PLC protocol are stated. Section IV summarizes the main features of the HomePlug Green PHY standard and its applicability to the in-car communication is discussed. A proposal for the reduction of the overhead is given. In section V simulation results with the modified protocol are presented. A conclusion is given in section VI.

II. TODAY'S IN-CAR COMMUNICATION

Today's cars include highly interconnected ECUs covering a wide range of applications. Upper class vehicles - depending on the configuration of the features - may contain up to 50-100 ECUs. Besides these ECUs there are also sensors and actuators which have to be connected. The cabling needed for the interconnection of the ECUs easily reaches summed up length of more than 1-2 km. Widely spread bus-systems which can be found in modern vehicles are CAN, LIN and FlexRay.

Multiple buses of each type are typically deployed in a car and most of them are interconnected by a gateway.

As a high rate of change from vehicle generation to vehicle generation is in the body and comfort area where the LIN-bus is dominant, the focus of this paper is on the replacement of the LIN-infrastructure by a PLC-system. The low safety-requirements on the LIN-bus are another reason for choosing its replacement as a first step.

The LIN-bus is a comparatively cheap master-slave bus-system mostly used for non-safety-critical communication in the body and comfort domain. The maximum data rate is 20 kbit/s. However the possible savings of replacing the LIN-infrastructure by a PLC-system is quite high, because of the high number of LIN-buses in a car covering almost every part of the vehicle. A today's upper class vehicle already includes over 20 LIN-buses, which comes along with a high effort in wiring.

The high number of LIN-buses leads to various challenges. The increasing complexity enlarges the effort in planning, installation, management and the diagnosis of the harness. The weight is also increasing which directly translates into higher fuel consumption.

III. PLC FOR THE IN-CAR COMMUNICATION

The requirements for the in-car communication differ from typical Home Area Networks (HANs). Typical HANs consist of a small number of nodes interconnected by a network with comparably high data rate. For a high throughput, the frames typically carry some hundred bytes in order to reduce the overhead. For the in-car communication in contrast, a high number of nodes has to be connected while the exchanged information is typically in the range of only 1-8 byte. Thus a protocol for the in-car communication should be able to support a high number of nodes while minimizing the overhead for short payloads.

Most protocols for the home and wide area networks are non-deterministic. Regarding diagnosis, determinism is a desirable property for in-car communication systems.

Another point is the Quality of Service (QoS). Most protocols support a kind of QoS (i.e. prioritized queues), but for the in-car communication some functions have strict requirements regarding the latency. Functions like control loops often produce periodic traffic either continuously or at least for a period of time.

Thus a protocol should be suitable to transport periodic traffic on the one hand and support event based traffic on the other hand. In addition the costs have to be minimized. Costs can typically be reduced by using off-the-shelf components. For this reason the HomePlug Green PHY standard is analyzed in section IV.

The requirements for an in-car communication system can roughly be summarized as follows:

- Scalability: The protocol has to be able to handle a high number of nodes (one hundred or even more)

- The overhead has to be minimized; in case of the LIN-bus the payload is less or equal to 8 byte

- Support for periodic traffic and additionally support for event based traffic

- Determinism is a desired feature (at least for periodic traffic)

- Minimize costs: prefer the usage of off-the-shelf components

IV. IEEE 1901 AND HOMEPLUG GREEN PHY

The IEEE 1901 standard [6] has been defined to provide coexistence between different PLC standards. Thus IEEE 1901 includes the HomePlug standard. In the following, a brief overview on the HomePlug standard is given.

The HomePlug 1.0 standard has been released in 2001 by the HomePlug Alliance. HomePlug AV as an extension has been released in 2005. HomePlug AV currently supports data rates up to 500 Mbit/s. The new HomePlug AV2 standard already supports data rates of approximately 1 Gbit/s.

For Smart Grid applications the HomePlug Green PHY (HP GP) standard has also been released in 2005. HP GP is a subset of the HomePlug AV standard with the focus on a robust communication. The robustness is reached by using only QPSK (Quadrature Phase-Shift Keying) modulation and using the so-called ROBO (robust OFDM) modes. With the ROBO modes, multiple copies of the same signal are simultaneously transmitted over the OFDM (Orthogonal Frequency Division Multiplex) carriers.

A. MAC layer

For the organization and synchronization of a network, one node is selected to act as a Central Coordinator (CCo). The CCo periodically sends beacons to synchronize the network. In between the beacons, two medium access methods are defined. The first one is the mandatory CSMA/CA (Carrier Sense Multiple Access with Collision Avoidance) which uses a backoff-mechanism for collision avoidance. The optional TDMA is the second access method. The structure of the beacon period is depicted in figure 1.

	beacon period			
Beacon	CSMA	TDMA (optional)	Beacon	· · · ·

Fig. 1. HomPlug beacon period

For the CSMA/CA access, four Channel Access Priorities (CA0 to CA3) are defined. Before a stations starts a backoff, the priority resolution is performed. Two Priority Resolution Slots (PRS) are used to communicate the highest priority of a pending frame to all nodes in the network. Depending on the priority, a station either transmits a priority resolution signal or it listens on the channel. The states of the two Priority Resolution Slots (PRS0 and PRS1) for the four Channel Access Priorities are listed in table I.

After the priority resolution, every station in the network knows about the currently highest frame priority. Only stations with a frame of the highest priority are allowed to start a backoff afterwards.

The medium access is depicted in figure 2. After the priority resolution, the time-slotted random backoff procedure

TABLE I.   HOMEPLUG CHANNEL ACCESS PRIORITIES

| Channel Access Priority	PRS0 State	PRS1 State
CA3	1	1
CA2	1	0
CA1	0	1
CA0	0	0

Fig. 2.   CSMA/CA Medium Access

is started. The backoff procedure uses different counters and parameters: the Backoff Procedure Event Counter (BPC), the Backoff Counter (BC), the Deferral Counter (DC) and the Contention Window (CW). The BPC is set to 0 for the first transmission attempt. The CW and DC are set according to table II.

TABLE II.   HOMEPLUG BACKOFF PARAMETERS

BPC	CA0 & CA1		CA2 & CA3	
0	CW=7	DC=0	CW=7	DC=0
1	CW=15	DC=1	CW=15	DC=1
2	CW=31	DC=3	CW=15	DC=3
>2	CW=63	DC=15	CW=31	DC=15

After CW and DC have been set, the BPC is increased by 1. The Backoff Counter is chosen uniformly in the range of [0, CW] and is decremented by 1 at every idle time-slot. When the BC is 0, the frame is sent out and the successful reception is acknowledged after a Response Interframe Space (RIFS). After a successful transmission the BPC is reset. If a transmission fails, new values for CW and DC are chosen according to table II. Afterwards the BPC is incremented by 1 and a new value for the BC is uniformly chosen.

Until now the procedure follows the backoff procedure known from IEEE 802.11. The only difference is the Deferal Counter. If a station senses the medium busy during its backoff, it freezes its counters. When the channel is idle for the duration of a Contention Interframe Space (CIFS), a priority resolution is initiated again. If the station is allowed to run a backoff afterwards (i.e. it has got a frame of the highest priority), it may resume its backoff depending on the value of the DC. If DC is not 0, DC and BC are decremented by 1 and the backoff resumes. If DC is 0, CW and DC are again set according to table II, BPC is incremented by 1, a new BC is uniformly chosen and the new backoff is started. The DC is used to extend the backoff when the contention is high, thus reducing the collision probability.

In the TDMA access period, stations can ask the CCo for a reserved transmission time. The reservation request is sent in the CSMA region. This may introduce an additional delay, because stations first have to wait for the CSMA region to be the active one. The CCo uses the beacon to communicate the reserved transmission time to all nodes. In the CSMA region collisions might occur. Sporadic in-car messages typically don't need a reserved TDMA slot. In this case the frame

could already be transmitted in the CSMA region. For periodic frames in contrast, a static reservation would be the right choice in order to reduce the delay and overhead.

As a dynamic reservation of transmission time according to the HomePlug AV standard is not a suitable solution for the in-car communication, only CSMA is considered in this paper.

In table III the HomePlug Green PHY ROBO modes are listed. In the last column the Physical Block (PB) size is shown. PB136 denotes a Physical Block with 136 byte in length, PB520 corresponds to a PB with 520 byte. These PBs are the smallest amount of data which is transmitted on the channel. Each PB consists of a PB header and a PB check sequence (each 4 byte in length). For a PB136 this leads to a PB payload (PB body) of 128 byte. In the PB payload the MAC Protocol Data Unit (MPDU) is transported. A MAC frame consists of a 2 byte header and a 4 byte checksum. In the case of a PB136 this leads to a MAC payload of 122 byte per PHY block.

TABLE III.   HOMEPLUG GREEN PHY ROBO MODES

mode	PHY rate	# copies	PHY block
Mini-ROBO	3.8 Mbit/s	5	PB136
Standard-ROBO	4.9 Mbit/s	4	PB520
High Speed ROBO	9.8 Mbit/s	2	PB520

Considering typical in-car frame length of up to 8 byte, even the usage of the smallest PB size of 136 byte introduces a huge overhead. In addition there is an overhead for the CSMA/CA medium access (backoff, interframe spaces and acknowledgment).

B. Protocol overhead

The Physical Protocol Data Unit (PPDU) is depicted in figure 3. A preamble is sent out, followed by a Frame Control (FC), which is either one or two OFDM symbols in length. After the FC, several OFDM data symbols are sent out which carry one or more PBs. The FC has got different tasks. The Delimiter Type field of the FC identifies the type of frame or the content of the data symbols (e.g. beacon, data or acknowledgment).

Fig. 3.   HomePlug Green PHY PPDU structure

One option to reduce the overhead introduced by the size of the Physical Blocks could be the definition of shorter PBs. The length of one block could be chosen to fit into one OFDM symbol, thus reducing the transmission time. A drawback of this option is the fact, that it breaks the current standard.

In protocol design when defining headers, a common way is to reserve fields or IDs for future extensions. In other cases, fields designed for a specific function are later on used for a different purpose. One example is the Explicit Congestion Notification (ECT), which is an extension to the IP- and TCP-protocol. ECT is used for the end-to-end notification of network congestion and it uses 2 bit of the Type Of Service (TOS) field in the corresponding header.

Another option to reduce the overhead is described in the following. The IEEE 1901 standard defines two MPDU frame formats. The long MPDU consists of the Frame Control followed by the MPDU payload, which is transported in the OFDM data symbols as illustrated in Fig. 3. The short MPDU in contrast only consists of the Frame Control. One example for a short MPDU is the acknowledgement indicated by the corresponding value in the Delimiter Type field of the FC.

When having a look at the Frame Control fields, two Delimiter Types can be found as currently unused. In addition there is a Variant Field whose type of content is indicated by the Delimiter Type. One of the currently unused Delimiter Types can be used to identify the transmission of a payload in the Variant Field. The Variant Field in the FC is 96 bit in length. This is enough to transport a LIN payload of 8 byte plus additional management information like an address.

The transmission time for one frame with a PB136 is 353.08 µs (including the transmission of the preamble, Frame Control and data symbols for the transmission of the PB136) when all 1155 OFDM carriers are available and Mini-ROBO mode is used. Note that the shortest PHY block size (PB136) is only used with the Mini-ROBO mode. When transmitting the payload in the Frame Control, the transmission time can be shortened to 110.48 µs (transmission time for the preamble plus Frame Control). But this is only the transmission time of the frame. Additional protocol overhead such as the priority resolution, backoff, interframe spaces and the transmission of the acknowledgement has to be considered.

When using the default values for the RIFS (140 µs), the CIFS (100 µs) and the slot duration (35.84 µs), the additional overhead in a collision-free environment can be determined. As the Contention Window in the first backoff-stage is 7, the average backoff length is 3.5 slots. The acknowledgement is sent in a short MPDU, i.e. only a FC is transmitted. With these values the duration of the additional average overhead without collisions can be calculated:

$$T_{overhead} = 2 \cdot T_{PRS} + 3.5 \cdot T_{slot} + T_{RIFS}$$
$$+ T_{ACK} + T_{CIFS} = 547.6 \, \mu s \qquad (1)$$

The total transmission time for a frame including protocol overhead and frame transmission sums up to 658.08 µs. This means more than 80% of the total transmission time is protocol overhead.

Regarding the requirements for the in-car communication, the transmission with the IEEE 1901 CSMA/CA protocol might be an option. The overhead can be reduced by using the FC for the payload transmission. The costs can be reduced when using standard HomePlug Green PHY hardware. But the protocol is non-deterministic and collisions might occur. In order to investigate this effect, simulations with the modified protocol are presented in the next section.

V. SIMULATIONS WITH THE MODIFIED PROTOCOL

In order to analyze the performance of the protocol, simulations are carried out. A framework for the OMNeT++ network simulator has been created which supports the IEEE 1901 standard. The simulation is modified for the payload transmission in the Frame Control. The scenario of the simulation is a network with 100 nodes. All nodes are transmitting frames with a configurable inter-arrival time with an exponential distribution.

The beacon period is set to 50 ms. On an AC line the beacon period is two line cycles in length (33.3 ms @60 Hz, 40 ms @50 Hz). On a DC line the beacon period has to be defined in the CCo. The Frame Control is used for the payload transmission, e.g. no OFDM data symbols are transmitted. The nodes are equally configured to generate frames with a given frame rate. For the arrival process an exponential distribution is used.

Figure 4 shows the simulation results when only frames of priority CA1 are sent. The results for priority CA0 are the same, because CA0 and CA1 are using identical backoff parameters. On the x-axis the number of total generated frames of all 100 nodes in the network is given. The green solid line corresponds to successfully transmitted frames. The dashed red line corresponds to collided frames. The number of collided frames/s can exceed the number of generated frames/s because collided frames are transmitted (and may collide) again. Vertical lines have been added to mark the traffic-values where 3% and 5% of the transmissions collide. At a frame generation rate of approximately 1075 frames/s the collision probability exceeds a threshold. From this point on the channel usage rapidly increases due to re-transmitted frames. This in return increases the collision probability, which is a self-amplifying effect.

The presented figures are the mean value of 500 simulation runs. The variation from run to run in the low load case is minimal, beacuse the collision probability is low. In the area where the collisions shoot up, huge variations can be observed because of the random backoff procedure.

Fig. 4.   Simulation results with priority CA1 only

Saturation is reached at a frame generation rate of approximately 1200 frames/s. From this point on stations always have frames in their TX queue. For the simulations the retry counter has been set to 7, e.g. after 7 unsuccessful transmission attempts a frame is dropped. When the offered load is increased further, the TX queue is filled more and more, but this has no observable effect on the channel usage. The system should be operated below a frame generation rate of 1000 frames/s, because otherwise the collision rate is way too high.

If a higher priority is used (CA2 or CA3) the average backoff length is shortened. This in turn increases the collision probability. One option to reduce the collision probability is the usage of all four priorities. The priority resolution ensures that only stations with a frame of the highest priority in the network start a backoff. For low loads this reduces the average number of stations simultaneously running a backoff, thus reducing the collision probability.

In figure 5 the simulation results are shown, when all four priorities are used. For every generated frame, a random priority is chosen with an uniform distribution, i.e. the probability of each priority is 25%.


Fig. 5. Simulation results with uniformly distributed priorities

The overall number of collisions is lower and the throughput is higher when compared to figure 4. Using the priorities penalizes low priority traffic in favor of higher prioritized frames. This might be a desired behavior, but priorities have to be chosen with care.

VI. CONCLUSION

An overview on current in-car communication systems has been given. The requirements for a PLC protocol for the in-car communication have been summarized. These requirements commonly differ from the ones for example in Home Area Networks. A PLC standard designed for Smart Grid applications - the HomePlug Green PHY standard - has been analyzed regarding to its applicability as a replacement for the LIN-infrastructure. Even though the HomePlug Green PHY standard has been designed for Smart Grid applications, the minimum MAC payload is 122 byte when using the smallest Physical Block size (PB136). Currently the payload length of a LIN or a CAN frame is limited to 8 byte, which is far below the minimum payload of a Green PHY frame. In order to reduce the overhead, the transmission of the payload in the Variant Field of the Frame Control has been suggested. Thereby the transmission time of one frame can be reduced from 353.08 μs to 110.48 μs. Calculations show that the additional protocol overhead (backoff, interframe spaces and acknowledgments) is still huge.

Simulations have been carried out in order to analyze the protocol performance with a high number of nodes. The collision probability increases with increasing traffic load. The four Channel Access Priorities of the HomePlug standard can be used to lower the collision probability. But they have to be chosen carefully, because low priorities might starve when the load is too high. In addition the energy consumption has to be investigated. When using CSMA/CA, nodes have to listen on the channel most of the time. Thus using a pure CSMA/CA medium access might not meet the requirements for some applications. If this is the case, the combination of CSMA/CA and TDMA might be a solution, as defined in the HomePlug AV standard. A modification to the TDMA mechanism might be necessary to fit automotive requirements.

REFERENCES

[1] "CAN specification version 2.0," *Robert Bosch GmbH*, 1991.

[2] "LIN specification package revision 2.1," *LIN Consortium*, 2006.

[3] "Flexray protocol specification v3.0.1," *Flexray Consortium*, 2010.

[4] M. Blesinger, E. Biebl, T. Gehrsitz, J. Eberspächer, P. Fertl, O. Klemp, and H. Kellermann, "Angle-dependent path loss measurements impacted by car body attenuation in 2.45 ghz ism band," in *Vehicular Technology Conference (VTC Spring), 2012 IEEE 75th*, 2012, pp. 1–5.

[5] M. Lienard, M. Carrion, V. Degardin, and P. Degauque, "Modeling and analysis of in-vehicle power line communication channels," *Vehicular Technology, IEEE Transactions on*, vol. 57, no. 2, pp. 670–679, 2008.

[6] "IEEE standard for broadband over power line networks: Medium access control and physical layer specifications," *IEEE Std. 1901-2000*, 2000.

An evaluation framework for pre-distribution strategies of certificates in VANETs

Michael Feiri
Services, Cybersecurity and Safety
University of Twente
The Netherlands
Email: m.feiri@utwente.nl

Jonathan Petit
Services, Cybersecurity and Safety
University of Twente
The Netherlands
Email: j.petit@utwente.nl

Frank Kargl
Institute of Distributed Systems
University of Ulm
Ulm, Germany
Email: frank.kargl@uni-ulm.de

Abstract—Security and privacy in vehicular communication are expected to be ensured by the pervasive use of pseudonymous certificates and signed messages. The design and establishment of necessary public key infrastructure and hierarchies of certificate authorities is ongoing in industry consortia, such as the Car-to-Car Communication Consortium. The privacy preserving dissemination of pseudonymous certificates is however still expected to be limited to single-hop exchanges between vehicles. This limitation to one-hop strategies might not be ideal, especially considering the importance of ensuring trustworthy stateless information exchange upon reception of the very first communication packets. We propose to investigate multi-hop pre-distribution strategies for certificates to significantly reduce this first encounter problem.

I. INTRODUCTION AND RELATED WORK

A core requirement for effective vehicular communication is secure information exchange between vehicles. As the volatility of vehicular networks makes stable secure channels impractical, the commonly accepted solution for inter vehicle communication (V2V) is to use authenticated messages, based on a common vehicular public key infrastructure (PKI) and accredited certificate authorities (CA). Relevant standardization efforts at ETSI [1] and IEEE [2] have proposed appropriate protocols to implement such architectures. In order to additionally assure privacy for the passengers of vehicles in such architectures it is foreseen to issue multiple pseudonymous identities to vehicles, which can be switched according to given rulesets. The goal of providing multiple pseudonymous identities to vehicles is to provide location privacy to passengers.

The dissemination of certificates to vehicles is indirectly specified in ETSI and IEEE by means of direct inclusion of certificates in signed messages or through the possibility to explicit requests the inclusion of certificates. There are no universal rules that define the inclusion or omission of certificates, but the fact that certificates represent a significant amount of data makes it attractive to minimize the amount of certificate inclusions while maximizing the service quality. Service quality in this context is the ability to verify incoming messages of previously unknown vehicles as fast as possible. Waiting periods for exchanges of certificates after the first encounter of a vehicle in communication range need to be minimized. Existing research and suggestions in relevant standardization effort limit themselves to one-hope dissemination.

PKI systems, such as S/MIME [3], rely on similarly bootstrapping secure communication through an initial exchange of certificates between communication partners over not-yet-secured channels or messages. In case a communication partner is not available for the initial exchange of certificates it is possible to use a third party cache of certificates. This is possible because the trust in certificates does not depend on secure delivery but instead is solely hinged on (a chain or set of) trust anchors that certify the authenticity of the enclosed public keys. The concept of public cache server is popular for the related PGP/MIME [4] system, where keyservers are commonly relied on to deliver public key material along with cross certifications added by third party entities for the formation of a web-of-trust.

Communication system operating on public internet infrastructure can be expected have relatively stable connections and to tolerate varying amounts of latency to complete such an initial exchange of certificates to bootstrap secure communication. No matter if it is done end-to-end between endpoints or using caches located in keyservers. Vehicular communication on the other hand operates under more constrained conditions and under stricter requirements. Hidden station effect can prevent two way communication, the reachability of caches in RSUs or backend infrastructure (V2I) can not be assumed at all times, and the availability of certificate material must be guaranteed within reasonable latencies with respect to the beacon frequency and the relative trajectories of the moving vehicles.

Within the specific constraints of vehicular communication patterns, it is still possible to envision protocols that use caches for more effective dissemination of certificates. Also, another way to investigate dissemination of certificates beyond 1-hop neighborhood is to adapt protocols that propose to disseminate neighbor information through 2-hop piggybacking [5]. Such protocols have been specifically proposed for purposes of distributed congestion control (DCC) in vehicular communication networks. Both approaches will be investigated using analytical models to predict the suitability of these approaches to enhance the efficiency and effectiveness of certificate distribution in VANETs.

II. REQUIREMENTS

Pseudonymous certificates are the foundation for both security and privacy in vehicular communication networks. Vehicles are expected to locally cache at least a small set of certificates to have the capability to immediately verify

the messages of all nearby vehicles even if the (chain of) relevant certificates are not included in every message. This can occur when a signed message only includes the digest of a certificate instead of the full (chain of) certificates in order to save bandwidth.

For the purpose of our analysis we assume that the size of a single certificate including compressed NIST P-256 ECDSA keys is 140 bytes [6]. One GiB of storage space could store approximately 7.669.584 certificates. Considering that in Europe alone an estimated number of more than 250 million vehicles are in use it becomes immediately obvious that pre-distributing all certificates is not a realistic option. A global estimate of motor vehicle registrations indicated about 1 billion vehicles in 2010. Assuming each vehicle would be equipped with a set of 100.000 pseudonyms to cover every possible 5 minute period over the course of a full year yields and upper estimate of 100 trillion pseudonymous identities. Storing this amount of pseudonyms would require about 12 EiB of storage. We assume as a requirement for our analysis that a vehicles on-board unit (OBU) are limited to a maximum of 1GiB of storage space for caches of certificate material.

Another requirement exists relative to a baseline of bandwidth consumption for the simple inclusion of certificates in every message sent by vehicles. For the purpose of our study we only investigate periodic beacon messages. Combined with a synthetic channel model and estimates of average vehicles densities this allows us to calculate probabilities for channel load and packet loss under various protocols. To derive meaningful metrics of application level service quality, we calculate the expected awareness quality (AQ) [7] for the awareness of nearby vehicles in a safety relevant area of 300 m around vehicles. A penalty for increased latency between the first encounter of a new vehicle and the first successfully verifiable secured message is implicitly included in the AQ metric. Any pre-distribution scheme should improve the AQ compared to the baseline case of including full certificates in all periodic messages. Ideally, a new pre-distribution scheme would also provide improved AQ over certificate omission schemes such as Neighbor-based Certificate Omission [8] and Congestion-based Certificate Omission [6].

A third requirement for a new pre-distribution scheme is to not influence the privacy of passengers by not compromising the unlinkability of pseudonyms used by a vehicle. This is a relevant aspect due to the basic fact that pre-distribution will leak the expected future location of vehicles and their pseudonym. Even under hierarchic, geographic, or temporal scoping rules, it should be hard for an attacker to link pseudonyms as belonging to a single vehicle. We assume an attacker model that matches common expectations of not being too powerful in terms of network coverage and not being mobile enough to outright follow a tracked vehicle.

III. SOLUTIONS AND FUTURE WORK

The solution space for certificate pre-distribution can be partitioned into three broad techniques, all of which will be augmented with directional scoping:

- n-hop dissemination.
- store-and-forward dissemination.

- probabilistic dissemination.

These techniques emphasize different forwarding methods, but are essentially all based on pushing information where it is expected to be needed. The opposite approach of pulling information that might be needed along a given trajectory would require the ability to accurately predict the positions of vehicles that are far outside communication range and to broadcast pull requests beyond the intended trajectory. Both aspects are deemed prohibitive for efficient use of the available communication channels.

To ensure the unlinkability of pseudonyms it seems intuitively plausible to not pre-distribute multiple pseudonymous certificates that belong to the same vehicle, as this creates a link between the two identities. On the other hand, the model of an attacker against location privacy assumes that a local mobile follower can watch pseudonym changes anyway. It is even possible to argue that the pre-distribution of valid pseudonym certificates without claiming the actual presence of users of these certificates enlarges the size of the k-anonymity group without having a negative effect on the service quality through an introduction of phantom vehicles.

An attractive opportunity to implement pre-distribution of certificates in VANETs presents itself in the context of proposed DCC systems [5]. Such systems optimize channel utilization by aiming to maintain a constant channel busy ratio (CBR). A CBR value below the ideal target value can be interpreted as an indication of unused bandwidth. An optional service quality enhancement scheme such as pre-distribution of certificates could be a useful consumer of such unused bandwidth.

As future work we will develop an analytical model that allows us to model the AQ of different push dissemination strategies for certificate pre-distribution, and compare the effectiveness and efficiency against existing inclusion and omission schemes. Additionally, we will investigate the impact of certificate pre-distribution on location privacy and possible interactions with DCC systems.

REFERENCES

[1] ETSI TC ITS, "ETSI TS 102 731 v1.1.1 - intelligent transport systems (ITS); security; security services and architecture," Standard, TC ITS, 2010.

[2] IEEE, "IEEE 1609.2v2 - Standard for Wireless Access in Vehicular Environments (WAVE) - Security Services for Applications and Management Messages," 2011.

[3] B. Ramsdell, "Secure/multipurpose internet mail extensions (s/mime) version 3.1 message specification," 2004.

[4] M. Elkins, "Mime security with pretty good privacy (pgp)," 1996.

[5] T. Tielert, D. Jiang, Q. Chen, L. Delgrossi, and H. Hartenstein, "Design methodology and evaluation of rate adaptation based congestion control for vehicle safety communications," in *VNC*, 2011, pp. 116–123.

[6] M. P. Feiri, J. Y. Petit, and F. Kargl, "Evaluation of congestion-based certificate omission in vanets," in *Proceedings of the IEEE Vehicular Networking Conference (VNC 2012), Seoul, Korea.* USA: IEEE, November.

[7] R. K. Schmidt and T. Leinmüller, "A spatio-temporal metric for the evaluation of cooperative awareness," in *18th World Congress on Intelligent Transport Systems*, 2011.

[8] E. Schoch and F. Kargl, "On the efficiency of secure beaconing in vanets," in *Proceedings of the third ACM conference on Wireless network security*, ser. WiSec '10. New York, NY, USA: ACM, 2010, pp. 111–116.

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

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

Identifying LTE Connectivity Hot Spots in Vehicular Environments: A Learning Approach

Christoph Ide, Kai Piontek and Christian Wietfeld

Communication Networks Institute

TU Dortmund University

44227 Dortmund, Germany

e-mail: {Christoph.Ide, Kai.Piontek, Christian.Wietfeld}@tu-dortmund.de

Abstract—**Due to the increasing demand of mobile Machine-Type Communications (MTC), the interaction between MTC and human services is a recent problem for cellular communication systems like Long Term Evolution (LTE). In order to reduce the negative impact of MTC on human communication, a Learning-based Channel-Aware Transmission (L-CAT) scheme will be introduced in this paper. The algorithm bases on a learning process of cellular connectivity hot spots and is designed for non-time-critical vehicular data applications like extended Floating Car Data (xFCD) transmissions for traffic forecast systems. The results based on real-world measurements show that L-CAT leads to a much faster data transmission that correlates with a more resource efficient MTC.**

I. Introduction and Related Work

Many mobile devices follow the routes of driving vehicles. This routes are not randomly, but follow special patterns. For example the same routes are taken by a vehicle many times: the way to work or to good friends. This fact can be used in order to predict mobility of a cellular communication device that is mounted on a car. The mobility prediction is a feasible input for traffic management systems. It can be predicted how many vehicles will pass a certain highway and if this would cause a traffic jam, some of them could be rerouted. But the mobility estimation can also be useful for communication issues. In [1] a method to improve handover quality of cellular communication systems by means of a mobility forecast algorithm is shown. However, the same routes correspond also to a similar communication connectivity. This fact will be used in this paper to improve non-real-time vehicular data applications. We present Learning-based Channel-Aware Transmission (L-CAT), a decentralized communication approach that uses a learning process of cellular connectivity in vehicular environments. L-CAT detects good communication locations, stores them, provides them to others users and uses them for the next drive at the same route to improve the communication efficiency.

The transmission scheme can be used to provide vehicular sensor data (so called extended Floating Car Data (xFCD) [2]) very efficiently to a traffic management server. This Machine-Type Communication (MTC), which can be carried e.g. by a Long Term Evolution (LTE) network, should interfere as less as possible with other human users in the communication network. LTE MTC is a recent topic in the evolution of the standardization of LTE technology follower [3] [4]. An overview about vehicular LTE data communication can be found in [5].

The usage of the channel quality for efficient cellular communication is a common approach on different layers of the communication system. Channel-dependent scheduling [6] (also for the uplink [7]) is one famous example for using the channel quality, aggregated at the base station of many active LTE users, to schedule the users to the best resources in time and frequency domain. In contrast to channel-dependent scheduling, L-CAT works on the application layer and decentralized without the need of an active communication link. In [8] a channel-aware transmission scheme is presented for random access networks. This scheme works without learning process and uses a threshold function for the transmission decision. A forecast algorithm for mobile connectivity in the area of wireless and cellular communication systems is presented in [9]. The scheme takes active performance indicators (e.g. data rate) into account for the prediction of good communication locations for WiFi and cellular communication networks.

The key contributions of this paper are:

- The L-CAT scheme that uses a learning process of the cellular connectivity (cf. Sec. II).
- L-CAT implementation for the performance evaluation by real world measurements (cf. Sec. III).
- L-CAT leads to a much faster data transmission due to a local transmission decision that is based on the LTE connectivity (cf. Sec. IV).
- The faster data transmission is correlated with a more resource efficient communication, so that L-CAT needs less physical resources for the same MTC payload size.

II. L-CAT Concept

In [10], a channel-aware transmission scheme is proposed for xFCD transmissions. The transmission policy is extended by a learning component of the cellular connectivity in this paper. Fig. 1 illustrates the L-CAT concept. The transmission

decision is done locally in the LTE device mounted on a vehicle. It depends on the data priority and cellular connectivity. The LTE connectivity estimation is based on two components, historical data and live measurements of passive indicators like Reference Signal Receive Power (RSRP) and Reference Signal Received Quality (RSRQ). The historical data contains passive as well as active connectivity indicators (e.g. local dependent data rate) and the historical mobility. The live measurements are only focused on passive performance indicators so that no active connection to the cellular communication system is needed before the transmission decision is done locally in the LTE device.

The cellular communication carries the xFCD in the uplink as well as the traffic state and forecast from the traffic server in the downlink. Connectivity maps that are generated by the devices are also stored in the server. We suggest a bidirectional update of the map in the vehicular and server via WiFi (e.g. ones in a week).
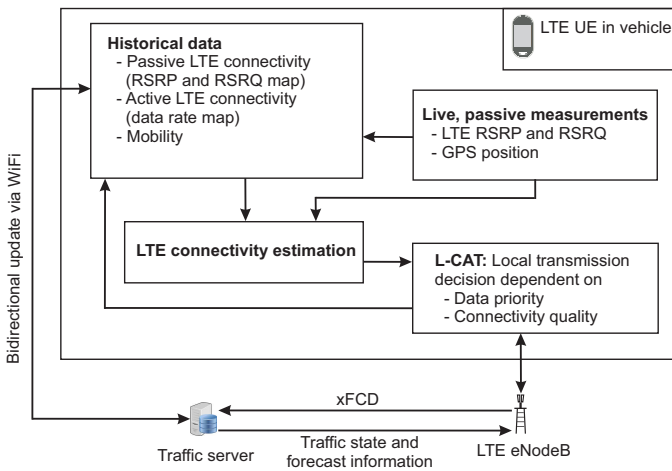

Fig. 1: Illustration of the L-CAT Concept.

III. IMPLEMENTATION OF L-CAT FOR LTE FIELD TESTS

This section summarizes the so far implemented functionalities of L-CAT. According to Fig. 1, we implemented a first version in an Android app. The app contains four basic components that are needed for L-CAT in the device:

- A look up map, where the historical connectivity map is stored. In the app, local dependent connectivity is simplified by good connectivity hot spots. These are Global Positioning System (GPS) locations with a radius that describes regions with a very good connectivity (we use RSRP > -90 dBm).
- A measurement function that monitors passively the current LTE parameters.
- The LTE connectivity estimation is simplified in a function that calculates the time until the next connectivity hot spot is reached by the current GPS position and velocity.

- The transmission decision in our app depends on good connectivity hot spots. We transmit a User Datagram Protocol (UDP) packet of 100 kByte if we are in such a hot spot or if we predict that the next hot spot is more than five minutes away.

The measurements are performed in the public LTE network of Deutsche Telekom. As performance indicator we measure the transmission time. We will show in the result section that this indicator is strongly correlated with the number of LTE Resource Blocks (RBs).

IV. RESULTS

For the performance analysis of L-CAT, we performed field tests in Dortmund, Germany. A map with measured LTE RSRP values of one example test drive is illustrated in Fig. 2. The RSRP over time is shown in Fig. 3. On the one hand, the driving route includes regions where no LTE coverage is given: e.g. at location B, a tunnel shadows the LTE signal. On the other hand, there are locations where the LTE signal quality is very good. Nearby to position D, an LTE base station is located. This results in RSRP values to up to -75 dBm.

The average RSRP is -99 dBm. This is also the average value if no CAT is applied, which means that the data is transmitted periodically. By applying L-CAT, the average RSRP
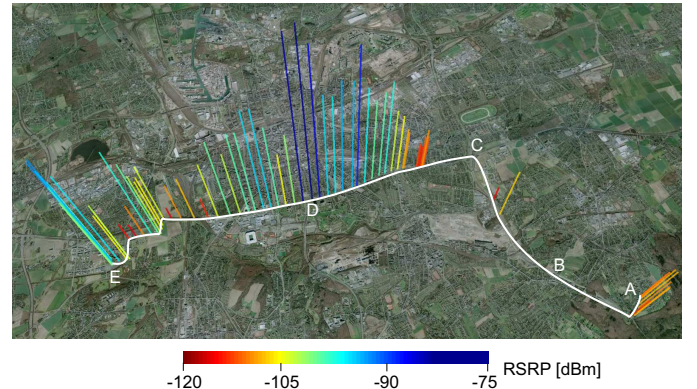

Fig. 2: Map of Example Measurement with RSRP Measurements. © Google.


Fig. 3: RSRP over Time for one Example Measurement incl. RSRP Improvement due to CAT and L-CAT.

for triggering an LTE communication, increases to -85 dBm. This is due to the fact that locations D and E are identified as hot spots. For CAT, an average RSRP of -89 dBm can be achieved. The communication under better channel conditions leads to a much faster data transmission. Fig. 4 illustrates the transmission time of 100 kByte uplink data as box plots. These results are gained from a measurement campaign of 2 h duration. It can be seen from the figure that a very wide range of transmission times is given for a transmission without CAT. The very high values result from LTE communications at location A or between D and E, where the connectivity is very bad. For CAT, the average transmission time decreases to 540 ms in contrast to 720 ms for a periodical transmission. The fact that many transmissions are also very fast (< 1 s) without CAT is caused by the good LTE coverage in the surroundings of location D. By applying L-CAT, the average transmission time can be deceased to 430 ms. Furthermore, peaks with a very high transmission time can be avoided by the usage of these hot spots. In our measurement campaign, the maximum transmission time with L-CAT is 730 ms. These results show that the LTE connectivity based on a passive indicator (we use the RSRP) recorded in previous measurements can be used in order to optimize an active LTE communication in the presence.

We have shown in [11] that the same LTE deployment provides a scheduler that is similar to a max rate scheduling for small packets. This means that users with a very low RSRP value are assigned less RBs in the frequency domain than users with good channel conditions. This leads to the fact that for a low RSRP, the duration of the transmission is not only longer because more total resource are necessary (due to a more robust modulation and coding and retransmissions), but also because the resources in the frequency domain are restricted. However, there is still a strong correlation between transmission time and total number of RBs given. This relationship is shown in Fig. 4. These LTE field measurements are gained with the same setup used in [11]. By means of a real-time spectrum analyzer, the RBs in time and frequency domain are captured at seven different locations with various channel conditions. The correlation coefficient between transmission time and total number of RBs is 0.82.

V. Conclusion and Future Work

In this paper, we have presented L-CAT, a decentralized transmission protocol for vehicular data applications. The transmission decision is based on a learning process of the cellular communication connectivity. The performance of L-CAT is analyzed by real-world measurements. It is shown that the communication policy leads to a much faster vehicular data transmission that is correlated with a much higher LTE resource efficiency. In the future, we will validate the L-CAT scheme by protocol simulations and extent the transmission decision by taken application characteristics into account. In
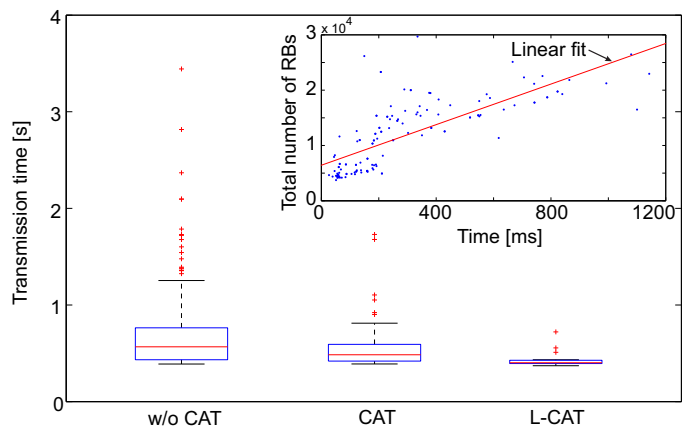

Fig. 4: Comparison of the Transmission Time for without CAT, CAT and L-CAT as well as Correlation between Transmission Time and Total Resource Allocation (100 kByte).

addition, we will develop more complex forecasting algorithms in order to predict the mobility in challenging scenarios.

References

[1] S. Michaelis and C. Wietfeld, *Comparison of User Mobility Pattern Prediction Algorithms to Increase Handover Trigger Accuracy*, The 63rd IEEE Vehicular Technology Conference 2006, Melbourne, Australia, May 2006.

[2] M. L. W. Huber and R. Ogger, *Extended Floating-Car Data for the Acquisition of Traffic Information*, World Congress on Intelligent Transport System, Toronto, Canada, 1999.

[3] 3GPP TR 23.888 V11.0.0, *System Improvements for Machine-Type Communications*, Sep. 2012.

[4] 3GPP TS 22.368 V12.2.0, *Service Requirements for Machine-Type Communications*, Mar. 2013.

[5] G. Araniti et al., *LTE for Vehicular Networking: A Survey*, IEEE Communications Magazine, vol. 51, no. 5, May. 2013.

[6] J. Jang and K. B. Lee, *Transmit Power Adaptation for Multiuser OFDM Systems*, IEEE J. Sel. Areas Commun., vol. 21, pp. 171-178, Feb. 2003.

[7] J. Lim et al, *Channel-Dependent Scheduling of Uplink Single Carrier FDMA Systems*, IEEE 64th Vehicular Technology Conference (VTC-Fall), Montreal, Canada, Sep. 2006.

[8] S.-H. Wang, A.-D. Lin and Y.-W. P. Hong, *Channel-Aware Transmission Control for Cooperative Random Access Networks*, IEEE 76th Vehicular Technology Conference (VTC-Spring), Taipei, Taiwan, May. 2010.

[9] A. J. Nicholson and B. D. Noble, *BreadCrumbs: Forecasting Mobile Connectivity*, ACM International Conference on Mobile Computing and Networking, San Francisco, USA, Sep. 2008.

[10] C. Ide, B. Dusza, M. Putzke and C. Wietfeld, *Channel Sensitive Transmission Scheme for V2I-based Floating Car Data Collection via LTE*, IEEE International Conference on Communications (ICC), Ottawa, Canada, Jun. 2012.

[11] C. Ide et al., *Performance of Channel-Aware M2M Communications based on LTE Network Measurements*, 24th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC), London, UK, Sep. 2013.

Real World Privacy Expectations in VANETs

Michael Feiri
Services, Cybersecurity and Safety
University of Twente
The Netherlands
Email: m.feiri@utwente.nl

Jonathan Petit
Services, Cybersecurity and Safety
University of Twente
The Netherlands
Email: j.petit@utwente.nl

Frank Kargl
Institute of Distributed Systems
University of Ulm
Ulm, Germany
Email: frank.kargl@uni-ulm.de

Abstract—**Vehicular communication technology is nearing deployment in the market. We see initial plug tests in 2013 to confirm interoperability of multiple independent implementations. As the entrance into the market is coming closer it is time to consider the privacy expectations of the relevant standards. These expectations are built upon location privacy through unlinkable pseudonyms. In this paper we focus on the real world privacy expectations that can be fulfilled in the first generation of vehicular communication technology using pseudonymity. What level of privacy is really achievable and does the effort to achieve this level privacy justify the cost and complexity of introducing pseudonymity into vehicular communication?**

I. INTRODUCTION AND RELATED WORK

Privacy for passengers of cooperative vehicles was identified as a requirement for market acceptance quite early in the process of developing vehicular communication infrastructure. The SEcure VEhicular COMmunication project (SeVeCom) [1] collected relevant attacker scenarios and proposed pseudonyms as a useful approach to provide privacy in vehicular contexts. The use of pseudonyms does not imply anonymity, because short-term identities are still attached to vehicles to ensure accountability and non-repudiation. The key requirement for the effectiveness of pseudonyms is their unlinkability for attackers, while authorities may have the ability to resolve pseudonymous identities to the owner of a vehicle.

Major standardization efforts for vehicular communication systems at ETSI [2] and IEEE [3] consider pseudonymity in their security architectures. However, important details remain underspecified and subject to research. The biggest open question concerns the strategy for pseudonym changes, which has a large influence on the effectiveness of pseudonyms. Previous research efforts have already highlighted challenges of performing effective pseudonym changes. According to these efforts it requires drastic measures, such as silent periods [4], [5] or context sensitive collaborative operations in mix-zones [6], [7] to ensure meaningful k-anonymity. Only recently have researchers started to investigate the impact of pseudonym change strategies on service quality [8]. Nevertheless the full consequences and practicability of pseudonym change strategies in realistic environments are still not well understood.

The implications of these issues raise concerns about the practicability of meaningful privacy guarantees through pseudonym changes. Is effective location privacy attainable without severe penalties for service quality? Is it attainable if a significant amount of collaborative vehicles is simply unwilling or unable to participate in adequate pseudonym change protocols? What kind of attacker can pseudonym change protocols even protect against?

II. ASSUMPTIONS AND ATTACKER MODELS

A. Service quality assumptions

The potential for Sybil attacks has been identified in previous works related to security and privacy in vehicular networks [9], [10], which is a reason to strictly limit the validity of pseudonymous certificates. Recommendations for deployments of pseudonymous certificates suggest lifetimes of around five minutes [11]. However such configurations would prevent pseudonym change strategies that rely on unpredictable context sensitive and/or collaborative pseudonym change strategies. Any unpredictable pseudonym change strategy requires the availability of multiple valid pseudonyms. Proof-of-work systems might counter simple sybil attack scenarios, though fundamentally the risk of sybil attacks remains.

Recent research by Lefevre and Petit [8] has highlighted the severe impact of silent periods [4], [5] as part of a pseudonym change strategy on service quality of Intersection Collision Avoidance (ICA) applications. This observation is unlikely to be limited to ICA applications. Cooperative awareness is the fundamental building block of many safety applications in vehicular networks, such as ICA. An unfortunately timed pseudonym change could break the stability of cooperative awareness. The basic position beacons that all vehicles are expected to broadcast to announce their position and trajectory are sometimes even called Cooperative Awareness Messages (CAM) [12]. These are mandatory messages and the awareness of the exact position of surrounding vehicles is a key enabler for most safety applications. The need for awareness of surrounding entities is a fundamental requirement. Privacy preservation efforts must not interfere with this requirement. If a fully anonymous communication protocol was available, it would not be an applicable solution for vehicular communication networks. This is because it would make entities untrackable even in close proximity, thus breaking the correctness of the awareness of surrounding vehicles. Local trackability is the foundation of cooperative safety.

As pointed out by Lefevre and Petit [8], if pseudonym changes include long silent periods, it would become untenable to build services that provide safety critical services. It appears reasonable to only allow silent periods in situations without any safety relevant interactions with other vehicles. However, it is not predictable if and how frequently such situations

will occur. Furthermore due to hidden station effects even the detection of such situations is unlikely to be reliable enough for consideration in combination with safety critical applications.

Mix Zones [6] have been proposed as a way to collaboratively perform pseudonym changes. This technique can give a reasonable amount of expected k-anonymity even under the assumption that an attacker can observe the entire pseudonym change process. The Mix Zones concept achieves a considerable effectiveness in this scenario, however the attacker is considered to be a passive observer. The synchronization of pseudonym changes with other entities implies that privacy decisions depend on external input. Unavailability, inability or even malicious unwillingness to participate in a pseudonym change process might prevent vehicles from ever changing their pseudonyms. Additionally the adherence to as combined silent period would be problematic for the above mentioned reasons. This also applies to encrypting messages instead of stopping to send messages, as proposed by Freudiger (CMIX) [7]. The potential inability of nearby vehicles to process messages would have a similarly negative effect on service quality, while high resolution tracking would still allow for tracking of even encrypted beacons.

B. Attacker models

The natural upper bound for an attacker is an all seeing observer with the ability to perform active attacks. A combination of mix zones and silent periods could thwart an all seeing passive attacker, but as discussed previously, silent periods and reliance on cooperative pseudonym change protocols may not be realistic options in practice. A study about the effectiveness of an all seeing attacker using Multi Hypothesis Tracking (MHT) was performed by Wiedersheim et.al. [13] and shows high levels of success for the attacker, even under noisy data and extremely frequent pseudonym changes. We are not aware of effective countermeasures against an all seeing passive or active observer under realistic service quality requirements.

Even with gaps in the coverage of the attacker it is highly likely that an attacker can simply watch and match vehicles across pseudonym changes. Advanced tracking algorithms are very effective at tracking and predicting vehicle mobility. A set of studies examining plausibility checks of location claims reveals high success rates of vehicle tracking using Kalman filters [14], [15] and Particle filters [16]. Such local tracking of nearby vehicles also reveals an interesting lower bound on attacker capabilities. There is no effective way of defending location privacy against a single mobile attacker, which simply follows a vehicle. Such following can be based on the position beacons or alternatively on sensor readings that work in close proximity to the surveilled vehicle. The mobile attacker can then watch the surveilled vehicle, trivially observing and linking any pseudonym change in the vehicular communication channel.

III. FUTURE WORK

We see that the pseudonym changes are ineffective against powerful all seeing observers and ineffective against small but mobile observers. The protection level against medium sized attackers is subject to further research. It is likely dependent on attacker mobility and the coverage of the relevant area.

Covering for example intersections and considering knowledge of pseudonym change strategies should enable very effective tracking. Reliably privacy protection within the existing pseudonymity framework is only provided against small immobile attackers. We see that that pseudonym changes under realistic assumptions can only protect effectively against this kind of weak attacker. It becomes reasonable to wonder if pseudonym changes at the beginning and at the end of a trip might suffice.

Ultimately, it might be worth questioning if the cost and the complexity of implementing pseudonymous communication is justified by the limited level of attainable privacy. It might also be worth considering whether the concealment of vehicle identification is necessary to protect the privacy of passengers. Car sharing models might make it questionable to directly link vehicle ownership to the identity of the human driver. Moreover, what if an autonomous vehicle does not even carry a passenger? Meanwhile, passengers do expect privacy in the sense of being anonymous, while pseudonym change strategies can only offer unlinkability against small and medium sized immobile attackers.

REFERENCES

[1] P. Papadimitratos, A. Kung, J.-P. Hubaux, and F. Kargl, "Privacy and identity management for vehicular communication systems: A position paper," in *Workshop on Standards for Privacy in User-Centric Identity Management*, Zurich, Switzerland, July 2006.

[2] ETSI TC ITS, "ETSI TS 102 731 v1.1.1 - intelligent transport systems (ITS); security; security services and architecture," Standard, TC ITS, 2010.

[3] IEEE, "IEEE 1609.2v2 - Standard for Wireless Access in Vehicular Environments (WAVE) - Security Services for Applications and Management Messages," 2011.

[4] K. Sampigethaya, L. Huang, M. Li, R. Poovendran, K. Matsuura, and K. Sezaki, "CARAVAN: Providing location privacy for VANET," in *3rd Workshop on Embedded Security in Cars (ESCAR '05)*, 2005, pp. 1–15.

[5] L. Huang, H. Yamane, K. Matsuura, and K. Sezaki, "Towards modeling wireless location privacy," in *Privacy Enhancing Technologies*. Springer, 2006, pp. 59–77.

[6] A. Beresford and F. Stajano, "Location privacy in pervasive computing," *Pervasive Computing, IEEE*, vol. 2, no. 1, pp. 46–55, 2003.

[7] J. Freudiger, M. Raya, M. Félegyházi, P. Papadimitratos *et al.*, "Mix-zones for location privacy in vehicular networks," in *Proceedings of the first international workshop on wireless networking for intelligent transportation systems (Win-ITS)*, 2007.

[8] S. Lefèvre, J. Petit, R. Bajcsy, C. Laugier, and F. Kargl, "Impact of v2x privacy strategies on intersection collision avoidance systems," in *IEEE Vehicular Networking Conference*, 2013.

[9] T. Zhou, R. R. Choudhury, P. Ning, and K. Chakrabarty, "Privacy-preserving detection of sybil attacks in vehicular ad hoc networks," in *Mobile and Ubiquitous Systems: Networking & Services, 2007. MobiQuitous 2007. Fourth Annual International Conference on*. IEEE, 2007, pp. 1–8.

[10] F. Schaub, Z. Ma, and F. Kargl, "Privacy requirements in vehicular communication systems," in *Computational Science and Engineering, 2009. CSE'09. International Conference on*, vol. 3. IEEE, 2009, pp. 139–145.

[11] ETSI TC ITS, "ETSI TS 102 867 v1.1.1 - intelligent transport systems (ITS); security; stage 3 mapping for ieee 1609.2," Standard, TC ITS, 2012.

[12] ETSI, "Intelligent transport systems (ITS); vehicular communications; basic set of applications; part 2: Specification of cooperative awareness basic service," EN 302 637-2.

[13] B. Wiedersheim, Z. Ma, F. Kargl, and P. Papadimitratos, "Privacy in inter-vehicular networks: Why simple pseudonym change is not enough," in *7th International Conference Wireless On-demand Network Systems and Services (WONS '10)*, 2010, pp. 176–183.

[14] H. Stubing, A. Jaeger, C. Schmidt, and S. A. Huss, "Verifying mobility data under privacy considerations in car-to-x communication," in *17th ITS World Congress*, 2010.

[15] A. Jaeger, N. Bimeyer, H. Stbing, and S. Huss, "A novel framework for efficient mobility data verification in vehicular ad-hoc networks," *International Journal of Intelligent Transportation Systems Research*, vol. 10, no. 1, pp. 11–21, 2012.

[16] N. Bismeyer, S. Mauthofer, K. Bayarou, and F. Kargl, "Assessment of node trustworthiness in vanets using data plausibility checks with particle filters," in *Vehicular Networking Conference (VNC), 2012 IEEE*, 2012, pp. 78–85.

Implementation of Day One ITS-G5 Systems for Testing Purposes

Raphael Riebl
Technische Hochschule Ingolstadt
Research Centre
Email: raphael.riebl@thi.de

Christian Facchi
Technische Hochschule Ingolstadt
Research Centre
Email: christian.facchi@thi.de

Abstract—As the finalisation of several components of the *Intelligent Transport System (ITS)* network stack for communication in the 5.9 GHz spectrum *(ITS-G5)* is taking shape, it is about time to evaluate the performance of the European *Inter-Vehicle Communication (IVC)* communication variant as backed by the *CAR 2 CAR Communication Consortium (C2C-CC)* as a whole. In order to evaluate the performance of ITS-G5 stations, it is useful to study the interaction of the involved network layers. In this paper, progress made towards an ITS-G5 implementation capable to be used in embedded and simulation environments is presented.

I. Introduction

After several years of design and research in *IVC* some not yet finalised standards are now expected to get approved soon. For example, the *GeoNetworking (GN)* specification [1] is at the time of writing in the final approval procedure. Some other standards as the p amendment to IEEE 802.11, also known as pWLAN for short, are established for years already. In the context of this paper an *IVC* system is a network stack according to recent ETSI *ITS* specifications and configured as described by the *C2C-CC* Basic System Standards Profile [2].

While a few hundred milliseconds latency might not be critical for informational services, it can be fatal for safety-critical services, if data is not disseminated within a short time frame, depending on the particular use case. For that reason, it is necessary to assure, that the designed communication system as well as a particular implementation is capable to satisfy such time frames even under stressful conditions. As has been shown in [3], there arise problems like oscillating *Distributed Congestion Control (DCC)* states when channel load is just increased enough. Therefore, it is promising to investigate if it is still possible to maintain low-latency requirements for safety-critical messages.

Taking these considerations into account, testing of pWLAN based *Vehicular Ad Hoc Networks (VANETs)* should be possible in a cost-efficient, reliable and replicable manner. Furthermore, a testing setup shall be able to handle scenarios with a lot of vehicles and therefore a crowded radio channel, so a network stack's behaviour can be studied with suboptimal communication premises.

Though large-scale field tests have been conducted in the past [4], these tend to be expensive and it is hard to repeat those. As the system evolves, tests have to be carried out again though. An appropriate way out of this dilemma are simulations. In the following, an extended version of the Veins framework 2.1 [5] is presented. Veins already supports the coupling of network communication and realistic vehicle mobility simulation, which are the fundamental aspects of *IVC*. The extension covers an complete ITS-G5 stack and enables the integration of real communication devices. Since there are already experts working on measurements of the radio signal itself, e.g. as part of ETSI ITS Plugtests in November 2013, the following simulation-based test setup focuses on protocol interactions from access layer upwards, but not on physical radio propagation phenomena like multipath, reflections et cetera. Instead, a simplified model supplied with Veins [6] is used to determine if vehicles are within communication range.

Simulations have already been conducted for various aspects of *VANETs*, which are either limited to certain layers or single use cases though. Since the simulations presented in [3] are also utilising the Veins framework, they are similar to those designated in this paper. However, they had a focus on the physical and access layer. Other aspects like multi-hop transport types and its implications are not considered.

II. Design

One problem inherent to simulations is to make sure, that the simulation model and its assumptions match real world conditions appropriately. This issue is approached in the following with the conception of an ETSI ITS-G5 compliant stack named Vanetza, which is intended to be usable for event discrete simulations like Veins and is nonetheless executable on embedded systems. It can be handy to run the network stack on a real device for testing the interoperability with other implementations and thereby validate the simulation model, at least partly. Furthermore, an outline is given of a possibility to test a *Device Under Test (DUT)* with channel load generated by a simulated environment, as it is sketched in Figure 1. For this purpose a *DUT* can interact with the simulation over its radio interface. The radio interface is chosen, because it is common for all ITS-G5 devices and all layers of the *DUT* are included in this way.

As a guideline for the features required to implement for Day One usage as propagated by the *C2C-CC*, refer to its Basic System Standards Profile [2] and the accompanying document concerning *DCC* [7].

A. Simulation

Simulations enable investigation of scenarios with potentially lots of participating vehicles. Veins [6] employs OMNeT++ [8] for the simulation of network communication and SUMO [9] for vehicle mobility. Since OMNeT++ is based on the concept of discrete event simulation, time progress is controlled by an event scheduler maintaining an event list. Therefore, the network stack's awareness of time progress is caused by increasing time stamps provided by the simulation. In order to handle timeouts internal to the network stack, e.g. the BEACON timer of *GN* [1] or expiration of location table entries, the network stack can propose the next invocation time stamp. Hence an appropriate event can be scheduled. Furthermore, the Vanetza network stack has to be capable to get instantiated multiple times in the same simulation system process. This is not achievable with the proprietary network stack implementations known to us.

B. Embedded System

For demonstration of the Vanetza communication stack on embedded systems, the well-known MK2 of Cohda Wireless[1] is deployed. This device allows injection and interception of custom packets at the *Medium Access (MAC)* layer through a raw packet socket *Application Programming Interface (API)*. Since there are only constrained hardware resources available on embedded systems like the MK2, it is mandatory to keep resource usage by the network stack low. In addition, this requirement holds true for powerful machines where the simulation is executed, because these have to run possibly hundreds of network stack instances in parallel. Events from outside, i.e. data indications from the *MAC* layer or data requests from the application layer, are handled quite similar to the simulation mode. Occurrence and order of these events, however, is not scheduled in advance but caused by environmental stimuli. In contrast to the simulation mode, there is an additional timer component in embedded systems mode, which triggers the internal events of the network stack.

The *MAC* layer is accessed in both cases through a lightweight interface which unifies the network stack's communication to the simulated *MAC* interface in Veins and the raw packet socket on Cohda's MK2 device.

C. Combination of Simulation and Embedded System

Beside purely simulative or embedded use cases, a hybrid approach is useful for testing devices. Therefore, bridging the gap between both worlds by coupling simulation and a *DUT* is beneficial.

Figure 1 gives a conceptional overview of a simulation architecture, which integrates a real ITS-G5 device, so it can interact with the simulation. In the simulation a single vehicle is selected, which shall represent the Ego vehicle. It represents movement and communication behaviour of the real hardware *DUT* in the simulated world. This simulated vehicle is special, because its upper network stack is not
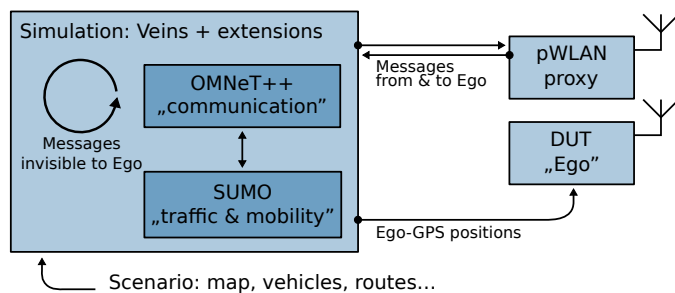
[1] http://www.cohdawireless.com/product/mk2.html

Fig. 1. Simulation architecture

part of the simulation, but executed on the *DUT*. When the Ego vehicle's simulated *MAC* layer receives a packet from its virtual environment, this packet is relayed to *DUT* through a component called pWLAN proxy. This pWLAN proxy is a MK2 device with a custom application connected to the simulation over Ethernet. At the same time pWLAN proxy and *DUT* are able to communicate over the air. However, not all packets sent by simulated vehicles are relayed. Only those packets with a signal strength above a predefined threshold are forwarded to the proxy. Signal strength degrades with increasing distance of signal path and number of obstacles along this path, as presented in [6]. The other way round packets emitted by the *DUT* are captured by the proxy and injected in the simulation and virtually transmitted to Ego vehicle's neighbours. There is also a data stream of *Global Positioning System (GPS)* positions going from the simulation to *DUT*, which allows us to fake *GPS* positions for the network stack running on *DUT*.

Since the simulation has to react on external stimuli, the event scheduler of the event discrete simulation tool OMNeT++ used by Veins has to be replaced with a custom implementation. While the default scheduler tries to make progress as fast as possible, i.e. the simulated time is not bound to wall clock time, the customized scheduler named `ProxyScheduler` tries to keep simulated and wall clock time synchronized. In Figure 2 an example of `ProxyScheduler`'s behaviour is depicted as sequence chart. Instead of continuing with the following event immediately as the default scheduler does, `ProxyScheduler` listens for messages from outside until the wall clock time catches up. If a message is received in the meantime, a new simulation event is generated and put into OMNeT++'s event data structure `cMessageHeap` appropriately. The communication connection to the outside environment is provided through the `CohdaProxy`, which incorporates the data exchange between simulation and proxy. Of course, this mechanism works only as long as the computer is able to process simulation events faster then the wall clock time progresses. If there are too many simulated vehicles, this is no longer the case. A speedup in simulation execution might be achievable through parallelization though. At the moment all events are processed in strict sequential order.

In order to keep hardware costs low, a single MK2 device has to handle the packets of several simulated vehicles. Thus
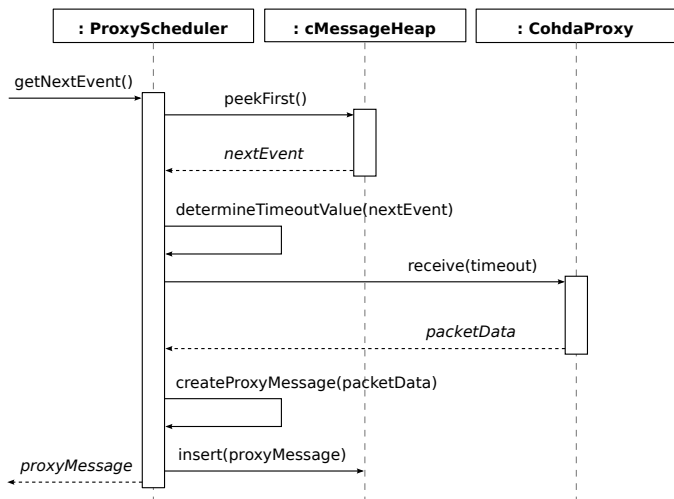
Fig. 2. Scheduling of a simulation event due to an incoming *DUT* packet

Fig. 3. Latencies for 2 ms intervals and 500 bytes payload

latency experiments were conducted to see if such a device is capable to do so. The measurement series in Figure 3 shows the round-trip-time for a packet with 500 bytes of *MAC* payload sent by one MK2 (ping) and returned by another (echo). A lot of round-trip-times are located around 4.9, 6.0 and 6.5 ms. Interestingly, there are also repeated peeks up to 10 ms with uniform intervals. Since there is no increase in packet loss at the same time, this is possibly caused by a temporarily increased processor load, e.g. a periodic kernel task delaying packet processing slightly. Channel load is controlled through an artificial pause of 2 ms between two pings, whereas echoes are replied as fast as possible. Reducing the pause by 1 ms caused an immense increase of measured round-trip-time. Increasing the pause, however, did not affect the reliability of the transmission significantly. A 2 ms pause means 500 messages per second just sent by one MK2 and the double number of messages occupying the channel in the same time. Because of size constraints, presentation of further results for other packet lengths and timings are omitted. If channel load regulations are considered as enforced by *DCC* [7], which restricts the average maximum message rate in relaxed state to less than 20 messages per second per station, the load of at least 25 neighbour stations could be represented with one MK2 device. Consequently, it is possible to reduce the number of required communication devices significantly.

III. IMPLEMENTATION

For the basic services, namely the *Physical (PHY)* and *MAC* layers of IEEE 802.11p, it is possible to rely on already available implementations. Since the interfaces provided by Veins' implementation of these layers were too limited to a certain use case, some minor customisations were necessary to make them more generic.

In the following, some implementation decisions are highlighted and explained, how these simplify reasoning about specifications. For demonstration purposes the *GN* layer as specified in [1] is used. The motivation to employ C++

for the network stack implementation stems from the desire of reasonable performance and easy interaction with *APIs* written in C like the Cohda *Software Development Kit (SDK)*. Furthermore, simulation tools like OMNeT++ are written in C++ too.

To simplify reasoning about the specification the source code should be kept as close to the specification language as possible. Consequently, maintenance code should be kept at a bare minimum, i.e. it is unfavourable to clutter code aligned to the specification text with memory management and low-level maintenance of data structures. Instead, it is favourable to have compact and expressive code.

A. Byte order

Though byte order is not a new difficulty, it can still cause some headaches. If a field in *GN* headers consists of multiple bytes, it has to be transmitted in big endian byte order [1, p. 12]. The host byte order of e.g. x86 processors, however, is litte endian. There are a number of C functions for swapping bytes, but there is no way to distinguish which kind of order is present in a particular variable. Developers have to remember the used byte order, which causes subtle errors when values with different byte orders are mixed carelessly. Since it is syntactically valid code, the compiler will not produce any warnings though. In Linux kernel development a dedicated tool is used for detection of semantic errors as e.g. illegal assignment of little endian values to big endian variables and vice versa [10]. This, however, requires an additional step during the build process. Vanetza's way to deal with byte order, however, is based on C++'s type system and therefore enforced by the C++ compiler. Types as e.g. `uint16le_t` and `uint16be_t` are distinct types based on standard C++'s `uint16_t` type with the associated byte order encoded as part of the type information. Sequence numbers within *GN* headers for example are of type `uint16be_t` and thus stored

in big endian order, whereas for arithmetic manipulation in program logic they are available in host byte order. At runtime these custom types just occupy as much memory as their underlying type. Through byte swapping they can be converted into each other seamlessly while retaining appropriate byte order information. The accompanying byte swapping functions are left out if two variables' byte order matches already and result in no performance loss at runtime. In the end, the described approach helped to get rid of byte order confusions and hence data fields transmitted in wrong byte order.

B. Encoding of Quantities

Quantities, consisting of a magnitude and a unit, are encoded in various ways in *GN* header fields. For example, heading values have a precision of 1/10 degree in *Long Position Vectors* whereas the angle in a *GeoBroadcast (GBC)* header has a precision of 1 degree. Thus, a value of 10 has to be interpreted as 1 degree and 10 degree respectively, though both fields are of 16 bit unsigned integer type. It might also be favourable to use another unit internally, because the C/C++ math library expects angle values in radian measure instead of angular degree. Fortunately, there is already a library called *Boost.Units* available [11], which allows to enrich scalar types like `double` with unit information and even provides dimensional analysis at compile time. Based on this library, special types for header field quantities can be built, which relieve developers from the burden to remember the currently involved scaled unit and convert appropriately. With the introduction of these types it was possible to track down several implementation errors. Without these types careful source code analysis and time-consuming debugging would have been necessary. Furthermore, there is now a lower risk to introduce bugs during source code changes, e.g. due to missing scaling of values.

IV. Conclusion

During the implementation of the *GN* layer, some unsound parts were discovered and reported to our project partner and after internal review to ETSI for further discussion. As one example, an unnecessary delay when a source station is going to send a *GBC* message utilising *Contention Based Forwarding (CBF)* [1, Annex E.3] has been found. While for *forwarder operation* the *CBF* timers are essential, they are meaningless for *source operation* because the message originates from source station. In the worst case this introduces a delay of 100 ms, the maximum duration of a packet in the *CBF* buffer. This reduces the total time available for a receiving vehicle to react on a possibly safety-critical message.

Another example concerns the buffering of *GBC* packets. Since the *CBF* algorithm might execute the greedy forwarding algorithm, a *GBC* packet might be buffered in a unicast buffer by greedy forwarding. Consequently, the intended behaviour of the *GN* layer is unclear at this point. In this case a proposal has been made to buffer it in the semantically correct broadcast buffer and re-execute the forwarding algorithm which caused the buffering in first place when flushing the buffer later on.

As the presented implementation has not yet been finished, possibly more feedback can be given soon for elimination of remaining glitches, so a reference system as envisaged by *C2C-CC* can be built with less effort. Furthermore, the authors are looking forward to simulate large scenarios with *C2C-CC* Day One compliant communication stacks as soon as all Day One features are covered. When the network stack running on an embedded system is adjudged to be conformant to the specification, it can be concluded due to the common code base, that its simulation counter-part is most probably conformant as well. This increases the confidence into future simulation results.

References

[1] European Telecommunications Standards Institute, *Draft ETSI EN 302 636-4-1*, 1.2.0, Oct. 2013.

[2] B. van der Kluit, T. Buburuzan, *et al.*, *C2C-CC Basic System Standards Profile*, 1.0, 2013.

[3] D. Eckhoff, N. Sofra, and R. German, "A Performance Study of Cooperative Awareness in ETSI ITS G5 and IEEE WAVE," in *Wireless On-demand Network Systems and Services (WONS), 10th Annual Conference on*, 2013, pp. 196–200.

[4] sim^{TD} Konsortium, *sim^{TD} | fakten*, Jun. 2013. [Online]. Available: http://www.simtd.org.

[5] C. Sommer, R. German, and F. Dressler, "Bidirectionally Coupled Network and Road Traffic Simulation for Improved IVC Analysis," *IEEE Transactions on Mobile Computing*, vol. 10, no. 1, pp. 3–15, 2011.

[6] C. Sommer, D. Eckhoff, R. German, and F. Dressler, "A Computationally Inexpensive Empirical Model of IEEE 802.11p Radio Shadowing in Urban Environments," in *Wireless On-Demand Network Systems and Services, Eighth International Conference on*, 2011, pp. 84–90.

[7] J. Rey, N. Sofra, and K. Sjöberg, *C2C-CC Distributed Congestion Control (DCC) for Day One*, 1.0, 2013.

[8] A. Varga and R. Hornig, "An overview of the OMNeT++ simulation environment," in *Proceedings of the 1st international conference on Simulation tools and techniques for communications, networks and systems & workshops*, 2008, pp. 1–10.

[9] D. Krajzewicz, J. Erdmann, M. Behrisch, and L. Bieker, "Recent Development and Applications of SUMO - Simulation of Urban MObility," *International Journal On Advances in Systems and Measurements*, vol. 5, no. 3&4, pp. 128–138, Dec. 2012.

[10] L. Torvalds, P. Machek, and B. Copeland, *sparse*, Oct. 17, 2013. [Online]. Available: https://git.kernel.org/cgit/linux/kernel/git/torvalds/linux.git/tree/Documentation/sparse.txt?id=refs/tags/v3.10.

[11] M. C. Schabel and S. Watanabe, *Boost.Units 1.1.0*, Boost 1.55.0, Nov. 2013.

A Modular Approach to Platooning Maneuvers

Michele Segata*[†], Falko Dressler*, Renato Lo Cigno[†]

*Computer and Communication Systems, Institute of Computer Science, University of Innsbruck, Austria
[†]Department of Information Engineering and Computer Science, University of Trento, Italy
{segata,dressler}@ccs-labs.org, locigno@disi.unitn.it

Abstract—**Driving vehicles in platoons has the potential to improve traffic efficiency, increase safety, reduce fuel consumption, and make driving experience more enjoyable. A lot of effort is being spent in the development of technologies, like radars, enabling automated cruise control following and ensuring emergency braking if the driver does not react in time; but these technologies alone do not empower real platooning. As platoons will initially share the road with human-driven vehicles, interesting new questions regarding the interactions between the two categories of vehicles arise. In this paper we briefly describe the focus of our research, i.e., the analysis of interferences caused by non-automated vehicles during a maneuver. As an example, we consider the JOIN maneuver. We define the application layer protocol to support the maneuver, together with situations that can prevent successful termination, and describe how they can be detected. We then show the validity of the idea by simulating some sample scenarios, showing either that the maneuver can successfully be performed, or safely be aborted. As final contribution, we describe our idea toward a modular approach, i.e., the development of complex maneuvers by combining smaller sub-maneuvers, aiming to ease development and safety analysis.**

I. INTRODUCTION

Better road usage and increased safety will pass through the capability of vehicles to implement cooperative driving, *platooning* for short. Albeit recently there has been a strong focus on autonomous or semi-autonomous driving [1], where Inter-Vehicular Communication (IVC) is not needed, only platooning, which requires fully developed IVC, can guarantee improved road safety, while increasing the infrastructure usage and reducing fuel consumption [2].

Platooning is much more than simple car following. Platoons must be built and split, vehicles must be able to join and leave, the platoon leader must be changed, e.g., because the driver is tired or has reached his destination. All the possible maneuvers must be supported by a proper application level protocol, providing the communication primitives or Application Programming Interface (API) needed to implement them. Indeed, this is only the starting point, as the API must provide also the means to cope with impairments, unexpected situations, partially failing communications, interfering vehicles, and finally also the emergency maneuver to relinquish the vehicles' control to all the drivers safely in case there are no more the conditions to operate the platoon.

In this paper, we briefly analyze the application layer protocol for a join maneuver, which is able to handle interferences by human driven vehicles. We implement the protocol into our platooning extension for Veins [3], [4], and show how it performs in two sample scenarios. Finally we describe how we intend to further tackle this problem in our future research, i.e., moving toward a modular approach.

II. RELATED WORK

The scientific community investigated different ways to perform maneuvers in an Automated Highway System (AHS), both with and without the infrastructure cooperation. One approach assuming infrastructure cooperation tackles the problem from a control theoretic point of view, defining the laws to control the vehicles during the maneuvers, together with higher layer mechanisms to the cars involved [5], [6].

Another high level approach is presented in [7]. The authors describe a set of different communication patterns that can be used in order to exchange data while performing a maneuver. Moreover, they define a set of controllers each of those responsible for a different situation. For instance, there is a controller dedicated to obstacle avoidance. However, not much details on how a particular fault should be detected and communicated to other parties are given.

Other works focus on mechanical and network fault handling, investigating how to detect and to react to them in order to minimize risks [8] or performance loss [9].

Recently, mixed highway scenarios have gained more attention [10]. The aim is to make platoons able to travel on public roads, avoiding the deployment of dedicated infrastructure. This poses new challenges that need to be addressed due to the presence of human drivers which might interfere with platooning operations. In [11], for example, the authors study mechanisms to perform cooperative maneuvers (e.g., a lane change by an entire platoon) to avoid dangerous situation.

In this context, the challenge of defining an application level protocol that support the different maneuvers, seen as different applications, has not been tackled to the best of our knowledge. The identification of external events due to the presence of other road users, or to other impairments as communication faults, and the algorithms that the applications deploy to react to the situation are extremely important to make platooning safe and acceptable by the broad public.

III. MANEUVERS AND SCENARIOS

To properly support platooning, a set of required maneuvers needs to be implemented. The first and most studied one is the FOLLOW maneuver, i.e., standard cruising, where interesting issues on multi-body control have to be solved and that represent the steady-state of a platoon. From a communications perspective FOLLOW can be realized with standard DSRC/WAVE beacons; a working version implementing the controller defined in [6] is available in an extension of Veins simulator we use for evaluation [3], [4].

From a protocol point of view the maneuvers to form and to manage a platoon are more challenging, e.g., JOIN, LEAVE,

(a) Normal procedure


(b) Interference by slower vehicle in front

Figure 1. Graphical sketch of different situations for the JOINMIDDLE maneuver. Automated cars shown in dark color.

MERGE, and SPLIT require a more sophisticated coordination among cars than simply receiving beacons from the other cars in the platoon. Moreover, they have additional parameters, e.g., the position in the platoon where a car wants to JOIN.

Here, we are instead interested in shedding some insight on platoons management and the communication challenges they pose, specially in face of "external threats", such as human driven vehicles interfering with the maneuvers.

As a representative of management maneuvers in this paper we focus on the JOIN procedure, assuming that one car joins the platoon in the middle, which is clearly more challenging than joining at the head or at the tail of the platoon. Besides considering the plain procedure, we also include in the protocol "escape" procedures, to handle cases when there are interferences by human-driven vehicles. For the sake of simplicity the escape is just aborting the maneuver and returning to normal platooning.

An example of a JOIN maneuver is shown in Figure 1. In the standard setup (Figure 1a), a vehicle creates a gap to let another one in. A slower human-driven vehicles may however be encountered while approaching the platoon which prevents the joiner to conclude the maneuver (Figure 1b). This situation must be detected and reported to the high layer logic which should decide what is the best action to undertake.

In this paper we consider two specific scenarios. In an extended version of this work we considered more situations, but for the sake of brevity we have chosen a subset and tried to focus on the idea:

- Scenario 1 (far truck interference): the joining vehicle encounters a truck on the lane where it is trying to join, but the truck does not prevent the conclusion of the maneuver as it is far enough.
- Scenario 2 (close truck interference): as for Scenario 1, a slow truck obstructs the joining vehicle, but this time it is forced to abort to avoid a collision.

We assume that vehicles are controlled and travel as envisioned in the SARTRE project [10]: drivers instruct the vehicle, which are otherwise entirely autonomous, through a Human Machine Interface (HMI). Actions like steering or touching the brakes disengage the Cooperative Adaptive Cruise Control (CACC) and lead to the platoon split. How this happens, however, is out of the scope of this paper.

All platooning-capable vehicles are equipped with an IEEE 802.11p compliant device, a GPS receiver, and a radar. The CACC, in order to safely perform automated close-following, needs acceleration and speed values of a subset of vehicles in the platoon. Such subset depends on the design of the controller itself. We adopt the controller designed during the PATH project [6], where each vehicle requires acceleration and speed of the platoon's leader and the vehicle in front. Other designs, with different characteristics, are possible [9], but do not influence the maneuvers we focus on.

All protocols are implemented on top of standard broadcast beacons transmitted at $10\,\mathrm{Hz}$ as commonly required [12]. "Unicast" messages are obtained by identifying the intended recipient at the application level, with a proper tagging in messages, which however can be read by any other vehicle, adding redundancy and reliability to the system.

The number of events that can interfere with platoon maneuvering are humongous, but here we only consider the one envisaged in the scenarios already described: the goal of this work is verifying the feasibility of automatic maneuvering controlled via a standard DSRC/WAVE vehicular network environment in a mixed scenario, and we do not pretend to make an exhaustive study.

We think that the scenarios we consider (Figure 1b) can be very common in case of platooning cars, which travel faster than trucks. Note that whether the truck is equipped with communication devices or not is irrelevant: it will in any case interfere with the maneuver. We want to explore if implementing proper reactions to this situation, i.e., completing the maneuver if the truck is far enough or abort it if the truck is too close, is feasible and if the situations are distinguishable with the on-board sensing (the radar).

There are other situations which can prevent the successful termination of the maneuver, e.g., when an unauthorized human-driven car enters the platoon, or when a network fault occur. We addressed these problems but we intentionally omit them in here, as the aim is mainly to show future research directions.

IV. JOIN APPLICATION PROTOCOL

Consider the JOIN maneuver, in particular with a car entering in the middle of the platoon. Three vehicles are "actively" involved in the procedure. The joiner M sends a join request to the leader L, which replies back with the position at which M is supposed to join. M then moves into position on one of the two lanes adjacent to the platoon. When M is in the position indicated by L, car F opens a gap to let M in. M and F close their gaps and the procedure terminates.

All packets sent for notification, i.e., to perform state changes, must be reliably transmitted at least to the vehicle that has the active role in the maneuver. This is obtained including in the broadcast beacons the identity of the intended recipient, which will return an application layer acknowledgement enabling the
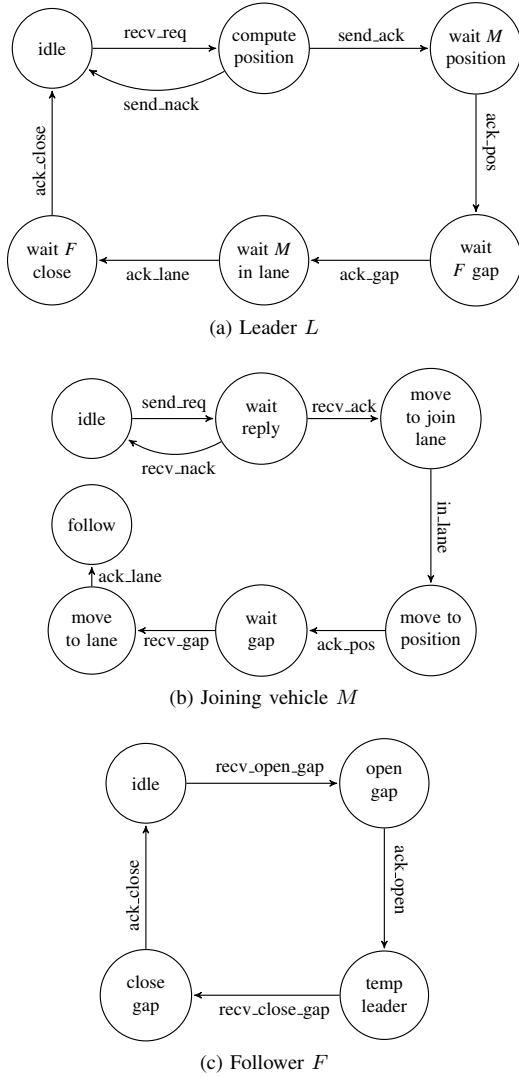
(a) Leader L

(b) Joining vehicle M

(c) Follower F

Figure 2. State machines for the vehicles involved in a JOINMIDDLE maneuver. No fault/misbehavior detection included.

Figure 3. State machine for the detection of a slow vehicle in front.

detection of lost packets and possibly triggering retransmissions. It is conceivable to achieve the same goal by using IEEE 802.11p unicast frames. This possibility, however, has the drawback that the other cars do not receive this message, so they miss part of the information about the maneuver status.

The state machines at the different vehicles that define this JOIN protocol are shown in Figure 2. We only represent the maneuver itself for the sake of clarity, without including all the details to detect faults and impairments and the actions taken to counter them: considering every possible fault or impairment is more a task for a standard specification than for a proof-of-concept prototype. In our implementation when the maneuvers cannot be completed as intended, it is simply aborted, i.e., M does not join the platoon. The 'idle' state corresponds indeed to the steady state platooning for all the cars but M, which until has received the positive join acknowledgement from L remains human driven. At the end of the procedures all car return to the steady state 'idle' platooning, thus for the joiner M entering the 'follow' state of this procedure means
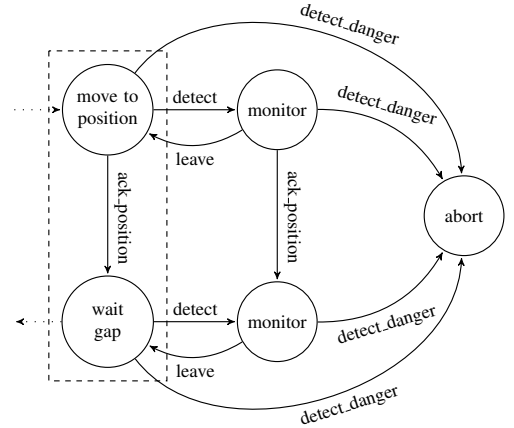
becoming a normal follower car.

We now extend the state machine to cope with the situation in which M detects a vehicle in front while trying to get in the correct position to join the platoon. This can happen during the "move to position" and the "wait gap" statuses of the state machine in Figure 2b. To handle this case we can extend the state machine of M as shown in Figure 3, where the two states enclosed in the dotted line are the same states of Figure 2b. When M detects a vehicle in front, it first switches to the "monitor" state. The radar can indeed detect objects which are up to $200\,\text{m}$ to $300\,\text{m}$ distant, which do not immediately interfere with the maneuver. Whenever a dangerous situation is detected, e.g., the Adaptive Cruise Control (ACC) is mandating to decelerate to avoid a collision, then the maneuver is aborted.

Notice that being in the "monitor" state does not prevent to continue the maneuver. If M is able to move to the join position, and the vehicle in front does not endanger maneuver's safety, it can continue waiting for F to open the gap and, in case, successfully complete the maneuver.

V. IMPLEMENTATION AND EVALUATION

For the protocol evaluation, we implemented it into the platooning enabled extension of Veins [3] and test it the aforementioned scenarios. To show the validity of our approach, we implement anomaly detection mechanism connected to basic countermeasure procedures. In particular, to detect the presence of a slow vehicle in front, we exploit data obtained from the radar, and compute the acceleration that the ACC would apply. If the deceleration becomes greater than $3\,\text{m/s}^2$, then the system issues a warning. The countermeasure connected to this warning is to make the joiner M send an abort message to the leader, disabling CACC and switching back to ACC.

We analyze the maneuver in the different scenarios from a vehicle dynamics point of view. Plots in Figure 4 show the dynamics of the vehicles in the platoon, plus the dynamics of the joiner M. The figures plot the distance from the vehicle in front as perceived by the radar.

Figure 4. Vehicles dynamics for the different scenarios showing radar distance measured by car x.

We start with the analysis of Scenario 1 (Figure 4a). The plot shows how the maneuver is correctly performed. The joiner M approaches the platoon from the side, and when in position, F and car 4 slow down to open a gap. The joiner M detects a truck in front, as shown by the radar trace, but it is far enough to let M in. The gap is then slowly closed, and the procedure terminates.

In Scenario 2 (Figure 4b) instead, M reaches the join position, F starts to open the gap, but M has to abort to avoid a collision. At this time, M switches to ACC and remains behind the truck, while F closes the gap and the platoon continues to drive as before.

These results show how the protocol can be easily extended to detect and react to anomalies in the procedure. Using the same approach, we can develop the state machines for other interferences, or for other maneuvers, analyze possible weaknesses, and study how to tackle them.

VI. CONCLUSION – TOWARD A MODULAR APPROACH

This work has proposed and analysed an application level protocol to support JOIN maneuver in two sample scenarios, showing that relatively simple logic can support complex maneuvers as letting a vehicle join a platoon in the middle of the same, while guaranteeing that in case of interference the maneuver can safely be aborted. The state machines of Figure 2, however, do not handle all possible situations that might occur, and as previously mentioned they need to be extended. When including all possible kind of interferences and network faults, state machines might become very large, and thus their verification become difficult.

Consider the JOIN maneuver we described beforehand. We can split it as follows: M performs a LANECHANGE followed by a LARGEDISTANCEFOLLOW of the car in front of F, then F should OPENGAP, M must JOINATBACK the first sub-platoon, and F should finally CLOSEGAP. Now imagine that M wants to leave the platoon. M and F should OPENGAP, M leaves invoking LANECHANGE, and the platoon can continue after F performs CLOSEGAP. Moreover, some of these sub-maneuvers can further be split, e.g., JOINATBACK or OPENGAP. The latter can be performed by combining LEADERCHANGE with LARGEDISTANCEFOLLOW.

We think that the smaller is the procedure, the easier is its design and its verification. Our future work thus includes a clear definition of a basic set of maneuvers that can be composed in order to perform more complex ones.

REFERENCES

[1] T. Luettel, M. Himmelsbach, and H.-J. Wuensche, "Autonomous Ground Vehicles–Concepts and a Path to the Future," *Proceedings of the IEEE*, vol. 100, no. 13, pp. 1831–1839, May 2012.

[2] A. Davila, E. Aramburu, , and A. Freixas, "Making the best out of aerodynamics: Platoons," in *SAE Technical Paper 2013-01-0767*, 2013.

[3] M. Segata, F. Dressler, R. Lo Cigno, and M. Gerla, "A Simulation Tool for Automated Platooning in Mixed Highway Scenarios," *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 16, no. 4, pp. 46–49, October 2012.

[4] C. Sommer, R. German, and F. Dressler, "Bidirectionally Coupled Network and Road Traffic Simulation for Improved IVC Analysis," *IEEE Transactions on Mobile Computing*, vol. 10, no. 1, pp. 3–15, 2011.

[5] R. Horowitz and P. Varaiya, "Control Design of an Automated Highway System," *Proceedings of the IEEE*, vol. 88, no. 7, pp. 913–925, July 2000.

[6] R. Rajamani, H.-S. Tan, B. K. Law, and W.-B. Zhang, "Demonstration of Integrated Longitudinal and Lateral Control for the Operation of Automated Vehicles in Platoons," *IEEE Transactions on Control Systems Technology*, vol. 8, no. 4, pp. 695–708, July 2000.

[7] F. Michaud, P. Lepage, P. Frenette, D. Letourneau, and N. Gaubert, "Coordinated Maneuvering of Automated Vehicles in Platoons," *IEEE Transactions on Intelligent Transportation Systems*, vol. 7, no. 4, pp. 437–447, December 2006.

[8] R. Rajamani, A. Howell, C. Chen, J. Hedrick, and M. Tomizuka, "A Complete Fault Diagnostic System for Automated Vehicles Operating in a Platoon," *IEEE Transactions on Control Systems Technology*, vol. 9, no. 4, pp. 553–564, July 2001.

[9] J. Ploeg, E. Semsar-Kazerooni, G. Lijster, N. van de Wouw, and H. Nijmeijer, "Graceful Degradation of CACC Performance Subject to Unreliable Wireless Communication," in *16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)*. The Hague, The Netherlands: IEEE, October 2013.

[10] C. Bergenhem, Q. Huang, A. Benmimoun, and T. Robinson, "Challenges of Platooning on Public Motorways," in *World Congress on Intelligent Transport Systems*, Busan, Korea, October 2010.

[11] S. Lam and J. Katupitiya, "Cooperative Autonomous Platoon Maneuvers on Highways," in *IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM 2013)*. Wollongong, Australia: IEEE, July 2013, pp. 1152–1157.

[12] J. Ploeg, B. Scheepers, E. van Nunen, N. van de Wouw, and H. Nijmeijer, "Design and Experimental Evaluation of Cooperative Adaptive Cruise Control," in *IEEE International Conference on Intelligent Transportation Systems (ITSC 2011)*. Washington, DC: IEEE, October 2011, pp. 260–265.

Vehicular Communications in the V-Charge Project

(Extended Abstract)

Julian Timpner, Stephan Rottmann, Lars Wolf

Institute of Operating Systems and Computer Networks

Technische Universität Braunschweig

Email: (timpner|rottmann|wolf)@ibr.cs.tu-bs.de

Abstract—Future requirements for drastic reduction of CO_2 production and energy consumption will lead to significant changes in the way we see mobility in the years to come. However, the automotive industry has identified significant barriers to the adoption of electric vehicles, including reduced driving range and greatly increased refueling times. The V-Charge Project, funded by the European Commission, seeks to address these problems simultaneously by developing an electric automated car, outfitted with close-to-market sensors, which is able to automate valet parking and recharging for integration into a future transportation system. In this paper, we describe the research challenges associated with providing a wireless Vehicle-to-X (V2X) communication infrastructure for V-Charge-enabled parking areas.

Fig. 1. The initial experimental platforms for the V-Charge project.

I. INTRO

As part of their "Europe 2020" program, the European Commission has outlined a number of ambitious targets for Europe to meet by the year 2020[1]. These targets address a wide range of social, environmental, and economic issues. Part of the strategy is to address the problem of climate change, to reduce greenhouse gas emissions, to move toward renewable sources of energy, and to increase energy efficiency. One aspect of this challenge will be the reduction in reliance on fossil fuels and the move to electric motor vehicle transport. However, the automotive industry has identified significant barriers to the electrification of vehicles, including reduced driving range and increased refueling times [1]. The European V-Charge Project seeks to address these problems simultaneously by developing an electric automated car, outfitted with close-to-market sensors, which is able to automate valet parking and recharging for integration into a future transportation system. [2]

To illustrate the vision of V-Charge, we use the following smart car system scenario: A traveling person seeks to catch a flight, possibly having a tight schedule. At large transportation hubs like airports, the process of finding comfortable parking close to one's departure terminal is usually quite cumbersome and time-consuming. V-Charge improves this situation significantly, allowing the driver to stop the car at a designated drop-off zone in front of the terminal and to directly proceed to the departure gate. Meanwhile, a back-end server, called the V-Charge ParkingManager [3], which is in charge of the efficient parking resource management, provides the vehicle with relevant mission information, such as an assigned charging station or parking spot, a map of the premises, etc., via Vehicle-to-Infrastructure (V2I) communications. These data allow the

vehicle to autonomously navigate and maneuver to its assigned target destination. Similarly, when the driver returns, he can remotely command the vehicle to the pick-up zone by issuing an according request to the back-end using his smartphone.

This implies three major fields of research: (i) *vehicle functionality*, onboard localization, detection of static and dynamic obstacles, and on-board planning using only close-to-market sensors, (ii) *logistics*, optimal scheduling of charging stations and assignment of parking spots, and (iii) *infrastructure*, development of a secure and reliable communication framework to store and share a database of information about the parking area. [2]

This paper will present an overview of the research fields (ii) and (iii) of the V-Charge project, with a focus on the communications aspect. The experimental platform consists of two VW Golf, shown in Fig. 1, which have been modified to support fully automated driving using only close-to-market sensors.

II. PARKING MANAGEMENT

Since the number of charging stations at large parking areas, due to cost reasons, will be limited, the search for an available (and charging-capable) parking spot will be typically even more complicated and time-consuming for electric vehicle (EV) drivers than for drivers of internal combustion engine (ICE) cars. V-Charge therefore provides an automated parking and charging system, based on a central back-end server which is in charge of an efficient parking resource management. It also provides each vehicle with relevant mission information allowing it to navigate to its assigned target destination. Since the developed management algorithms are generally applicable, however, they are also beneficial for non-autonomous (e.g., human valet parking) or non-electric usage scenarios.

To enable this system functionality, two main contributions to the management and infrastructure part of the project are

---

[1] http://ec.europa.eu/europe2020/europe-2020-in-a-nutshell/targets/index_en.htm

made. First, the abovementioned concepts for efficient parking management are developed. Based on driver requirements, e.g., prospective parking time, current battery charging level and required travel distance, the V-Charge ParkingManager [4] assigns (schedules) available parking resources, such as regular parking spots and, in particular, scarce charging stations to connected vehicles. Requirements for charging station scheduling as well as a short overview of first evaluation results are given in [5]. Several scheduling algorithms have been developed and evaluated [3] in detail in a dedicated simulation environment, considering different usage scenarios. For the simulation setup, real-world parking statistics obtained from Hamburg Airport and the City of Braunschweig, Germany, are used.

Second, a framework for V2I and Vehicle-to-Vehicle (V2V) (both terms are often subsumed as V2X) communications is developed and described in Section III. This framework enables the distribution of mission information to connected vehicles. Mission information includes assigned charging stations or parking spots, a digital map of the premises with marker and docking positions at charging stations, etc. Sensor data aggregated by roaming and parking vehicles, e.g., road conditions and parking occupancy (relevant in mixed-mode scenarios), are transferred back to the server, where they are merged with the central map. Further, the framework provides global system monitoring and remote debugging. Of course, state-of-the-art security and trust concepts are factored in, as described in Section IV. Driver interaction (status check, drop-off, pick-up) is realized via mobile user devices (smartphones). In Section V, we show how to collect network connectivity information and how to use it for advanced parking management decisions.

III. V2I Communications Architecture

All in-vehicle modules, such as sensors, localization and on-board planning, share information in the form of broadcasts or point-to-point connections using the Data Distribution Service (DDS) middleware. Therefore, we examined the feasibility of exploiting the middleware on the wireless link for V2I communications (as suggested in [4]). This would afford a homogeneous protocol stack and possibly synergy effects between project partners. Our evaluations, however, have shown that in our usage scenarios, DDS on the wireless link has some drawbacks such as lack of addressing specific nodes, high overhead for implementing security, and lack of multi-hop support. Thus, DDS will be used both on the vehicle and server side, but not between them.

Instead, an efficient and powerful framework based on a Delay-/Disruption-Tolerant Network (DTN) [6] is being developed for V2X communication. Because of the limited communication range of wireless radios and the highly dynamic structure of Vehicular Ad Hoc Networks (VANETs), often there is no end-to-end path between any two network nodes willing to exchange data. DTNs overcome this intermittent connectivity with a store-carry-forward approach and therefore do not require stable links. In a DTN, messages are stored at a node as long as there is no next hop available for forwarding a message. The node carries the message along its way and forwards it as soon as a new connection for forwarding becomes available. Additionally, as an overlay network, DTNs can bridge different network technologies (e.g., IEEE 802.11


Fig. 2. Mission control architecture. A disruption-tolerant network and the DDS middleware are used for V2I and intra-component communications, respectively.

and IEEE 802.15.4), ensuring information exchange across network limits and thus reaching a larger set of potential communication partners.

Although originating from interplanetary communications, DTN seems more practical than DDS for the transmission of mission-related traffic. For this purpose, we utilize IBR-DTN[2] [7], a lightweight and modular Bundle Protocol [6] implementation. Due to the integrated routing modules, no direct connection between the source and the individually addressable sink of the data is required—the routes can be determined dynamically. While assuming direct connections between nodes in the parking lot as the default, vehicles or stationary nodes can also be used as relay stations to extend the network coverage, as described in Section V. Moreover, data can easily be encrypted, but still be routed via other nodes. An overview of the resulting communication architecture is given in Fig. 2.

IV. V2I Security

An interesting research challenge in the context of vehicular DTNs is the question how drivers can securely register their vehicle with the V-Charge service and deploy keys for securing V2I communications [8]. By means of a smartphone-based registration and key deployment process for V2I communications, we are able to achieve a high degree of user independence from third parties, since nobody but the owner (not even the OEM) ever possesses the vehicle's private key. Further, our open and easily auditable protocol warrants user trust in the underlying cryptographic principles. Although we propose a solution aiming at the V-Charge project, our concepts are more generally applicable. For instance, any vehicular cloud service relying on a Public Key Infrastructure (PKI) could use the same process, since we provide an application-independent means for secure key deployment without entrusting the service provider or a third party with the private key. Moreover, the proposed solution is applicable to vehicles that come pre-deployed with the required communication technologies as well as refitted ones. Besides local ParkingManagers (PMs) communicating securely with vehicles via DTN, the V-Charge service consists of several central components: (a) the CustomerManager, (b) the Authorization Server, and (c) the DTN Certificate Authority (CA).

(a) The CustomerManager [4] provides registration functionality and all necessary service methods (pick-up, drop-off, status checking) via RESTful Web services that are being interfaced by the V-Charge smartphone application. All methods of
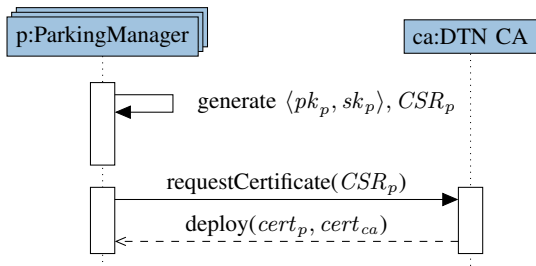
---

[2]http://www.ibr.cs.tu-bs.de/projects/ibr-dtn/

Fig. 3.   Deployment of certificates issued for local ParkingManagers.

Fig. 4.   Overlay heatmap of the RSS in an indoor scenario.

the interface are protected by the OAuth 2.0 standard [9]. (b) A central OAuth Authorization Server handles all OAuth sessions and provides verification methods. The connection between smartphone and the RESTful Web services itself is secured by SSL. (c) The DTN CA manages certificates deployed to all participants in the DTN, i.e., PMs and vehicles.

As depicted exemplarily in Fig. 3, a public/private key pair $\langle pk_p, sk_p \rangle$ is generated by each PM p. The DTN CA processes p's Certificate Signing Request (CSR) after it physically received it from p and sends back an certificate $cert_p$ signed with the CA's secret key sk_{ca}. After $cert_p$ has been deployed on p, vehicles equipped with the DTN CA's root certificate $cert_{ca}$ can verify if Bundles are affiliated to a PM (i.e., signed with sk_p) and are thus allowed to issue control commands (i.e., drop-off, pick-up). For more details, the interested reader is referred to [8].

V.   CONNECTIVITY MAP

In order to allow the central parking management system to efficiently and safely communicate with vehicles and to send them to assigned parking spots and charging stations, it is paramount to ensure that vehicles will not be sent to areas of the parking area/garage where the reception of the wireless signal is too weak or where there is no reception at all, as we have previously described in [10]. Otherwise, a vehicle sent to such an area might not be able to communicate with the parking management system anymore and will thus not be able to receive further mission commands—it would require manual intervention to retrieve it. Further, depending on the operational use case, there are extended requirements to the communication link that go well beyond of what a simple "reception/no reception" evaluation is able to achieve. For instance, some use cases (emergency stop, operating data monitoring, etc.) require a certain minimum Quality of Service (QoS) level in terms of reliability, latency or throughput.

The concept of the Connectivity Map [11] is used to meet these requirements. The basic idea is that vehicles in the parking area perform measurements in terms of current communication properties as a byproduct of V2I communications. The collected data, exemplarily depicted in Fig. 4, can be transferred and stored on a central back-end, allowing the estimation of the network characteristics for other vehicles. Therefore, areas where the QoS requirements cannot be met or where "white spots" occur can be determined in advance to avoid them in future mission planning or to use the network characteristics for sophisticated scheduling.

In the following, we will present some use cases for the connectivity map which allows us to create a cost- and energy efficient system.

A. Dynamic Optimization of Utilization of Infrastructure Components

The installation of (wired) infrastructure for providing network coverage of parking garages or large parking areas near fairgrounds is quite cumbersome and expensive. Solar-powered and battery-supported access points with a wireless (directional) uplink can be used to simplify the installation. Energy restrictions, however, require taking efficiency and energy consumption into account. Options are to power those nodes up only when they are absolutely needed or switching between different links for communication, depending on the application's needs. For example, a radio with low energy requirements, e.g., IEEE 802.15.4, may have a small (but in some cases sufficient) data rate, while IEEE 802.11 can achieve a higher throughput and larger communication range. For each radio, power control settings should be used to ensure a stable communication link while keeping the transmission power as low as possible.

The PM [4] makes decisions on the vehicles' positions in the car park and thus has the possibility to position them in a cluster, for instance, which allows wireless infrastructure components in deserted areas to be kept powered off.

B. Multi-Hop

Data on the network quality like signal strength or bandwidth sensed by vehicles in the parking area can be used to collect real-time, real-life information on the network status. Radio failures of single infrastructure components can be detected and according countermeasures can be taken. For example, a supervisor can be informed and the scheduling modules can handle the area as a (temporarily) white spot and navigate vehicles around the affected region.

In traditional V2I systems, vehicles in white spot areas could not be addressed by the back-end. In our system, however, the multi-hop features of the DTN software can be used to establish a communication with those cars. Thus, they can be requested to move to another parking spot with network coverage, for example.

C. QoS-based Parking Management

In [4], it was shown how a parking management system can accommodate different (and contradicting) user requirements

such as estimated parking time, state-of-charge for electric vehicles, etc. The Connectivity Map enables us to extend this system by making advanced parking management and scheduling decisions based on the abovementioned QoS requirements.

For parked vehicles, different QoS requirements can be defined. If a user wants large amounts of data such as media files or digital maps uploaded to his vehicle during the stay, it should be parked in high throughput areas. When all data transfers are finished, the car can be automatically moved in order to free high-connectivity parking spots. At all times, it has to be ensured that a basic connection to each vehicle can be established. Using multi-hop communications, this may be even ensured for white spots by adeptly parking vehicles nearby and using them as relays. Overall, this approach enables us to minimize infrastructure costs, e.g., Access Points (APs) installations and maintenance effort.

VI. CONCLUSION

In this paper, we have presented an overview of research challenges in the V-Charge project, with a focus on the communications aspect. This includes a DTN-based V2X communication architecture, that allows us to facilitate IBR-DTN's built-in features in order to increase the stability and reliability of the wireless link. Moreover, security challenges in the V-Charge system have been addressed. We have described our vision of advanced parking management decision support.

ACKNOWLEDGMENT

REFERENCES

[1] European Council for Automotive R&D, "The Electrification of the Vehicle and the Urban Transport System," http://www.eucar.be/publications.

[2] P. Furgale, U. Schwesinger, M. Rufli, W. Derendarz, H. Grimmett, P. Mühlfellner, S. Wonneberger, J. Timpner, S. Rottmann, B. Li, B. Schmidt, T. N. Nguyen, E. Cardarelli, S. Cattani, S. Brüning, S. Horstmann, M. Stellmacher, H. Mielenz, K. Köser, M. Beermann, C. Häne, L. Heng, G. H. Lee, F. Fraundorfer, R. Iser, R. Triebel, I. Posner, P. Newman, L. Wolf, M. Pollefeys, S. Brosig, J. Effertz, C. Pradalier, and R. Siegwart, "Toward Automated Driving in Cities using Close-to-Market Sensors, an Overview of the V-Charge Project," in *Proceedings of the IEEE Intelligent Vehicle Symposium (IV '13)*. Gold Coast, Australia: IEEE, Jun. 2013, pp. 809–816.

[3] J. Timpner and L. Wolf, "Design and Evaluation of Charging Station Scheduling Strategies for Electric Vehicles," *IEEE Transactions on Intelligent Transportation Systems*, 2013.

[4] ——, "A Back-end System for an Autonomous Parking and Charging System for Electric Vehicles," in *Proceedings of the IEEE International Electric Vehicle Conference (IEVC '12)*, Greenville, SC, Mar. 2012, pp. 693–700.

[5] ——, "Efficient Charging Station Scheduling for an Autonomous Parking and Charging System," in *Proceedings of the 9th ACM International Workshop on VehiculAr Inter-NETworking, Systems, and Applications (VANET '12)*, ser. VANET '12. Low Wood Bay, Lake District, UK: ACM Press, 2012, pp. 145–148. [Online]. Available: http://doi.acm.org/10.1145/2307888.2307918

[6] K. Scott and S. Burleigh, "Bundle Protocol Specification," RFC 5050 (Experimental), Internet Engineering Task Force, Nov. 2007. [Online]. Available: http://www.ietf.org/rfc/rfc5050.txt

[7] S. Schildt, J. Morgenroth, W.-B. Pöttner, and L. Wolf, "IBR-DTN: A lightweight, modular and highly portable Bundle Protocol implementation," *Electronic Communications of the EASST*, vol. 37, Jan 2011.

[8] J. Timpner, D. Schürmann, and L. Wolf, "Secure Smartphone-based Registration and Key Deployment for Vehicle-to-Cloud Communications," in *Proceedings of the ACM Workshop on Security, Privacy and Dependability for Cyber Vehicles (CyCAR '13)*. Berlin, Germany: ACM Press, Nov. 2013, pp. 31–36.

[9] D. Hardt, "The OAuth 2.0 Authorization Framework," RFC 6749 (Proposed Standard), Internet Engineering Task Force, Oct. 2012. [Online]. Available: http://www.ietf.org/rfc/rfc6749.txt

[10] T. Pögel, J. Timpner, S. Rottmann, and L. Wolf, "Estimation of Vehicular Connectivity in Autonomous Parking Scenarios," *PIK - Praxis der Informationsverarbeitung und Kommunikation*, vol. 36, no. 4, pp. 243–248, 2013.

[11] T. Pögel and L. Wolf, "Prediction of 3G network characteristics for adaptive vehicular Connectivity Maps (Poster)," in *Vehicular Networking Conference (VNC), 2012 IEEE*, Seoul, Korea, Nov. 2012, pp. 121–128. [Online]. Available: http://ieeexplore.ieee.org/xpl/articleDetails.jsp?tp=&arnumber=6407420

Improving Traffic in Urban Environments

Lara Codecà, Raphaël Frank, Thomas Engel

Interdisciplinary Centre for Security, Reliability and Trust

University of Luxembourg, 2721, Luxembourg

lara.codeca@uni.lu, raphael.frank@uni.lu, thomas.engel@uni.lu

Abstract—The vehicular traffic in the cities is increasing every year. The road infrastructure in many metropolitan areas is not able to sustain the rush-hour traffic demand and the extension of the road network cannot easily be done. There are some solution proposed to improve the traffic situation, among them, the optimization of the resources already available by means of collecting real time Floating Car Data (FCD) from the vehicles and use them to suggest dynamic routes in order to minimize travel delays. The centralized infrastructure able to achieve this goal has already been presented in "Improving Traffic in Urban Environments applying the Wardrop Equilibrium" (Codeca, L. et al., 2013). In this extended abstract we present the decentralized version of the system and the preliminary results of its evaluation.

Keywords—*Intelligent Transportation Systems, Urban Traffic, Traffic Flow Optimization, Wardrop Equilibrium*

I. Introduction

The vehicular traffic demand in metropolitan areas is increasing and the road infrastructure has often reached its limit; alternative means of transportation will play a major role in relieving traffic congestion, but it is likely that with the economic growth and its flexibility requirements, the demand for individual mobility must be taken into account.

The level of complexity of a model able to describe the metropolitan environment is high due to its characteristic: short road segments, different types of intersections, different types of vehicles, pedestrians, and unexpected obstacles. In the past, civil engineers, mathematicians and physicists proposed many models to describe and analyse vehicular mobility [1] and [2]. In order to accommodate the growing traffic demand, we propose to optimise the usage of the resources already available through a collaborative traffic management system. The solution proposed in this paper is the extension of what has been proposed in [3]. Several papers have recently proposed methods for efficiently sensing traffic-relevant information, among them [4], [5], and [6].

To optimize the resources already available, the participating vehicles receive a dynamically adapted route to their destination based on the Real Time Traffic Situation (RTTS). In this work the best route is defined in terms of the Estimated Travel Time (ETT) from source to destination and is based on the first Wardrop principle [7]. This principle is known as the user-optimal equilibrium and provides the best solution for each individual user.

The evaluation of the protocol is done by means of simulation using well-known tools such as SUMO [8], OMNet++ [9] and Veins [10]. In this work, to model the urban traffic dynamics we use the Krauss car-following model [11] that is implemented in SUMO.

II. Wardrop's First Principle of Equilibrium

In this work we apply the selfish user-centred policy defined by Wardrop's first principal of equilibrium [7]. In case of network games based on congested transportation or telecommunication networks, an accepted solution is provided by the Wardrop equilibria. The first principle states that every player selects a route that minimises the travel cost between source and destination. We shall call $R_{s,d}$ the set of all the possible routes from the source s to the destination d and the route chosen by the vehicle i as $route^i$.

$$route^i_{s,d} = \min(R_{s,d}).$$

Since users selfishly choose their routes to minimise their costs, this principle is known as the user-optimal equilibrium. It must be taken into account that the solution is not necessarily system-optimal. Wardrops second principle is the one related to the optimality of the whole system. It states that users minimise the overall total travel time of the system. Due to the fact that our aim is to find an efficient solution deployable in a city, the selfish solution implies an immediate benefit for the single user and thus increases the chance of a greater initial acceptance.

III. Protocol description

The topology we used is a 10x10 Manhattan grid. Each road segment measures 500 meters, has the same priority, and has only one lane in every direction. Each intersection follows the right-before-left priority rule, and deploys a 802.11 access point. The architecture of the system is based on a Vehicle-to-Infrastructure (V2I) communication network in which the vehicles in the monitored area communicate with a local Traffic Coordination Point (TCP) using their On-Board Unit (OBU). The communication protocol is the same as presented in [3] and is divided in two parts, Information Beaconing and Route Management. Both of these parts need to use a reliable overview of the system.

Real Time Traffic Situation (RTTS): With RTTS we refer to all the different data collected and aggregated to estimate the traffic conditions. We use this information to compute the user equilibrium and decide the best route for a vehicle.

Information Beaconing: The participating vehicles move through the road network, while the OBU collects the traffic metrics (location, direction and speed) and sends them every 60 seconds to the local TCP located at the intersection. The TCPs aggregate these metrics dynamically to update the RTTS.

Fig. 1. Simulation scenario. 10x10 Manhattan grid topology with access points on the intersections.

Fig. 2. Comparison of the average velocity of all the vehicle in the simulation with external information.

Route Management: The OBU sends the route request containing the current location and desired destination to the local TCP. The optimal route computed using the RTTS is sent as a reply to the OBU. If the optimal route changes due to varying traffic conditions, the OBU receives a route update.

The minimum delay route is computed using Dijkstra's shortest path algorithm [12] with dynamic edge costs using as weight the Estimated Travel Time (ETT). The ETT_s for a road segment is computed using its length, l_s, divided by the average speed $\overline{v_s}$ of all vehicles currently travelling along that segment.

$$\overline{v_s} = \frac{1}{n_s} \sum_{i=1}^{n_s} v_i,$$

$$ETT_s = \frac{l_s}{v_s}.$$

where n_s is the number of vehicles on the segment and v_i is the velocity of the vehicle i. A route is composed of multiple segments and its overall cost ETT_r is obtained as follows:

$$ETT_r = \sum_{i=1}^{m_r} ETT_{s_i}.$$

where m_r is the number of segments composing the route. The minimum delay route mdr between the origin o and the destination d is the one that minimises the ETT.

$$mdr_{o,d}^i = \min(ETT_{R_{o,d}}).$$

The $mdr_{o,d}^i$ is the route provided to the vehicle i from the system.

IV. SIMULATION SETUP AND PRELIMINARY RESULTS

Simulation environment: To build the simulation environment for this scenario we used Veins, a simulation framework able to provide the bi-directional coupling mechanism to link OMNet++ for the wireless network model, and SUMO for the vehicular mobility model. SUMO provides different mobility models, for all our simulations we use the Krauss car-following model. Each trip has a source and destination chosen

at random among all the edges of the Manhattan topology and the route for each trip is computed using `duarouter`, a tool provided by SUMO that uses Dijkstra's algorithm to choose the path having the minimum cost. Here the cost is determined by the length and priority of the segment.

Evaluation parameters: Based on the results obtained previously [3], we fixed some parameters such as the beaconing time and the route update to 60 seconds. The simulation started with 10,000 vehicles and the aim was to measure the time necessary to all of them to reach their destinations. For every vehicle, we measured the *average velocity*.
In this work we evaluated the performances of the protocol with different participation rate for the users and different coverage of the intersection in terms of access points. We compared the performances of it with two different sets of information aggregated as the Real Time Traffic Situation (RTTS). In the first scenario, the RTTS is computed using the information collected form the OBU of the participating vehicles, aggregated with external information provided by the SUMO global knowledge. This can be motivated with the use of external information sources such as inductive loops and traffic cameras. In the second scenario, the RTTS is computed using only the information collected form the OBU of the participating vehicles.

Figures 2 and 3 show the average velocity in m/s for the different participation rate of users and access point coverage in the two different scenarios. In Figure 2 the RTTM used to compute the user equilibrium contains external information. By keeping in mind that the current number of finalized simulations is not yet statistically representative, we can see that the protocol is always improving the traffic situation when external information is used to maintain the RTTM, and this is true even with 25% of access point coverage. In Figure 3 the RTTM contains only the information provided by the active users through the OBU; here we can see that if there is no external information available, the coverage plays an important role. Furthermore, with 80% or more coverage, the performance is similar to the one with external information, even with low user participation. When the coverage is 60%

## Overall Average Velocity - RTTS with only OBU information

Fig. 3. Comparison of the average velocity of all the vehicle in the simulation with information collected only through the active OBU.

less, the information collected is misleading, which results in increasing the congestion instead of mitigating it.

This behaviour can be explained by the fact that when there is no network coverage, the overview of the situation is incomplete and the user equilibrium is computed with only partial information. Further investigations are needed to fully understand this phenomenon and how it can be avoided. The use of heuristics to compute the RTTS could be used mitigate the errors in absence of reliable information. Another possibility is to use more sophisticated models to compute the user equilibrium and the route updates. Finally, further investigations are required to understand and explain the anomaly in Figure 2 with 50% user's participation, and 100% coverage.

V. Conclusions and future work

In this paper we presented the preliminary results of the decentralized version of a traffic management system to coordinate vehicular flows in urban environments. It is based on real time traffic information gathered by an On-Board Units and transmitted via wireless network to the local Traffic Coordination Point (TCP) for aggregation. The routing algorithm is based on the first Wardrop principle and computes the minimum delay route using Dijkstra's shortest path algorithm with dynamic edge costs. We evaluated the system using Veins, on a 10x10 Manhattan grid topology.

The preliminary results show that using the proposed technique increases the average velocity for a network coverage greater than 80% and a user participation greater then 10%. In case of partial coverage, the protocol fails to relieve congestion and due to incomplete information it even makes it worse (possible attacker scenario). This phenomenon requires further investigations. Future work we will also consist in evaluating more parameters such as average waiting per vehicle, number of messages exchanged and number of route update per vehicle.

References

[1] Wong, G. C. K. et al. (2002). A multi-class traffic flow model: an extension of LWR model with heterogeneous drivers. Transportation Research Part A: Policy and Practice.

[2] Geroliminis, N. et al. (2008). Existence of urban-scale macroscopic fundamental diagrams: Some experimental findings. Transportation Research Part B: Methodological.

[3] Codeca, L. et al. (2013). Improving Traffic in Urban Environments applying the Wardrop Equilibrium. In Vehicular Communication and Application (VCA)

[4] Frank, R. et al. (2013) Energy-Efficient Rate-Adaptive Passive Traffic Sensing using Smartphones. Proceedings of the 12th Annual Mediterranean Ad Hoc Networking Workshop, MedHocNet 2013

[5] Kjaergaard, M. B. et al. (2009). Entracked: energy-efficient robust position tracking for mobile devices in Proceedings of the 7th international conference on Mobile systems, applications, and services, ser. MobiSys 09.

[6] Paek, J. et al. (2010). Energy-efficient rate-adaptive GPS-based positioning for smartphones in Proceedings of the 8th international conference on mobile systems, applications, and services, ser. MobiSys 10.

[7] J. G. Wardrop. (1952). Some theoretical aspects of road traffic research. Proceedings of the Institute of Civil Engineers, Part II.

[8] Behrisch, M. et al. (2011). SUMO - Simulation of Urban MObility: An Overview In: SIMUL 2011, The Third International Conference on Advances in System Simulation.

[9] Varga, A. (2010). OMNeT++. In Modeling and Tools for Network Simulation (pp. 35-59). Springer Berlin Heidelberg.

[10] Sommer, C. et al. (2011). Bidirectionally coupled network and road traffic simulation for improved IVC analysis. Mobile Computing, IEEE Transactions on.

[11] Krauss, S. (1998). Microscopic modeling of traffic flow: Investigation of collision free vehicle dynamics. Diss. Universitat zu Koln.

[12] Chen, J. C. (2003). Dijkstras shortest path algorithm. Journal of Formalized Mathematics 15.