



ulm university universität
uulm

Proceeding of the 3rd GI/ITG KuVS Fachgespräch Inter-Vehicle Communication (FG-IVC 2015)

**Raphael Frank, Christoph Sommer, Frank Kargl
Stefan Dietzel, Rens W. van der Heijden**

Ulmer Informatik-Berichte

**Nr. 2015-03
März 2015**

Proceedings of the 3rd GI/ITG KuVS Fachgespräch Inter-Vehicle Communication (FG-IVC 2015)

March 19-20, 2015, Ulm, Germany

Preface

On March 19th and 20th, twenty motivated researchers working on inter-vehicular communication (IVC) gathered in Ulm, Germany, for the third GI/ITG KuVS Fachgespräch on Inter-Vehicular Communication. On both days, they engaged in intensive discussions about the state of the field and future research directions.

This year, the Fachgespräch's scope has been extended beyond communication, now also including cooperative driving, which is seen as one of the biggest upcoming challenges in automotive research. In cooperative driving, intelligent, self-driving vehicles use communication to provide advanced features, such as platooning. Communication can play an essential role in extending the sensing range of intelligent vehicles.

We were also honored by an academic keynote on the topic of cooperative driving, given by Dr. Jonathan Petit from the University College Cork, Ireland, who addressed security challenges and solutions in smart vehicles and cooperative driving.

While we are approaching day-1 deployments, for which manufacturers have expressed their commitment for 2015, research also looks at the development of more future-looking applications involving large-scale, multi-hop cooperation of vehicles. Other important areas of on-going research include security, privacy, reliability, scalable testing environments, and simulation models. The upcoming deployment also raises questions regarding the potential of standardized yet only partially implemented features such as multi-channel beaconing and heterogeneity. These and other topics were all extensively discussed among the Fachgespräch participants.

Feedback to the organizers confirms that the concept of the IVC Fachgespräch resonated well with the participants, who all enjoyed the open atmosphere, lively discussions, and the setting in the lovely Villa Eberhardt in Ulm. We look forward to a continuation, as we are sure that also 2016 will hold many interesting topics and challenges to be discussed.

March 2015

Raphael Frank
Christoph Sommer
Frank Kargl
Stefan Dietzel
Rens W. van der Heijden

Contents

R. Frank, C. Sommer, F. Kargl, S. Dietzel, R. W. van der Heijden: Preface.....	1
P. Fuxjaeger, S. Ruehrup: A Software-Defined Radio Testbed for Car-to-X Communication Devices and Protocols	5
W. Bronzi: Bluetooth Low Energy Robustness Analysis for V2V Communications.....	9
F. Klingler: Improving Multi-Channel Beaconing in Vehicular Networks.....	12
S. Mauthofer: Robust Usage of V2X Data in Autonomous Driving.....	14
L. Kristiana, C. Schmitt, B. Stiller: Investigating a Reliable Inter-vehicle Network in a Three Dimensional Environment.....	16
F. Hagenauer: Approaches to Heterogeneous Vehicular Networks.....	20
J. Balen: Spatio-temporal distributed background data storage and management system in VANETs.....	22
R. Riebl, S. Neumeier, C. Facchi: Inter-Vehicle Communication on the Run – Experiences From Tweaking Veins Runtime Performance.....	25
D. Förster: Discussing Different Levels of Privacy Protection in Vehicular Ad-Hoc Networks	29
M. Feiri, J. Petit, F. Kargl: The case for announcing pseudonym changes.....	31
L. Dolberg, J. François, T. Engel: Assessing Artificially Caused Congestion on Urban Scenarios: A Case Study on Luxembourg Sumo Traffic.....	34
S. Schellenberg, R. Berndt, D. Eckhoff, R. German: How Electric Vehicles Can Benefit from Vehicular Networking.....	38

A Software-Defined Radio Testbed for Car-to-X Communication Devices and Protocols

Paul Fuxjaeger, Stefan Ruehrup

FTW – Telecommunications Research Center Vienna
Donau-City-Str. 1, 1220 Vienna, Austria
{fuxjaeger, ruehrup}@ftw.at

Abstract—Testing and development of communication equipment and protocols for vehicular ad hoc networks (VANETs) is often subject to a trade-off between the cost of evaluation (simulation or field trial) and the achievable accuracy of the analysis. Whenever the impact of interference and concurring transmissions (high channel load and many hidden terminals) needs to be investigated, complex scenarios with large number of nodes need to be built. Therefore, such scenarios are often only evaluated using simulation, while expensive field tests typically cover only a limited number of vehicles. Moreover, reproducibility of results from field tests is very hard to achieve since the channel conditions are almost impossible to recreate in detail.

In this paper we show how radio frequency signals can be generated using software-defined radio transmission that include effects of fading, doppler shifts and interference. By having complete control over these characteristics and feeding the corresponding signal to a real device under test, the impact of line-of-sight versus non-line-of-sight conditions, relative velocities and multiple timely-overlapping frames at the receiver input, can be evaluated. We provide concrete results from a controlled interference measurement and show possibilities as well as limitations of this hardware-in-the-loop testing approach.

I. INTRODUCTION

Vehicular ad hoc networks (VANETs) allow the exchange of information between vehicles via wireless communication. This information can be the input to driver assistance and driver information systems, which is relevant for road safety and traffic efficiency. European and US standardization for Car-to-X communication defines a whole protocol stack from the medium access layer based on the IEEE 802.11 standard [6] up to data elements for vehicle status, such as the Cooperative Awareness Message (CAM) [4] or the Basic Safety Message (BSM) [9], which are used by applications.

Our work is motivated by the observation that channel access in VANETs is not centrally coordinated, and system performance is highly impacted by the way transceivers can cope with interference (multiple CAM messages overlapping in time at a single receiver input) and specific propagation conditions (frequency selective and time-varying channels). Therefore, the question arises repeatedly if a VANET application that has been successfully simulated also works as expected under real life conditions. Large-scale field test with hundreds of vehicles would be necessary to answer this question exhaustively, but are very expensive to maintain. If




Fig. 1. A simple example: Beaconsing messages are continuously transmitted from surrounding vehicles at an intersection. Reproducing this exact scenario and comparing the performance at the receiver over different implementations or application configurations in a field trial is extremely costly.

done at all, field tests often consist only of a few vehicles, and they cannot be replicated easily.

In order to address the cost and reproducibility issues, we have created a measurement system that allows to synthesize a highly customizable radio frequency signal which is then fed to a device under test (DUT) receiver. The goal is to create an input signal for the DUT that is indistinguishable from the signal it would receive during an actual large-scale field trial such as depicted in Figure 1.

The key element to build this measurement system is using software-defined radio (SDR) elements in the signal generation stage. This principle allows to do two things that would be infeasible using ordinary ITS transceivers: First, we can generate receive signals that reflect specific wireless channel conditions (e.g. based on collected mobility traces) by filtering the digital baseband signal prior to transmission. Second, we are able to re-create receive signals that reflect situations known to occur regularly in uncoordinated medium access networks (e.g. hidden terminal cases or MAC layer collision events). SDR tools have gained a lot of attention in the recent years, and VANET-specific IEEE802.11p transmit and receive chains have been presented in [2] and [5]. Our work is partially based on these results.

The contribution of this paper is as follows:

- We propose a method to synthetically create highly realistic baseband signals, mimicking complex VANET communication scenarios.
- We show how this hardware-in-the-loop method can be

used for testing transceivers, protocols and applications without executing field trials.

- We present a simple validation measurement showcasing the benefit of this approach over a pure simulative analysis.

II. TESTBED OVERVIEW

The testbed currently consist of the following elements: baseband IEEE802.11p signal encoder, channel impulse response generator, radio frequency up-conversion units, additive analog signal combiner, the device under test (DUT) and a monitoring station to record results. An example lab setup with a connected on-board unit on the right side is shown in Figure 2. The detailed measurement methodology we use is as follows:

- 1) Multiple baseband streams containing standard compliant IEEE802.11p frames reflecting specific transmission events (messages overlapping in time) and channel conditions (LOS, NLOS, varying relative speeds, etc) are generated offline – mimicking the situation of a number of ITS stations in the vicinity of the DUT (see [8] for an overview on channel characterization). In this stage we make extensive use of the encoder that has been developed at FTW and released under a open source license [5]. The channel models we apply have also been developed in house [1].
- 2) Up to four USRPN210 front-ends are used in parallel to individually up-convert those baseband streams to the ITS-G5 band (5.9 GHz) [3]. For optimum emulation accuracy each front-end re-creates the signal of one individual vehicle. If more than four transmitting vehicles need to be emulated at the same time, multiple transmit signals can be aggregated in one digital baseband stream prior to transmission at the cost of emulation accuracy (because signals from different vehicles then have to be up-converted using the same local oscillator and the identical LNA gain setting).
- 3) All streams are additively superimposed in radio frequency domain via a passive power combiner. The sum-signal is attenuated by fixed power attenuators to produce a given average target SNR. The resulting interference-signal is directly fed to a device under test (DUT) – without being distorted by a wireless channel.
- 4) Measurement-points that include low-SNR frame receptions are repeated to derive statistically sound results, since thermal noise at DUT being the dominant cause of frame errors in these cases.

This measurement setup allows to run reproducible experiments in a lab environment and expose a device under test to signals that have been distorted by a specific wireless channel realization and contain interfering frames. In the following section we present a simple example to demonstrate the application of this principle.




Fig. 2. Example configuration for testing an on-board unit (OBU)

TABLE I
SDR PLATFORM DETAILS

SDR Frontend Model	Ettus USRPN210
Daughterboard Model	XCVR2450
Tuning Range	2.4-2.5GHz, 4.9-5.9GHz
Reproducible Dynamic Range	-100 dBm ...-20 dBm
Local Oscillator Type	internal TCXO (2.5ppm)
DAC resolution	2x16 bits
Baseband Data-Rate	800 Mbit/s (single stream)
Baseband Bandwidth	25 MHz
Interpolation Factor	2.5
Effective Signal-Bandwidth	10 MHz

III. AN EXAMPLE APPLICATION

In this example setup, we want to answer the question whether a specific VANET transceiver implementation is able to decode IEEE802.11p frames that are overlapping in time but differ in average received signal strength at the receiver input. This feature is known as *physical layer capture* and sometimes also called *message-in-message reception*. Very few concrete hints are documented in data sheets or related work and few simulation models even implement it – mostly because it is often not thought to be relevant for VANET simulation accuracy.

During our analysis of the medium access in scenarios with high vehicle density, we realized that this feature can have *significant impact* on the overall network performance as frames-overlapping-in-time events are very common in CAM *broadcasting* scenarios as soon as hidden terminal situations start to emerge.

Thus, the question at hand is: are current receiver implementations able to successfully decode frames that are overlapping in time, and if so, under which timing and signal-to-interference ratios can this be observed.




Fig. 3. Configuration used for measurement-based analysis of the capture effect. Two individually delayed IEEE802.11p frames are superimposed in radio frequency domain and fed to the receiver input of the DUT.




Fig. 4. Exemplary snapshot of the baseband signal (magnitude only) after the combiner. The measurement system allows to reproduce arrival delay differences with a timing resolution of 40 ns.

TABLE II
MEASUREMENT PARAMETERS

TX Carrier-frequency	5.88 GHz (Channel 176)
MPDU Size	500 bytes
Modulation Rate	QPSK R=0.5
Frame length	800 μ s
Transmit Interval	10 ms
Channel model	AWGN (time invariant)
High-Power-Frame RSSI	-80 dBm
Low-Power-Frame RSSI	-80 ... - 90 dBm
Delay Step size	25 μ s
Repetitions per Delay-Step	10000

The setup used in this experiment is depicted in 3. IEEE802.11p compliant OFDM frames from two interfering stations are additively combined and fed to a device under test, replicating a simple hidden-terminal scenario. Clocks at both SDR front ends are started with a common timing strobe (PPS) and the individual baseband signals are delayed by a discrete number of baseband samples at the signal encoder host. In this example, the main advantage of using SDR front ends over IEEE802.11p transceivers for transmission is that we can control the delay between interfering frames with a resolution of 40 ns (at 25MS/s baseband rate). Note however, that local oscillators in the USRP front ends side are not synced – thereby we mimic a real world interference situation as close as possible. In this case, a simple AWGN channel model is applied at the encoding host since effects of time variability and frequency selectivity of the wireless channel shall be excluded in this first experiment. See Table II for more details about the parameters used.

Figure 4 shows two exemplary baseband signal snapshots for different timing conditions after the additive signal combiner. This interference experiment is repeated 10k times for each delay and signal-to-interference (SIR) setting.

The red curve in Figure 5 presents the final outcome. In order to visualize their significance these empirical results are contrasted by the what is modeled by the current IEEE802.11p interference model in the network simulator NS-3¹ (blue curve). The y-axis shows the frame error ratio of the higher-power frame. The x-axis depicts the arrival delay of the higher-power frame with respect to the lower-power frame. The

¹<http://www.nsnam.org>




Fig. 5. Outcome of a simple frame-overlapping-in-time-experiment for four different signal-to-interference settings. The blue line shows the results from a corresponding simulation using the current NS-3 IEEE802.11p interference model. The results differ significantly from the measurement observations and motivate a redesign of the simulation model in order to accurately predict real world behavior.

signal-to-interference ratio is noted in dB on the top-right and increases from top to bottom rows. Depending on the SIR value, the observed behavior of the DUT clearly differs from the simulation result. The most significant differences are:

- The accumulation of interference is modeled too pessimistic in the simulation model. At an SIR of 4.8dB (second row) the FER of the higher-power frames arriving earlier (negative delay) does not increase until close to when both frames *fully overlap*.
- The capture effect is clearly visible but *not modeled at all* by the NS-3 model. The fourth row indicates that this particular receiver can successfully capture as soon as the SIR exceeds 10 dB.
- The capture effect only depends on the SIR and is independent of the delay of the higher-power-frame. This contradicts what has been claimed in related work [7]: that the delay must not exceed a certain limit in order for capturing to occur.
- A subtle detail is the peak at 800 μ s in the last row – it indicates that the frame-start detector is blind and misses the next frame if there is no interframe space of at least 25 μ s after a previous frame has ended.

As can be seen from this simple example, controlled interference experiments are indispensable in order to remove uncertainty when it comes to real-world behavior of current hardware implementations. The lessons learned from this initial experiment motivates the refinement of the simulation model which eventually will enable the research community to make more accurate predictions of final system performance.

IV. CONCLUSION AND FURTHER WORK

We have shown how radio frequency signals can be generated using software-defined radio transmission that include

effects of fading, doppler shifts and interference. These baseband signals can then be fed to devices under test in order to validate certain receive and interference model assumptions. As a motivating example the packet capturing effect is investigated and demonstrates the difference between actual receive behaviour and generic simulation models. It shows that the assessment of final system performance requires a thorough study of the properties of the individual receiver implementation. The proposed SDR testbed can provide the necessary insights to build such a tailored simulation model.

Beyond validating and improving interference simulation models, this SDR-based testbed can be used for the following applications:

- Systematic comparison and benchmarking of receiver and protocol implementations
- Input generation for system-level tests and validation, where the receiving on-board unit is connected to embedded control units or driver assistance systems.
- Replay of baseband recorded in field trials in order to optimize a given receiver implementation without having to repeat the measurement.

ACKNOWLEDGMENT

The Austrian Competence Center “FTW Forschungszentrum Telekommunikation Wien GmbH” is funded within the program COMET – Competence Centers for Excellent

Technologies by BMVIT, BMWFJ, and the City of Vienna. The COMET program is managed by the FFG.

REFERENCES

- [1] L. Bernadó, T. Zemen, F. Tufvesson, A. Molisch, and C. Mecklenbräuker, “Delay and doppler spreads of non-stationary vehicular channels for safety relevant scenarios,” *IEEE Transactions on Vehicular Technology*, vol. 63, no. 1, pp. 82–93, 2014.
- [2] B. Bloessl, M. Segata, C. Sommer, and F. Dressler, “An iee 802.11a/g/p ofdm receiver for gnu radio,” in *Second Workshop on Software Radio Implementation Forum (SRIF '13)*, 2013, pp. 9–16.
- [3] “ETSI EN 302 663 – Intelligent Transport Systems (ITS); Access layer specification for Intelligent Transport Systems operating in the 5 GHz frequency band,” 2013.
- [4] “ETSI EN 302 637-2 – Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Part 2: Specification of Cooperative Awareness Basic Service,” 2014.
- [5] P. Fuxjäger, A. Costantini, D. Valerio, P. Castiglione, G. Zacheo, T. Zemen, and F. Ricciato, “IEEE 802.11p Transmission Using GNURadio,” in *6th Karlsruhe Workshop on Software Radios (WSR)*, 2010, pp. 1–4.
- [6] “IEEE Std 802.11-2012. Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications,” Mar. 2012.
- [7] J. Lee, W. Kim, S.-J. Lee, D. Jo, J. Ryu, T. Kwon, and Y. Choi, “An experimental study on the capture effect in 802.11a networks,” in *2nd ACM International Workshop on Wireless Network Testbeds, Experimental Evaluation and Characterization (WinTECH '07)*, 2007, pp. 19–26.
- [8] C. F. Mecklenbräuker, A. F. Molisch, J. Karedal, F. Tufvesson, A. Paier, L. Bernadó, T. Zemen, O. Klemp, and N. Czink, “Vehicular channel characterization and its implications for wireless system design and performance,” *Proc. IEEE*, vol. 99, no. 7, pp. 1189–1212, 2011.
- [9] “Dedicated Short Range Communications (DSRC) Message Set Dictionary,” SAE J2735, 2014.

Bluetooth Low Energy Robustness Analysis for V2V Communications

Walter Bronzi

Interdisciplinary Centre for Security, Reliability and Trust
University of Luxembourg, L-2721, Luxembourg
walter.bronzi@uni.lu

Abstract—Bluetooth Low Energy (BLE) is a low-power wireless technology operating on the 2.4 GHz Industrial Scientific Medical (ISM) band. As it gains more and more momentum for monitoring applications it is interesting to analyse in practice how BLE handles other interference on the same frequency band for a possible deployment in a vehicular context. By means of experiments, we first build and evaluate an interference testbed based on the IEEE 802.11 technology, we then establish a BLE communication between two devices and analyse what effects such configuration has on packet rate and round-trip time (RTT). We found that by generating UDP traffic on three separate access points set to non-overlapping Wi-Fi channels (1,6,11) there is a noticeable increase of bluetooth retransmissions and RTTs following the increase of interference. The association time between the two BLE devices will also get increasingly higher the fuller the frequency band gets. Finally we discuss BLE suitability for Inter-Vehicular Communications (IVC).

Index Terms—Bluetooth Low Energy, IEEE 802.11, Testbed, Experiments, Interferences, IVC

I. INTRODUCTION

One of the latest additions to the communication technologies is *Bluetooth Low Energy* (BLE), also called *Bluetooth Smart*. Although similar in some regards, BLE is not backwards compatible with previous Bluetooth versions as it uses a different controller (i.e. Physical and Link Layer). However most devices that support BLE implement both protocol stacks in dual-mode.

This low energy and low latency communication protocol has been developed to facilitate monitoring small devices that require little power to run. Common application areas include fitness, healthcare and smart homes. The protocol defines several upper layer functionalities that allow fast and easy message exchange between devices.

In our previous work [1] we investigated how BLE could be a suitable candidate for Inter-Vehicular Communications (IVC). During our tests we measured performance in terms of delivery radio and round-trip time for multiple vehicular scenarios. We concluded at the time that BLE provides a fast, low power and reliable solution for IVC for non-delay sensitive applications (e.g. dense traffic situations, platooning).

Taking into account our previous findings on BLE in its current state, we realised how important a study on interferences would impact even further its possible deployment scenarios. For this reason this work is focused on showing how co-existence between BLE and Wi-Fi looks in practice




Fig. 1. BLE data and advertising channels co-existence with Wi-Fi non-overlapping channels

and in investigating how resilient BLE communications are to interferences on the same radio frequency (RF) band coming from IEEE 802.11 enabled devices operating at the maximum of their capacity.

For the purpose of this experiment we built a static interference testbed composed of 6 Raspberry Pis, 3 of which in AP mode assigned to the only three Wi-Fi non-overlapping channels (1,6,11).

We observe in Fig.1 how by manually assigning the Wi-Fi channels we drastically reduce the amount of BLE data channels available to the protocol while still leaving unperturbed the advertising frequencies (Ch. 37, 38, 39).

By the means of experiments, we then evaluate the impact of sequentially enabling access points to the packet rate and RTT of a communication link between two BLE enabled mobile devices situated close-by. We also observe how the association time between the two mobile phones varies with the different scenarios. We thus identify and discuss BLE strengths and shortcomings that make the current implementation of BLE suitable only for a certain kind of applications.

The remainder of this paper is organised as follows. In Section II we provide a literature review. Next, in Section III, we present an overview of our testbed. Preliminary results are discussed in Section IV. In Section V we conclude and provide directions for future work.

II. RELATED WORK

BLE has been standardised by the Bluetooth *Special Interest Group* (SIG) under the Bluetooth 4.0 specification [2]. Gomez et al. [3] provided a concise overview of the BLE protocol




Fig. 2. Raspberry Pis Interferences Testbed

stack and investigated the impact of several critical parameters on its performance. They identified that there exists a trade-off between energy consumption and network performance that depends on several configuration parameters.

As BLE is mainly used for low data communications, the achievable throughput has not been previously of main concern. Regardless it is important to understand the limit of the hardware used as it plays an important role in setting up the scenarios. (For further information on modelling the maximum throughput of BLE please refer to [4]).

Most works evaluate on paper the co-existence of Wi-Fi and Bluetooth Classic (BC) [5] [6] and very few introduce BLE in the equation [7] [8]. BLE hardware is able to reuse existing BC coexistence features such as passive interference avoidance schemes like adaptive frequency hopping (AFH) [9]. BLE channels also have a different spacing compared to BC's (2MHz for BLE and 1MHz for BC) and are of two kinds: data and advertising channels. Advertising channels are specifically chosen to be in the least congested zone of 2.4 GHz band. (In orange in Fig. 1)

In our previous work [1] [10] we brought BLE to a vehicular context and analysed its performance as communication alternative for IVC.

III. BLE INTERFERENCES TESTBED

For the purpose of benchmarking the behaviour of a BLE communication between two mobile phones with a sufficient amount of interferences we build a relatively compact (2 m x 2 m) testbed.

As shown in Fig. 2 the setup is composed of 6 Raspberry Pis, 3 of which acting as access points and the other 3 as clients. For the first 3 Pis we installed *hostapd* and *dnsmasq* to configure the access points in mode IEEE 802.11g and assigned them different channels (1,6,11). To generate a sufficient amount of traffic to perturb communications we utilised *iperf* from the clients to the APs to create a UDP maximum throughput stream (≈ 35 Mbps).

For the BLE communication the mobile devices were running modified software from our previous experiments

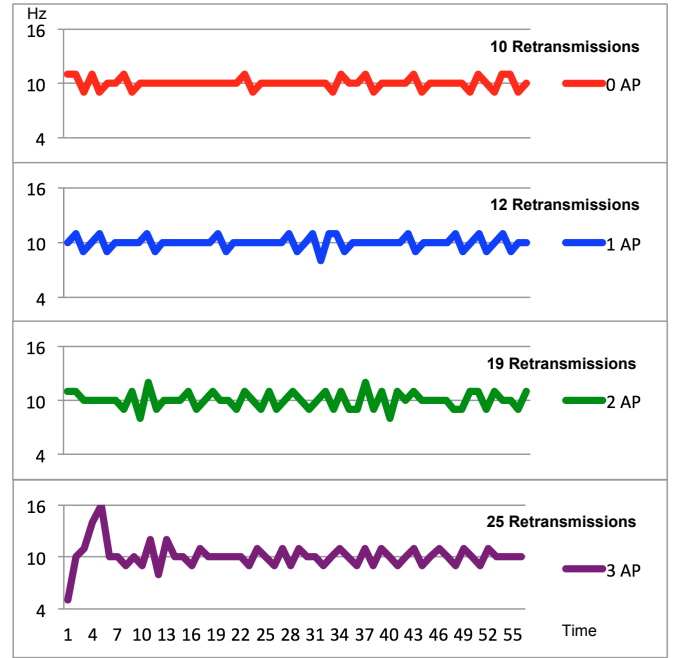


Fig. 3. Packet rate for different interference scenarios

[1] with the purpose of logging packet exchange during a timeframe of one minute (Further details in the following Section). The mobile phones were physically placed in the middle of the testbed to keep the scenario relevant to our purpose and maximise interferences.

We optimised to our best effort the test environment to ensure no other interference human or environmental would spoil the results.

IV. PRELIMINARY RESULTS

To maximise as much as possible the throughput of the BLE communication we established a constant packet size of 158 Bytes (3 for the header and 155 for the payload) at a rate of 10 Hz with which we can achieve the best reliability and stability given our hardware (≈ 12.6 Kbps). The packet size is the current maximum size that is achievable with our combination of hardware and software (iPhone6 - iOS8.1). The results are based on four different interference scenarios:

- No Access Point
- 1 Access Point with iperf traffic (Channel 1)
- 2 Access Points with iperf traffic (Channel 1,6)
- 3 Access Points with iperf traffic (Channel 1,6,11)

For each run the BLE communication between the phones was left to run for one minute. Fig. 3 shows the packet rate for each configuration. By sequentially adding more access points we notice a growing number of retransmissions with the same average packet rate. The fluctuations, even if minimal, show that our testbed is indeed having an impact to the communication.

It is noticeable by looking at the packet rate and the RTT in Fig. 4 how the first "batch" of packets is affected by the extensive interferences on the spectrum. The backlog of messages is due to the two devices having difficulties finding

a free data channel (only 9 are free - Fig. 1) and most likely the AFH algorithm doing its job in finding the best one to use.

We also noticed during our tests that, even though the BLE advertising channels are not congested, the initial association time takes longer the more interference we introduce. With no interference we recorded ≈ 1.5 s for the first packet to be sent whilst for 1 AP we have ≈ 1.8 s, for 2 AP ≈ 2.3 s and finally with 3 AP we have ≈ 3.9 s before we can initiate the communication. Again this might be an implementation behaviour and will have to be tested further.

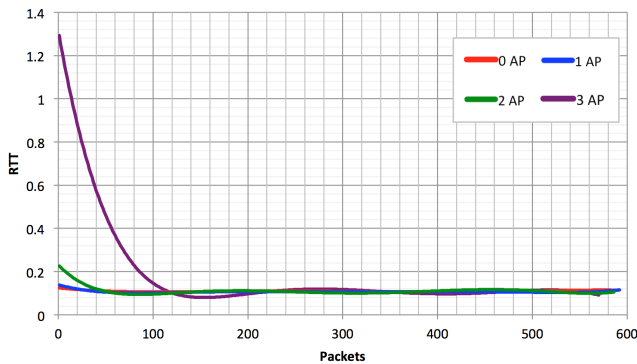


Fig. 4. Round Trip Time for different interference scenarios

V. CONCLUSION AND FUTURE WORK

In this paper we tested the robustness of Bluetooth Low Energy in a co-existence scenario with IEEE 802.11g. We established a static testbed to be used to generate interferences on the 2.4 GHz ISM band and proceeded to analyse the effects of such on a BLE packet exchange between two mobile devices.

We measured performance in terms of packet rate and round-trip time for different scenarios with varying interference levels.

We conclude by saying that BLE seems to be very resilient and although there was no packet loss in any of our scenarios a higher interference seems to impact slightly its performance. From our previous tests [1] we further proved how BLE weak spots consist in the initial pairing. The association time and first few packets following are highly affected by our worse interference scenario and can further reduce the likelihood of BLE being used for delay sensitive vehicular scenarios.

As future work, our interference testbed can be improved by observing the impact of a full coverage of all BLE channels (including the advertising ones) either by adding another pair of Raspberry Pis and adapting channels or by switching to IEEE 802.11n in 40Hz mode although this can prove to be challenging.

We are also planning on running all future test in a vehicular context with the help of two USRPs (Universal Software Radio Peripheral) that can provide us with continuous frequency coverage throughout the whole spectrum.

REFERENCES

[1] W. Bronzi, R. Frank, G. Castignani, and T. Engel, "Bluetooth low energy for inter-vehicular communications," in *Vehicular Networking Conference (VNC), 2014 IEEE*, Dec 2014, pp. 215–221.

[2] Bluetooth Special Interest Group, "Bluetooth Specification Version 4.0," URL: https://www.bluetooth.org/docman/handlers/downloaddoc.ashx?doc_id=229737.

[3] C. Gomez, J. Oller, and J. Paradells, "Overview and evaluation of bluetooth low energy: An emerging low-power wireless technology," *Sensors*, vol. 12, no. 9, pp. 11 734–11 753, 2012. [Online]. Available: <http://www.mdpi.com/1424-8220/12/9/11734>

[4] C. Gomez, I. Demirkol, and J. Paradells, "Modeling the maximum throughput of bluetooth low energy in an error-prone link," *Communications Letters, IEEE*, vol. 15, no. 11, pp. 1187–1189, November 2011.

[5] Sue White, "Wi-Fi and Bluetooth coexistence," URL: <http://www.ecnmag.com/articles/2012/03/wi-fi-and-bluetooth-coexistence>.

[6] Jeff Paynter, "Can Bluetooth and 802.11b/g/n Wi-Fi Devices Coexist?" URL: <http://intelligenthospitaltoday.com/can-bluetooth-and-802-11bgn-wi-fi-devices-coexist/>.

[7] Phil Smith, "Comparing Low-Power Wireless Technologies," URL: <http://www.digikey.com/en/articles/techzone/2011/aug/comparing-low-power-wireless-technologies>.

[8] Raajit Lall, "a se bouscule dans le spectre 2,4 GHz !" URL: <http://www.ni.com/newsletter/51685/fr/>.

[9] N. Golmie, O. Rebala, and N. Chevrollier, "Bluetooth adaptive frequency hopping and scheduling," in *Proceedings of the 2003 IEEE Conference on Military Communications - Volume II*, ser. MILCOM'03. Washington, DC, USA: IEEE Computer Society, 2003, pp. 1138–1142. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1950654.1950739>

[10] R. Frank, W. Bronzi, G. Castignani, and T. Engel, "Bluetooth low energy: An alternative technology for vanet applications," in *Wireless On-demand Network Systems and Services (WONS), 2014 11th Annual Conference on*, April 2014, pp. 104–107.

Improving Multi-Channel Beaconing in Vehicular Networks

Florian Klingler

Distributed Embedded Systems Group,
Dept. of Computer Science, University of Paderborn, Germany
klingler@ccs-labs.org

I. INTRODUCTION

One-hop broadcasts, termed beacons, are nowadays the main communication primitive for a wide range of Inter-Vehicle Communication (IVC) applications. They have been standardized as Cooperative Awareness Messages (CAMs) and Basic Safety Messages (BSMs).

Adapting the beacon interval has been identified as the most critical parameter to allow CAMs/BSMs to be exchanged in all possible scenarios, e.g., traffic jams with hundreds of cars within communication range or very sparse scenarios. The main reason is to not overload the wireless channel and thus avoid packet collisions while at the same time minimizing the communication delay. Presented adaptive beaconing concepts rely on a single wireless channel, and thus have the limiting factor of channel capacity.

This is in contrast to current standardization, which reserves multiple wireless channels for vehicular networking. Initially, seven channels were allocated in the U.S., later five channels in Europe as well – and there is a clear trend towards the availability of even more channels: More recently, the European ITS standard ETSI ITS-G5 moved to define up to seven channels, with optional use of IEEE channel 94 to 145 in traditional WiFi bands as well [1].

We study the feasibility of multi-channel beaconing and show how this improves message dissemination performance. Our approach builds on our previous work presented in [2], adding a novel concept for channel scheduling for different message priorities. First results show that the use of multiple channels leads to substantial performance improvements, while at the same time lowering the channel utilization per individual channel.

II. RELATED WORK

SOTIS [3] pioneered the exchange of information for traffic efficiency applications: knowledge bases (one being maintained on each vehicle) integrate received traffic information items; a subset of a vehicle’s local knowledge base is periodically assembled into beacons and broadcast to neighboring vehicles. Yet, as discussed earlier, it was found that static periodic beaconing is not suitable for every road traffic scenario.

To the best of our knowledge, REACT [4] is the first protocol which proposed a dynamic beaconing approach. The interval between two consecutive beacons is adapted according to the density of the road network.

Adaptive Traffic Beacon (ATB) [5] extends this approach. It proposes a novel prioritization scheme. Its overarching goal is to exchange as much information as possible, but avoid overloading the wireless channel at any time. Each knowledge base entry includes a priority based on the entries’ information, and the beacon interval is based on this priority and the channel quality. In [6], [7] similar concepts have been investigated, as well as in the ETSI ITS-G5 standardization group [8].

The IEEE 1609 DSRC/WAVE series of standards [9] describes how to operate a single-radio multi-channel system using a dedicated Control Channel (CCH), but leaves scheduling decisions to applications. Our work fills this gap and proposes channel scheduling algorithms to provide multi-channel operation for IVC.

III. MULTI CHANNEL BEACONING

We are working on a multi-channel beaconing extension to ATB, which is specifically designed to take advantage of the additional Service Channels (SCHs) available in the DSRC band. We evaluated our multi-channel approach in a Single-Radio Multi-Channel (SR-MC) split phase scenario, but the presented concept can easily be extended to Multi-Radio Multi-Channel (MR-MC) environments without using split phase channel switching.

As in WAVE, time is divided into CCH and SCH intervals, each with a duration of 50 ms and having a small guard interval in front to minimize the probability of lost messages during channel switching. WAVE follows the principle to broadcast data announcements on the CCH, advertising that SCH where data will be transmitted during the following SCH interval. Channel switching is only performed during guard intervals.

Our protocol’s operation is divided into four distinct steps: First, we regulate the beaconing rate by adapting the number of intervals to elapse before sending a data announcement. Second, when we selected an interval, we carefully determine the time within the interval to broadcast the announcement. This time t within the CCH interval is based on the priority of the payload information such that more important messages are sent earlier in the interval. Third, at t , a node selects the SCH to transmit the payload information by taking into account all received announcements up until this point, and sends the announcement. In a fourth step, during the guard interval the node tunes its radio to the announced SCH and broadcasts the data in that way to avoid synchronized collisions.




Fig. 1. Channel utilization and fraction of informed vehicles for a medium density freeway scenario.

IV. FIRST RESULTS

We show the performance of our multi-channel approach by using a medium utilized freeway scenario having two lanes in each direction, consisting of 90% cars and 10% trucks. Each vehicle periodically generates low priority dummy messages and fills its local knowledge base with it. After protocol execution has reached a steady state, we select a random vehicle in the middle of the freeway to generate a high priority message. We will track its dissemination in our evaluation. Each simulation is repeatedly executed for different random number seeds to get good confidence in the results. Data is recorded within a region of interest of 1 km to minimize border effects, and only after a steady state of the protocol has been reached. We compare our multi-channel approach to the baseline single-channel ATB protocol and focus on two different metrics, namely channel utilization as low level performance, and relative message dissemination speed to measure application level performance.

To investigate channel conditions we select the channel utilization experienced by each individual vehicle. This metric is calculated as the fraction of simulation time for which physical Clear Channel Assessment (CCA) of that vehicle would have considered the channel busy. Figure 1a shows the results split by channel. Both beaconing schemes, single-channel and multi-channel, keep the channel utilization at a very low level, following their aim of not overloading the channel. In particular, the CCH is lower utilized by the multi-channel protocol, because it uses the channel only for much shorter announcement beacons. This means that the multi-channel variant even would be able to send substantially more frames on the CCH than its single-channel counterpart. Payload transmissions across all SCHs can also be seen to be evenly distributed.

Looking at application layer performance, the second metric we select is the fraction of informed vehicles. We track how fast a single piece of information spreads through the network, by generating such a high priority item in the middle of a highway and feeding it to the vehicle's local knowledge base like described before. For each time step in the simulation we then track the fraction of all vehicles that already received this particular piece of information. The results are shown in Figure 1b, where we plot the mean fraction of informed vehicles in all simulation repetitions, normalized to $t = 0$. This metric is influenced by all previous mentioned factors, e.g., channel utilization and packet collisions. As can be seen, the multi-channel variant of the protocol is able to propagate the information through the network substantially faster than the single-channel variant. The results are even more interesting, since the multi-channel approach can only use 46 ms of each 50 ms channel interval to send and receive beacons, which is caused by the 4 ms guard interval in front of each slot.

V. CONCLUSION AND FUTURE WORK

We presented an improved version of our previous work for multi-channel beaconing, which is able to lower the channel utilization and thus observed packet collisions while at the same time increasing the relative message dissemination speed. First simulation results performed using a single-radio split phase multi-channel approach show the feasibility of multi-channel beaconing in vehicular networks. Since our approach is not limited to a single-radio system, we will extend our proposed protocol to work towards a multi-radio system using all available channel space specified by ETSI ITS-G5.

REFERENCES

- [1] European Telecommunications Standards Institute, "Intelligent Transport Systems (ITS); Access layer specification for Intelligent Transport Systems operating in the 5 GHz frequency band," ETSI, EN 302 663 V1.2.1, July 2013.
- [2] F. Klingler, F. Dressler, J. Cao, and C. Sommer, "Use Both Lanes: Multi-Channel Beaconing for Message Dissemination in Vehicular Networks," in *10th IEEE/IFIP Conference on Wireless On demand Network Systems and Services (WONS 2013)*. Banff, Canada: IEEE, March 2013, pp. 162–169.
- [3] L. Wischhof, A. Ebner, H. Rohling, M. Lott, and R. Halfmann, "SOTIS - A Self-Organizing Traffic Information System," in *57th IEEE Vehicular Technology Conference (VTC2003-Spring)*, Jeju, South Korea, April 2003.
- [4] E. Van de Velde and C. Blondia, "Adaptive REACT protocol for Emergency Applications in Vehicular Networks," in *32nd IEEE Conference on Local Computer Networks (LCN 2007)*. Dublin, Ireland: IEEE, October 2007, pp. 613–619.
- [5] C. Sommer, O. K. Tonguz, and F. Dressler, "Traffic Information Systems: Efficient Message Dissemination via Adaptive Beaconing," *IEEE Communications Magazine*, vol. 49, no. 5, pp. 173–179, May 2011.
- [6] R. K. Schmidt, T. Leinmüller, E. Schoch, F. Kargl, and G. Schäfer, "Exploration of adaptive beaconing for efficient intervehicle safety communication," *IEEE Network Magazine*, vol. 24, no. 1, pp. 14–19, January 2010.
- [7] M. van Eenennaam, W. Wolterink, G. Karagiannis, and G. Heijenk, "Exploring the solution space of beaconing in VANETs," in *1st IEEE Vehicular Networking Conference (VNC 2009)*. Tokyo, Japan: IEEE, October 2009.
- [8] European Telecommunications Standards Institute, "Intelligent Transport Systems (ITS); DCC Mechanisms for ITS operating in the 5 GHz range; Access layer part," ETSI, TS 102 687 V1.1.1, July 2011.
- [9] IEEE, "IEEE Standard for Wireless Access in Vehicular Environments (WAVE) - Multi-channel Operation," IEEE, Std 1609.4, February 2011.

Robust Usage of V2X Data in Autonomous Driving

Sebastian Mauthofer

Fraunhofer Institute for Secure Information Technology

Darmstadt, Germany

Email: sebastian.mauthofer@sit.fraunhofer.de

Abstract—The local sensors used for decision-making in autonomous driving are usually restricted to the line of sight. In principle, direct ad-hoc vehicular communication (V2X) could be used to virtually extend the field of vision. However, the security of V2X systems is primarily designed for warning messages where a driver is the ultimate decision-making authority. For autonomous driving, a source of information needs to be highly trustworthy or at least confirmable by another independent source. In this paper a method is proposed that allows an autonomous vehicle to benefit from V2X messages. By feeding V2X position data directly to local sensor systems like radar, lidar or video cameras, the time delay between the first visual contact with another vehicle and the confirmation of existence could be reduced and the tracking improved. This is achieved by optimizing the search patterns of the sensors and by preparing and calibrating the tracking algorithms. As autonomous decisions strongly depend on reliable real-time data, the proposed method might help to reduce the number and severity of accidents.

I. INTRODUCTION

Direct ad-hoc vehicular communication (V2X) and autonomous driving are two recent trends in the automotive industry aiming to increase the road safety and to optimize the traffic flow. While vehicular communication is designed as a supporting aid for drivers, the goal of autonomous driving is a driverless operated vehicle. V2X messages are sent between vehicles among themselves and between vehicles and road infrastructure. They are used to inform drivers about road conditions, approaching vehicles, traffic light phases or events like emergency braking. In case of a potentially dangerous situation, the drivers have to decide how to act. One of the greatest advantages of this technology is, that a driver can get real-time information about events outside his line of sight. For example, a driver who receives a warning that another vehicle is approaching an upcoming intersection with high speed and might violate his right of way can slow down his vehicle as a precaution, before he is actually able to see the other vehicle.

In contrast, vehicles equipped with autonomous driving systems have no driver as an ultimate decision-making authority. Therefore, autonomous vehicles depend on comprehensive knowledge about their surroundings. Using global data like street maps and real-time traffic information on the one hand and local sensors like radar, lidar or video cameras combined with highly accurate position data on the other hand, they perform autonomous driving decisions. The local sensors used for the major part of the environmental perception are usually restricted to the line of sight. This limits the basis of decision-making. In case of a right of way violation as described in the previous example, an autonomous vehicle can only react when the other vehicle comes into sight.

Even though V2X could be theoretically used to virtually extend the field of vision of autonomous vehicles, for now, this cannot be seen as granted.

One barrier to an integration of vehicular communication in autonomous driving is the trustworthiness of V2X data. The data is not generated by local sensors that the automotive manufacturers are responsible for and that they can secure and test thoroughly. Instead, it originates from unknown vehicles that could be less reliable or even manipulated. Furthermore, the security of V2X systems is primarily designed for warning messages where a human driver checks the plausibility and decides on suitable driving maneuvers.

This might not be sufficient for autonomous vehicles that require highly accurate and up-to-date information. Each source of information needs to be highly trusted or at least confirmable by another independent source.

As the safety and life of occupants depends on correct decisions, there are mainly two options to benefit from V2X data in autonomous driving: Either increase the level of V2X security, or use the data in a robust way, where the decision making vehicle would profit from correct messages and keep negative effects at a minimum, in case of wrong messages. In this paper a method is proposed to benefit from V2X data as an additional source of information in autonomous driving, largely independent from the trustworthiness of the sender.

mds

February 22, 2015

II. TECHNICAL BACKGROUND

Local vehicular sensors used to identify and track vehicles and other objects on the road are usually relying on Bayesian filters, like Kalman filters or particle filters, for tracking purposes. These filters are using a prediction and update cycle to keep track of moving objects. The more complete the description of a system or an object, in this case the information about a tracked vehicle, the easier it is to derive useful feed for the prediction of the next position.

The initial detection of an object (before it is recognized and tracked) takes the longest time, as a sensor has no previous knowledge about the existence and features of an object. First of all, the potential object needs to be physically detected by a sensor and confirmed using further sensor readings. Then it will be tracked using the previously mentioned Bayesian filters and analyzed if it is relevant or maybe a static element of the road side. This results in some delay between the first visual contact with an object and the time the object is recognized.

Modern radar systems are usually using beam forming technology. Depending on the type of system, the initial object detection time can be reduced if the position of an object is already known. The search pattern of the radar can be adjusted and the assumed object position can be directly targeted with a beam. Similar, in case of a video camera system, the initial detection of an object can be optimized by selecting specialized algorithms and focusing the processing power on the relevant section of the image.

III. PROPOSED SOLUTION

The proposed solution makes use of both effects and is intended for potentially untrustworthy V2X position data in the context of autonomous driving. In many cases, a V2X equipped vehicle is detected based on its beacon messages before it is detected by local sensors like radar or video cameras. This can be exploited as follows, to enhance the detection capabilities of the local sensors.

First of all, a central controller has to know the areas covered by the local sensors. When a new vehicle is detected by the V2X system its position should be monitored until it gets close to the detection area of a sensor. In parallel, the plausibility of the V2X position trace should be checked. If there is any doubt on the plausibility of the V2X-tracked vehicle or if the quality and update frequency of the V2X messages is not sufficient, the vehicle is ignored. Else, the corresponding local sensor is informed about the position, speed and heading of the potential vehicle and in case of sufficient data also about the predicted next position. As the position of the V2X-tracked vehicle is not yet within the detection area of the sensor, the sensor should not know and track this vehicle yet.

Using the given information about the expected vehicle, the sensor will start a new focused detection procedure. In case of a radar system, the vehicle could either modify the beam forming search pattern to increase the frequency at which the position of the expected vehicle is observed or optimize the beam itself with respect to shape, signal or power. In case of a video camera, the processing power can be focused on the relevant image section and optimized pattern search algorithms can be used.

After the vehicle is detected, the V2X data can be used to calibrate the Bayesian filter used for tracking. For sensor-confirmed vehicles the frequent V2X message updates could also be used for long-term tracking. When a tracked vehicle is shadowed by another vehicle or object and temporarily not visible for the sensor, the V2X data can be used to virtually continue the tracking and resume the sensor-based tracking as soon as the vehicle is visible again. That way, object specific tracking parameters would not be lost.

IV. CONCLUSION AND FUTURE WORK

In this paper a method is proposed to use V2X data in autonomous driving in a robust way. By feeding local sensors with V2X position data the autonomous decision-making process can benefit from the trustworthiness of the local sensors and the extended line of sight of the vehicular communication. In case of correct V2X messages the time delay between the first visual contact with another vehicle and

the confirmation of existence could be reduced. Additionally, the long-time tracking of vehicles would be supported and the tracking of vehicles could be maintained even if they would be temporarily shadowed by other vehicles or objects. In case of wrong messages there would only be a minimal negative effect on the overall system performance. As autonomous decisions strongly depend on real-time data, the proposed method might help to reduce the number and severity of accidents.

The next steps will be, after analyzing and comparing related approaches, to identify compatible radar and video camera systems and to estimate the potential reduction of the detection delay. Depending on the results, the focus of the research can be on the initial detection of vehicles or the following tracking optimization. Furthermore, advanced plausibility checks could be added in order to detect malicious or unusable V2X messages before they can disturb the local sensors.

Investigating a Reliable Inter-vehicle Network in a Three Dimensional Environment

Lisa Kristiana, Corinna Schmitt, Burkhard Stiller

Department of Informatics IFI, Communication Systems Group CSG, University of Zurich

Binzmühlestrasse 14, CH—8050 Zurich, Switzerland

[kristiana|schmitt|stiller]@ifi.uzh.ch

Abstract—Today inter-vehicle network offers a trend of mobile connectivity. It is inspired by the principle of Vehicular-to-Vehicular (V2V) network which tends to take into account of vehicles as mobile communication nodes. A major request of most end-users even when they are commuting with public transportation is Internet connection such as social media, messaging services, and news updates. The main challenge of inter-vehicle network is to provide a reliable connectivity in a large city environment which has particular characteristics. One of characteristics is the existence of overpass and tunnel which influences its reliability and stability in inter-vehicle network. Vehicles are distributed in three dimensional area, which can be under or on overpasses. The proposed forwarding method, termed Vehicle-to-Vehicle Urban Network (V2VUNet) filters out neighbor nodes with potential loss transmission. As this investigation currently is ongoing work this paper provides a brief overview of three dimensional challenges and evaluates existing forwarding method where a propagation and position-based routing are considered, which influences the design.

Index Terms—Vehicle-to-Vehicle Urban Network (V2VUNet), Vehicular Ad-hoc Network (VANET), Vehicular-to-Vehicular (V2V) networks, forwarding method, three-dimensional environment, position-based routing protocol

I. INTRODUCTION

The population in large cities rises almost faster than the development of public transportation and mobility needs of the population. A major problem of commuting with public transportation appears when traffic jams occur in busy hour. In turn, people spend too much time to commute, which leads to inconvenient trips. The Smartphone utilization (*e.g.*, for social media access, messaging services, or news updates) seem to determine a convenient activity to at least partially compensate these externally imposed conditions [2], mainly due to the fact that more than 80% of the people own a variety of handheld devices (*e.g.*, Smartphones or tablets) [4].

By observing people's Smartphone behavior, the basic idea of offering a wireless network access passenger-related applications within public transportation (*i.e.* busses or trains) can help to turn inconvenient road traffic conditions, such as congestion and long commuting time, into pleasure or even productive time. During this commute time passengers inside public transportations can check e-mail, surf the Internet, or perform social network activities [2]. In practical terms, the equipping of public transportation with Internet access defines a convenient solution for passengers.

Vehicular-to-Vehicular (V2V) networks build a subgroup of Vehicular Ad-hoc Networks (VANET) [4][7] and inspires the proposed inter-vehicle communication development Vehicle-to-Vehicle Urban Network (V2VUNet). V2V networks

today offer several applications, which can be classified based on their purposes into [5]:

(1) *Driving-related applications* assist a driver to reach traffic safety and a partial efficiency by aggregating road-related information and presenting it to the driver.

(2) *Passenger-related applications* emphasize on convenience and comfort of passengers on board.

(3) *Vehicle-related applications* do not have impact on the driver and passengers directly, they improve all operations of vehicles and their internal optimization.

A main challenge of implementing passenger-related applications (*i.e.* infotainment, social media access, or surfing the Web) is reliable and stable connectivity. Therefore, in V2V Network, the coverage issue is caused by three dimensional road topology which leads to frequently loss connection. While in a large city environment, the existence of overpass and bridges in many areas which leads to frequently loss connection. Vehicles in a real large city environment are distributed in three-dimensional area as may take place under or on the overpass as shown in Figure 1.




Fig. 1. The Architecture of V2VUNet in a Large City Scenario

This paper's focus is on the investigation of forwarding method based on greedy principle where three-dimensional road topology is considered. Assume that there are two layers of road, the upper road layer represents an overpass or bridge and the lower road layer represent regular road under overpass or bridge. Busses are distributed on both road layers and establish connectivity among them. Busses' connectivity has to deal with a hierarchy road topology *i.e.*, three-dimensional environment.

The remainder of this paper is structured as follows: Section II describes challenges in three-dimensional environment

including the theoretical background. While Section III presents implementation ideas of new forwarding method. Section IV discusses preliminary investigation results. Finally, Section V concludes and offers an insight into future steps.

II. CHALLENGES IN THREE DIMENSIONAL ENVIRONMENTS

In a large city environment, three-dimensional road topology (e.g., overpass, bridge, and tunnel) [11], buildings and dense road traffic are more available than in country side. This road topology hierarchy influences connectivity due to several reasons:

(a) Loss connection occurs because busses move into a tunnel, is considered as forwarding method issues

(b) The opportunity to connect with other busses is very rare, because of signal strength, is considered as propagation issue

(c) Even when it is connected, the duration is very short because high mobility of busses, is considered as routing protocol issues.

Those reasons lead to throughput of intervehicle communication. The literatures as following discuss about the problems mentioned in detail.

A. Forwarding Method Issues

Several three-dimensional forwarding methods based on Greedy forwarding method, have been developing to adapt the potential issues in inter-vehicle network. Greedy forwarding method has a fundamental to cope the three-dimensional environment [8]. Several variant of greedy forwarding method introduce a distance as the parameter. Those variant are Most Forwarding Progress within Radius (MFR), Nearest Forwarding Progress (NFP), and Compass routing. MFR choose the forwarder node which closest neighbor to destination [23]. NFP looks into intermediate node which closest to source [15]. Moreover, Compass forwarding method, which considers a minimum angle of intermediate node as a priority parameter [16]. Therefore, those variant of greedy forwarding method inspires the proposed forwarding method V2VUNet in order to search the potential intermediate nodes with optimal forwarding progress.

B. Propagation Issues

Propagation issues of wireless devices in large city environment become interesting topic, because there is, as far as the authors are aware of, no propagation model appropriate to the requirement at the moment. In a large city environment, there are several factors that must be considered as obstruction influencing propagation: building, road topology, and traffic condition [13][14]. Figure 2 illustrates a case where bus B moves into a tunnel, losing connectivity, and thus cannot reach bus A located on the upper road level anymore.

C. Routing Protocol Issues

Traditional routing protocols frequently used in Mobile Ad-hoc Network (MANET) (e.g., Dynamic Source Routing (DSR) and Ad-hoc On Demand Vector (AODV)) are not suitable for inter-vehicular network due to high mobility characteristic [17][18][19]. Interesting option for proposed solution are position-based routing protocols, because they are equipped with location information [18][22]. The location information provides location coordinates of a mobile nodes as x , y , and z . These coordinates are provided by Global Positioning System (GPS) and represented as longitude, latitude, and altitude. In three-dimensional environment, the altitude is

set as $z \neq 0$, to indicate that busses which have the value of $z \neq 0$, move on the upper road.

The classic position-based routing protocol for V2V network is Greedy Perimeter Stateless Routing (GPSR). GPSR implements perimeter mode to cope the void area problem [8][20]. Additionally, many improvements have been made as required in VANET characteristics [21]. Without adaptation and improvement, the position-based routing protocol have poor performances in three-dimensional environment since they are not designed to cope with this environment [11][12]. However, GPSR suitable to be used for network with dynamic topology

III. VEHICLE-TO-VEHICLE URBAN NETWORK

The proposed forwarding method, Vehicle-to-Vehicle Urban Network (V2VUNet) looks into the Greedy forwarding method with a modification addressing: (1) Three-dimensional Propagation Mode and (2) Intermediate Node Filtering.

Three-dimensional Propagation Mode.

At this first step, the propagation is set due to the hierarchical road topology as shown in Figure 2. Bus A is located on the upper road and bus B on the lower road.




Fig. 2. Scenario of Three-dimensional Environment

Transmission radius between two busses denoted as l , m , and n . Variable l represents the distance where bus B is out of reach by bus A, and m represents the distance where bus B has the optimal transmission range with bus A while n represents transmission with attenuation, because bus B will move into the tunnel and signal reception will be very weak or even lost. The l , m , and n radius are defined as follows assuming transmission radius of IEEE 802.11p typically is 1000m in ideal condition: Then $l > 1000m$, n is assumed 0 - 20m and m is 21m - 1000m. Variable n depends on the speed of bus. For instance if transmission radius is 1000m, speed of bus is 20m/s then the maximum connection duration is 5s. Therefore, if bus B reaches the bus A's transmission distance of 20m, with speed of bus B is 20m/s yields 0s connection duration. Hence, with the speed of 20m/s, n is set to 20m.

Intermediate Node Filtering.

At this stage, when bus B moves within transmission range of bus A, filtering mechanism is applied. Bus A determines the position of bus B and starts to transmit packet when bus B is located within m radius. Once bus B moves within n radius, bus A disconnects the transmission and looks into another bus within m radius. The proposed solution, V2VUNet based on the aforementioned given parameter is simulated using NS-3 [25] and first results are presented in the upcoming section.

IV. SIMULATION RESULTS

As the first step, the road hierarchical topology is set as shown in Figure 2. An overpass is located in the center of road. This overpass can be assumed as a tunnel, thus busses can move into it. Busses are placed randomly on both road layer as the initiation. Busses' movement model is defined as the bus trajectory in which busses move from one location to another and then return to the previous location. Moreover, the random velocity is also applied to generate uneven bus distribution as in real road traffic.

The second step, busses are equipped with GPS to get location information and interchange information. In this scenario, busses run with speed up to 60km/h, since it is considered as common case in large city roads. To evaluate V2VUNet, scenario is set as shown in Table 1. V2VUNet is compared with non-modified forwarding method, which does not implement intermediate node filtering. At this stage a position-based routing, GPSR is implemented in three-dimensional area. GPSR applies location service and greedy forwarding method [6]. Hence, GPSR is suitable to implement in this case.

TABLE I SIMULATION PARAMETERS

Parameter	
Transmission Range	1000m
Routing Protocols	GPSR, V2VUNet
Number of Nodes	20 - 50
Simulation Area	2000m x 2000m
Upper Road Height	10 m
Busses Velocity	0 -60km/h
Transfer rate:	2048 bps
Packet Size	64 bits
Simulation Time	400s
Number of lane	1

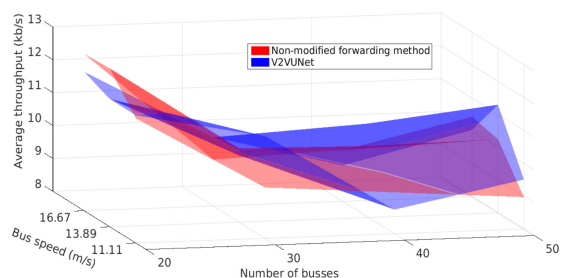


Fig. 3. Throughput vs Number of Busses and Speed of Busses

The preliminary result shows the throughput of V2VUNet is higher than non-modified forwarding method. The term of non-modified forwarding method is used to indicate that only GPSR protocol is implemented as a comparison of the proposed method at the moment. The throughput is indicated as a multi-hop performance since it is measured from source to destination or end-to-end.

Although a small improvement showed, it indicates that V2VUNet perform better in three-dimensional environment in some cases. In case of low speed, the throughput of V2VUNet is higher than GPSR. The reason is that the connectivity

among busses remain stable. However, at high speed, GPSR shows a better result than V2VUNet. On the other case, when a number of busses are added, V2VUNet performs better than GPSR both in low and high speed.

However, performance metric shows that V2VUNet yields a better result when number of busses and speed of busses increase. This means that V2VUNet performs better than the non-modified forwarding method, moreover, it is suitable for three-dimensional environment in case of high traffic density.

V. SUMMARY AND FUTURE WORK

To summarize, the position-based routing protocol requires modifications and adaptations to road topology. In case of three-dimensional area, the transmission range is influenced by obstruction (*i.e.*, tunnel). Moreover, in a large city environment where the uneven distributed vehicles occurs, another issue needs to be concerned is a sparse environment where there is no bus within a certain range. This condition in which exchanging data packet becomes impractical, therefore, requires a particular modification and adaptation. As for the sparse environment, the carry-to-forward approach can be very promising to obtain the effective connectivity.

For future work, other type of vehicles such as private cars, taxis, and motorcycles are added. These vehicles can form 'obstacles' causing traffic jam and reducing speed of busses. Since busses cannot move swiftly, which leads to loss transmission range among busses.

Unlike in large cities of developed countries (*e.g.* Jakarta, Indonesia), public busses tend to behave randomly both in timewise and frequency-wise: Busses do not even have schedules and official stops. Thus, it is difficult to predict the pattern of busses' arrival and, therefore, a random arrival model is applied [1].

Another aspect that frequently happens related to forwarding decision is the optimal area based on the angle between two mobile nodes. This will also be implemented. More complex and specific position-based routing algorithms will further be tested.

VI. REFERENCES

- [1] A. S. Alfa, M. F. Neuts: *Modelling Vehicular Traffic Using the Discrete Time Markovian Arrival Process*, Transportation Science, Vol. 29, No. 2, pp. 109–117, May 1995
- [2] K. Barry: *How Smartphones Can Improve Public Transit*, Wired, <http://www.wired.com/2011/04/how-smartphones-can-improve-public-transit/>, last access October 2014
- [3] T. Kosch, C. Schroth, M. Strassberger, M. Bechler: *Application Classification and Requirement*, Automotive Internetworking, John Wiley & Sons Ltd, pp 13–36, February 2012.
- [4] R. Meraihi, S. M. Senouci, D. E. Meddour, M. Jerbi: *Vehicle-to-Vehicle Communications: Application and Perspectives*, Wireless Ad Hoc and Sensor Networks, pp 285-308, Januari 2010
- [5] M. Mauve and B. Scheuermann: *VANET Convenience and Efficiency Applications*, VANET Vehicular Applications and Inter-Networking Technologies, John Wiley & Sons Ltd, pp 81–105, NOV 2009
- [6] B. Paul, M. Ibrahim, M. Bikas, and A. Naser: *Experimental Analysis of AODV & DSR over TCP & CBR Connections With Varying Speed And Node Density in VANET*, Vol. 24, No. 4, pp. 30–37, 2012.
- [7] S. Yousefi, M.-S. Mousavi, M. Fathy: *Vehicular Ad Hoc Networks (VANETs): Challenges and Perspectives*, 6th International Conference on ITS Telecommunications, Chengdu, China, pp 761-766, June 2006
- [8] B. Karp and H. T. Kung: *GPSR : Greedy Perimeter Stateless Routing for Wireless Networks*, International Conference on Mobile Computing and Networking, pp 243–254, 2000

- [9] K. Kastaros, M. Dianati, R. Tafazolli: *CLWPR - A Novel Cross-layer Optimized Position-based Routing Protocol for VANETs*, IEEE Vehicular Networking Conference, (VNC 2011), Amsterdam, Netherland, pp 139-146, November 2011.
- [10] K. Kastaros, M. Dianati, K. Roscher: *A Position-based Routing Module for Simulation of VANETs in NS-3*, Proceedings of the 5th International Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering (ICST) Conference on Simulation Tools and Techniques, (SIMUTOOLS 2012), pp 345-352, Desenzano, Italy, March 2012
- [11] Q. Lin, C. Li, X. Wang, L. Zhu: *A Three-Dimensional Scenario Oriented Routing Protocol in Vehicular Ad Hoc Networks*, Vehicular Technology Conference, (VTC Spring 2013), Dresden, pp. 1-5, June 2013.
- [12] Y. He, C. Li, X. Han, Q. Lin: *A Link State Aware Hierarchical Road Routing Protocol for 3D Scenario in VANETs*, Internet of Vehicles – Technologies and Services, Lecture Notes in Computer Science, (LNCS), Vol. 8662, pp 11-20, 2014
- [13] J. Walfisch, H. Bertoni : *Theoretical Model of UHF Propagation in Urban Environments*, IEEE Transactions on Antennas and Propagation, Vol. 36, No. 12, pp 1788-1796, December 1988.
- [14] Y. K. Yoon, J. H. Kim, M. W. Jung, Y. J. Chong: *Radio Propagation Characteristics in the Large City*, 16th International Conference on Advanced Communication Technology, (ICACT), Pyeongchang, South Korea, Februari 2014
- [15] T. C. Hou, V. O. K. Li, *Transmission Range Control in Multihop Packet Radio Networks*, IEEE Transaction on Communications, Vol. 34, No. 1, pp 38 - 44, January 1986
- [16] E. Kranakis, H. Singh, J. Urrutia: *Compass Routing on Geometric Networks*, 11 Th Canadian Conference On Computational Geometry (CCCG 1999), Vancouver, British Columbia, pp 51-54, 1999
- [17] C. Suthaputthakun, Z. Sun: *Routing Protocol in Intervehicle Communication Systems: A Survey*, IEEE Communications Magazine, Vol. 49, No. 12, pp 150 - 156, December 2011
- [18] F. Li, Y. Wang: *Routing in Vehicular Ad Hoc Networks : A Survey*, IEEE Vehicular Technology Magazine, Vol. 2, No. 2, pp 12 - 22, June 2007
- [19] T. L. Willke, P. Tientrakool, N. F. Maxemchuk: *A Survey of Inter-Vehicle Communication Protocols and Their Applications*, IEEE Communications Surveys & Tutorials, Vol. 11, No. 2, pp 3-20, June 2009
- [20] R. Kumar, M. Dave: *A Comparative Study of Various Routing Protocols in VANET*, International Journal of Computer Science Issues, (IJCSI), Vol. 8, No. 4, pp 643-648, July 2011
- [21] M. Jerbi, S. M. Senouci, R. Meraihi, Y. Ghamri-Doudane: *An Improved Vehicular Ad Hoc Routing Protocol for City Environments*, IEEE International Conference on Communications, (ICC 2007), Glasgow, Scotland, pp 3972 - 3979, June 2007
- [22] H. Ghafoor, K. Aziz: *Position-based and Geocast Routing Protocols in VANETs*, 7th International Conference on Emerging Technologies, (ICET), Islamabad, Pakistan, pp 1 - 5, September 2011
- [23] M. Abdoos, K. Faez, M. Sabaei: *Position Based Routing Protocol With More Reliability In Mobile Ad Hoc Network*, First Asian Himalayas International Conference on Internet, (AH-ICI), Kathmandu, Nepal, pp 1 - 4, November 2009.
- [24] C. Han, M. Dianati, R. Tafazolli, R. Kernchen, X. Shen: *Analytical Study of the IEEE 802.11p MAC Sub-layer in Vehicular Networks*, IEEE Transactions on Intelligent Transportation Systems, Vol.13, No. 2, pp 873 - 886, February 2012
- [25] Network Simulator (NS3) manual, <http://www.nsnam.org/docs/release/manual/html/index.html>, last access October 2014

Approaches to Heterogeneous Vehicular Networks

Florian Hagenauer

Distributed Embedded Systems Group, Department of Computer Science, University of Paderborn, Germany
hagenauer@ccs-labs.org

I. INTRODUCTION

IEEE 802.11p is currently the base for the major technologies used in the field of vehicular networking. This holds true for the IEEE 1609.4 standard family used in the U.S. (called IEEE WAVE) and for the European ETSI ITS-G5 standard. But when it comes to such WLAN based technologies problems might arise: First, it might be possible that IEEE 802.11p fails to cover dense urban areas due to too many parallel users. Second, the penetration rate of vehicles outfitted with the necessary equipment will be initially very low. With this in mind it might be an idea to use a cellular connection instead of using a WLAN link between cars. Indeed some car manufacturers are already equipping their cars with cellular technology, not just for voice calls, but also for value added services. But this also leads to various problems: If the used communication technology is based on LTE (which is currently the most advanced one in use) a high frequency of messages might overload the network [1]. Due to the involvement of a core network, naturally the delay will also increase. Finally, if LTE is used, other users might experience degraded network performance due to all the cars sending messages.

One solution to these problems are *heterogeneous vehicular networks* where a car is equipped with both a short range radio and a module for cellular communication. In 2005, Cavalcanti et al. [2] proposed a system which combines WLAN with cellular technologies. More recently this idea has been adapted to be used in vehicular networks. Among these approaches is Remy et al. [3] who proposed *LTE4V2X*. It uses WLAN to cluster cars which then exchange current positional data to allow for easier travel planning. Tung et al. [4] proposed a safety application using heterogeneous vehicular networks for collision avoidance. Again cars are managed in clusters; a central server, reachable via a cellular connection, warns clusters of incoming other clusters.

Currently simulation is the tool of choice for evaluating concepts in the area of vehicular networking. Such simulation frameworks have to take care of network simulation and of a realistic simulation of vehicles' mobility. Therefore, most of the time a dedicated network simulator is coupled with a mobility simulator. In such a scenario the simulators are able to react on each others events. Especially in cases of safety and efficiency applications this is a very important property. For example, simulating the traffic mobility with traces of real world traffic – without a feedback loop – makes it impossible for vehicles to change their routes dependent on information received via the network(s). Such information

could be about a traffic accident on the road ahead were the information was received via the network connection. There already exist frameworks which provide support for mobility-network feedback loops (*Veins*, *iTETRIS*, and *VSimRTI*), but all of them were missing support for full featured simulation of cellular networks. In December we released *Veins LTE*, an extension module for *Veins* which adds support for LTE [5]. For this, it uses *SimuLTE*, developed by Virdis et al. [6]. As *Veins* is based on OMNeT++, implementing, running, and evaluating new simulations is not a very complex task. This makes it also a good option for classroom use and teaching.

II. VEINS LTE

Veins LTE combines the features of *Veins* (network-mobility feedback loop [7], IEEE 802.11p based network stack) with *SimuLTE* for LTE support. The main purpose was to integrate LTE into the system to make it easy to develop new applications on top. At lot of parameters of these stacks can be changed without altering the source code which makes the whole framework accessible. On top of the two stacks multiple applications can be implemented and run. A special module performs protocol adaptation for the stacks and routes messages between them and the applications. Every application can annotate outbound messages to be sent via LTE only, via IEEE 802.11p only, or leave the decision to lower layers. For this, a *Decision Maker* module provides an interface for implementing useful algorithms to make smart choices – like using current channel conditions or the number of recently sent packages.

Integrating *SimuLTE* into *Veins* was not easy: *SimuLTE* relies on all nodes to be initialized at the start of the simulation. This is not useful for simulating vehicular networks as cars enter and leave the scenario throughout the simulation. Therefore, I had to add the capabilities of adding and deleting vehicles to the *SimuLTE* framework – from the Network Interface Controller (NIC) to the IP stack. The result is now a robust framework to simulate heterogeneous vehicular networks where users only have to implement the applications on top of the stacks. Afterwards, running and configuring the simulations to gather results is straightforward and well supported by OMNeT++.

An example screenshot of a running simulation can be seen in Figure 1. It shows cars exchanging messages via LTE, where the green lines represent uplink connections and the red ones downlink. The blue, dashed lines show cars communicating via Dedicated Short-Range Communication (DSRC) to form clusters. Finally, the red crosses show received warnings about other cars approaching.




Figure 1. A screenshot of a running *Veins LTE* scenario. The red crosses show warnings that cars receive of incoming other cars.

III. HETEROGENEOUS VEHICULAR NETWORKS

The development of *Veins LTE* is not just a *fire and forget* action, we released it as open source software¹ and use it actively in research. To evaluate it we performed a huge number of smoke tests. Such a validation allowed us to make sure that the integration of *SimuLTE* into *Veins* did not break any of the used components.

We already implemented one heterogeneous vehicular networking algorithm, namely the one proposed by Tung et al. [4], to further evaluate the simulation framework. The authors propose an algorithm for collision avoidance where cars use DSRC to exchange their current position and form clusters. The clusters are then sent via LTE to a central server which then is able to warn clusters about approaching vehicles. Beside collision avoidance, such data can be used to manage intersections or plan routes. To cluster the vehicles the authors used a decentralized clustering solution where vehicles are clustered in groups moving into the same direction. Therefore the vehicle closest to the next intersection becomes the cluster head – this is the vehicle which communicates via LTE with the central server. All vehicles on the same road and traveling towards the same intersection will be in the same cluster.

Remy et al. [3] take a completely different approach by letting a central server decide which cars are in which cluster. They argue that such a central server has a better overview of the overall situation and is able to generate better and longer living clusters. We are currently in the process of implementing this solution in *Veins LTE* as well.

As a third direction, we are actively using *Veins LTE* in our *Car4ICT* project.² Together with Toyota ITC we explore how cars can be used as a main Information and Communication Technology (ICT) resource in future smart cities.

In the literature, there are many more approaches how clusters of vehicles are created. The concept of [8], which was further improved in [9], focuses on providing a link to the Internet for cars without a cellular connection. This link is then shared with the clusters. To cluster vehicles they

use the direction of movement, the cellular signal strength, and the WLAN transmission range. Yet, how *best* to perform clustering is an open question. In addition, by running multiple applications which depend on clustering in parallel, the overhead for creating those clusters might get too big. To tackle this problem, our initial step is to compare the approaches from [3] and [4] to get a feeling if clustering benefits more from the advantages of a decentralized or from a centralized solution.

REFERENCES

- [1] C. Casetti, F. Dressler, M. Gerla, J. Gozalvez, J. Haerri, G. Pau, and C. Sommer, "Working Group on Heterogeneous Vehicular Networks," in *Dagstuhl Seminar 13392 - Inter-Vehicular Communication - Quo Vadis*. Schloss Dagstuhl, Wadern, Germany: Schloss Dagstuhl, Sep. 2013, pp. 201–204.
- [2] D. Cavalcanti, D. Agrawal, C. Cordeiro, B. Xie, and A. Kumar, "Issues in Integrating Cellular Networks, WLANs, and MANETs: A Futuristic Heterogeneous Wireless Network," *IEEE Wireless Communications*, vol. 12, no. 3, pp. 30–41, 2005.
- [3] G. Rémy, S.-M. Senouci, F. Jan, and Y. Gourhant, "LTE4V2X: LTE for a Centralized VANET Organization," in *IEEE Global Telecommunications Conference (GLOBECOM 2011)*. Houston, TX: IEEE, Dec. 2011.
- [4] L.-C. Tung, J. Mena, M. Gerla, and C. Sommer, "A Cluster Based Architecture for Intersection Collision Avoidance Using Heterogeneous Networks," in *12th IFIP/IEEE Annual Mediterranean Ad Hoc Networking Workshop (Med-Hoc-Net 2013)*. Ajaccio, Corsica, France: IEEE, Jun. 2013.
- [5] F. Hagenauer, F. Dressler, and C. Sommer, "A Simulator for Heterogeneous Vehicular Networks," in *6th IEEE Vehicular Networking Conference (VNC 2014), Poster Session*. Paderborn, Germany: IEEE, December 2014, pp. 185–186.
- [6] A. Virdis, G. Stea, and G. Nardini, "SimuLTE - A Modular System-level Simulator for LTE/LTE-A Networks based on OMNeT++," in *4th International Conference on Simulation and Modeling Methodologies, Technologies and Applications (SIMULTECH 2014)*, Vienna, Aug 2014.
- [7] 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, Jan. 2011.
- [8] T. Taleb and A. Benslimane, "Design Guidelines for a Network Architecture Integrating VANET with 3G & beyond Networks," in *IEEE Global Telecommunications Conference (GLOBECOM 2010)*, Miami, FL, Dec 2010.
- [9] A. Benslimane, T. Taleb, and R. Sivaraj, "Dynamic Clustering-Based Adaptive Mobile Gateway Management in Integrated VANET - 3G Heterogeneous Wireless Networks," *IEEE Journal on Selected Areas in Communications*, vol. 29, no. 3, pp. 559–570, March 2011.

¹<http://floxzy.at/veins-lte/>

²<http://www.ccs-labs.org/projects/car4ict/>

Spatio-temporal distributed background data storage and management system in VANETs

Josip Balen

Faculty of Electrical Engineering, J. J. Strossmayer University of Osijek, Croatia

josip.balen@etfos.hr

Abstract—With the advances in the development of embedded technology, wireless communication and VANETs, and the awareness that vehicles possess a vast amount of unused resources, new services and applications have been proposed recently. The emphasis of this paper is on describing an idea, related work, and open research challenges of using parked vehicles as a spatio-temporal distributed storage system for large amounts of data. We have briefly surveyed different data storage and management techniques in VANETs in order to get directions for the implementation of our strategy. Open research challenges and system requirements, as well as future work are pointed out.

I. INTRODUCTION AND MOTIVATION

The use of parked vehicles in Vehicular Ad Hoc Networks (VANETs) has gained attention in the recent years due to the awareness that they represent static roadside nodes which are available in large number, long-time stationary, well distributed among specific locations, and often a vast of unused resources. Therefore, parked vehicles have been included in VANETs as a complement to the driving vehicles and Road Side Units (RSUs). Moreover, new services and applications have been proposed recently, for improving Internet access throughout [1], extending RSUs' service coverage [2], and assisting Vehicle-to-Vehicle (V2V) communication [3].

Due to the dynamic and unbalanced nature of VANETs, topology management is one of the key issues while providing services. Here as well, VANETs benefit from leveraging parked vehicles since they remain static for a longer period of time and exist in large numbers. Furthermore, roughly 70% of vehicles are parked for an average of 23 hours per day [2], during the day in front of companies and in the city centers and during the night in front of home buildings and houses. Hence, by using parked vehicles as a backbone in VANETs it is possible to offer services and applications requiring certain stability and reliability like storage of large amounts of data.

Motivation for this paper also comes from the fact that parking lots are distributed throughout the urban areas at pre-known locations and occupied with parked vehicles on average more than 50% of the time [4]. Therefore, parked vehicles in urban areas represent a huge unused resource in terms of distributed storage and processing power that could be used as a spatio-temporal distributed data storage system. Furthermore, VANETs are ubiquitous, particularly in urban environments, hence users can easily access it and transfer data.

Another important point, as described in [5], is that persistent storage is becoming less expensive over time and it is expected

that computers in vehicles will have multiple types of storage attached. Furthermore, vehicles are not considered as resource constrained since they have plenty of space to accommodate multiple hard drives even with today's technology and sizes. This will result in a vast amount of unexploited and wasted storage resources.

A centralized approach to using mobile network infrastructure and 3G/4G for connecting to online servers is often expensive especially when roaming. Furthermore, it is unlikely to expect the Internet uplink capacity to scale faster than users' bandwidth needs. Although transferring bigger amounts of data using a mobile network infrastructure is costly in time and money, due to the global access capabilities using mobile network infrastructure it should still be considered for low data rate applications.

We envision the two following scenarios for using spatio-temporal distributed background data storage and management systems in VANETs. In the first scenario, data is generated by vehicles or various infrastructure; in the second, data is transferred by users. The amount and size of data generated by vehicles and infrastructure like traffic cameras, traffic lights, traffic signs or weather stations is constantly increasing. We can expect vehicles and infrastructure to generate multimedia data like photos and videos, and therefore storage and management of large amounts of data will be required. In the second scenario there are two types of users: single people and companies. An example use case for single people would be a situation where

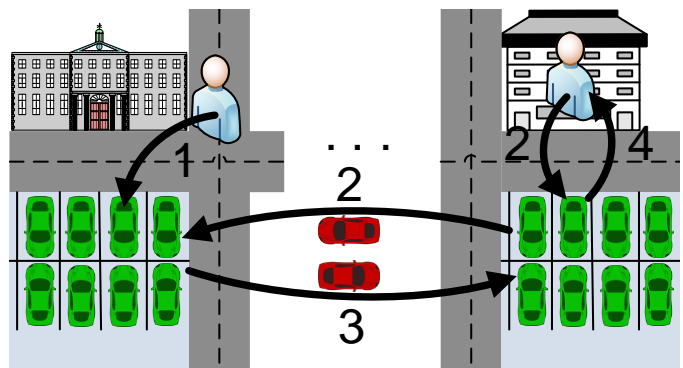


Figure 1. Spatio-temporal distributed data management example scenario. (1) User stores data in the VANET; (2) User starts data lookup; (3) Data is transferred from the source to the destination vehicle; (4) Data is transferred to the user.

they would like to temporarily store data from their smartphones and cameras in VANETs when they run out of space on their devices, as shown in Figure 1. Some interesting use cases for companies are mentioned in [5]. The authors propose using vehicles as data centers in parking lots of the companies, malls or airports. Drivers could rent their vehicle resources to the companies on a per-day, per-week or per-month basis. As a result parking lots could be turned into huge distributed data storage facilities.

In the rest of the paper, we survey different data storage and management techniques in VANETs in order to get directions for the implementation of our strategy. Furthermore, open research challenges are pointed out as well as system requirements and future work.

II. RELATED WORK

Using VANETs as a spatio-temporal distributed background data storage and management system has already been identified in the literature. We have briefly surveyed different strategies in order to get directions for the implementation of our strategy.

Location-based strategy: In [6], a location based data overlay strategy for Disruption-Tolerant Networks (DTNs) called Locus has been proposed. The basic idea is to maintain the locality of data by exchanging data between nodes that are close enough to the location where data was generated (home location). The authors argue that DHTs are not appropriate for the DTNs since the network model underlying DHTs assumes stable connectivity between all nodes participating in the overlay. Moreover, they point out that since a VANET is a highly dynamic network, a DHT based system will be unreliable and hence will have low probability of finding data in DTNs. Therefore, data objects are used as the primary entities. Data objects can contain any type of information that are tagged with an application-specific type indicator, the timestamp, and location of data creation. Data storage is based on an utility function that is based on the distance from the data location. It defines three zones: home area, drop-off zone, and the void zone. In order to keep data close to the home zone, data replication is used in the three following cases: (i) if a node is currently in the data's home area it will copy the data to surrounding nodes to increase the probability of keeping the replica in the home area; (ii) if a node is not in the drop-off zone the data is copied to the nodes closer to the home location; (iii) if a node is outside the data drop-off zone data should be replicated to the node that is moving towards the drop-off zone. Data access is done through queries and response messages. Since the primary matching condition of the query is location, nodes requesting data have to know the data location. Beside location, a querying node is sending its location and trajectory. After the data is found the node generates the response message and (based on received data and age of the query) it estimates the querying node's current location. This approach seems unrealistic, since a node's trajectory shouldn't be known due to privacy concerns. This could lead to errors in calculating the querying node's future location. Moreover, it is much easier (and loads the network much less) to keep the data close to

the place where it has been generated just by transferring it to the nearby parked vehicles.

Parked vehicles domains strategy: Using parked vehicles as a temporary network and storage infrastructure is presented in [7]. The idea is to have a vehicular cloud consisting of independent clusters of parked vehicles where the data will be stored. In order to establish spatio-temporal network infrastructure for providing connectivity as well as data storage and management capabilities, the Virtual Cord Protocol (VCP) is used. VCP provides Distributed Hash Tables (DHT) services, integrated greedy routing with guaranteed delivery and inter-domain routing. VCP is based on sending Hello messages and including new vehicles in an existing domain. If existing domains are not accessible in a defined period of time, creation of new a domain is initiated. Data management is based on PUBLISH and LOOKUP operations with greedy routing for message forwarding. Upon reception of a publish message, data is mapped to hash values. Although in [6], it is argued that a DHT based system will be unreliable and have low probability of finding data in DTNs here is shown that with using DHT for data management it is possible to get a high rate of successful data access operations. Future work is need in order to build efficient inter-domain routing strategy and connect multiple disconnected domains of parked vehicles with a store-carry-forward paradigm.

Vehicular cloud computing: This concept is based on cloud computing and considers each vehicle as a computation node. The main difference between traditional and vehicular cloud computing is due to the dynamic environment of VANETs and in terms of resource availability. An exhaustive survey on vehicular cloud computing can be found in [8]. There are several potential architectures based on using vehicles' processing, storage and network resources: Network as a service (NaaS), Storage as a service (STaaS), Cooperation as a service (CaaS), Information as a service (INaaS), Entertainment as a service (ENaaS), Computing as a service and Pictures on wheels as a service. For our work, the most interesting architecture is STaaS. Since it is predicted that vehicles will have huge storage resources, [5] proposed to use parked vehicles as datacenters at the parking lots of malls, airports, and companies since there are often several hundreds of parked vehicles. However, they conclude that users cannot immediately benefit from storage like they can from computation power or network access since vehicles will eventually leave the parking lot and take data with them. Therefore, they propose creating multiple replicas of the same data or decomposing data into smaller block sizes that can be quickly delivered on request so that big files can be reconstructed. However, the concept is not farther expanded on.

III. SYSTEM REQUIREMENTS AND RESEARCH CHALLENGES

The main differences between traditional distributed data storage and data storage based on parked vehicles are in availability and wide distribution of nodes, dynamic environment and resource availability. Since parked vehicles can randomly leave and join parking lots, it is necessary to have a data management

system on individual parking lot level. Furthermore, individual parking lots have to be connected with driving vehicles, and therefore it is necessary to have a data management system realizing a store-carry-forward paradigm on a more global level.

The desired spatio-temporal distributed background data storage and management system based on parked vehicles must be reliable and robust, enable data delivery in a timely manner in a big geographical area, cover sparse parts of a network, and prevent network overload. Furthermore, it has to provide on demand data storage and retrieval to users and applications. These requirements raise several research questions and challenges:

redundancy – keep in the network an optimal number of data replicas based on the type and priority of data in order to keep data safe,

storage locations – should be in proximity of the desired data delivery destination and should enable high data reachability. This could be based on prediction, user tracking (based on user request to store data on a different locations), or learning algorithms,

data validity – should be based on the type and priority of data and storage capacity,

buffer management mechanism – if storage capacity on an individual node is completely consumed data could be dropped or transferred to a different node,

area of relevance – define a geographic area based on possible data destination and data source,

timeliness – enable data delivery in a timely manner but at the same time prevent network overload. Data delivery delay directly depends on distance from the storage location. Expedited access for frequently queried data should be supported,

network bandwidth – since different QoS is handled on MAC layer, for transferring large amounts of data full network bandwidth could be used,

data transfer – fastest data transfer from source to destination could be achieved with either V2V greedy routing (for connected source and destination) or store-carry-forward routing (for disconnected source and destination).

There are several methods for meeting the above mentioned research challenges. First, all components could be predefined in the data management system. Second, based on network conditions all components could be dynamically handled. Third, applications and users could request specific quality of service.

In larger cities we can expect to have large scale vehicular networks with several hundred thousands of vehicles which represents huge distributed storage (in space and size) and pose new and unique challenges to efficient data management.

IV. CONCLUSION AND FUTURE WORK

We have scratched the surface of using parked vehicles as a spatio-temporal storage and management system for large amounts of data. Parked vehicles in urban areas represent a huge unused distributed storage and processing power resource that is currently wasted. Therefore, we consider using parked vehicles

as a distributed data storage system for large amounts of data. The goal of our system is to use as little infrastructure (mobile network or RSU) as possible since it costs money and time and is a single point of failure, in contrast with infrastructure-less systems that are mostly cheap and more fault tolerant. The purpose of an efficient data management strategy in VANETs is to provide synergy between data storage and data lookup, and consequently enable reliable storage for large amount of data in VANETs. This overview will generate a productive future research with the implementation and evaluation of spatio-temporal data storage and management system in terms of data delivery efficiency and delay, network traffic, data source and destination distance by using Veins simulator with a realistic scenario.

REFERENCES

- [1] R. Crepaldi, R. Beavers, B. Ehrat, M. Jaeger, S. Biersteker, and R. Kravets, "LoadingZones: Leveraging Street Parking to Enable Vehicular Internet Access," in *7th ACM International Workshop on Challenged Networks (CHANTS 2012)*, Istanbul, Turkey: ACM, Aug. 2012, pp. 23–30.
- [2] F. Malandrino, C. Casetti, C.-F. Chiasserini, C. Sommer, and F. Dressler, "The Role of Parked Cars in Content Downloading for Vehicular Networks," *IEEE Transactions on Vehicular Technology*, vol. 63, no. 9, pp. 4606–4617, Nov. 2014.
- [3] C. Sommer, D. Eckhoff, and F. Dressler, "IVC in Cities: Signal Attenuation by Buildings and How Parked Cars Can Improve the Situation," *IEEE Transactions on Mobile Computing*, vol. 13, no. 8, pp. 1733–1745, Aug. 2014.
- [4] J. Balen, G. Martinovic, K. Paridel, and Y. Berbers, "PVCM: Assisting multi-hop communication in vehicular networks using parked vehicles," in *4th International Congress on Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT 2012)*, Oct. 2012, pp. 119–122.
- [5] S. Olariu, T. Hristov, and G. Yan, "The Next Paradigm Shift: From Vehicular Networks to Vehicular Clouds," in *Mobile Ad Hoc Networking*, Wiley, 2013, pp. 645–700.
- [6] N. Thompson, R. Crepaldi, and R. Kravets, "Locus: A Location-based Data Overlay for Disruption-tolerant Networks," in *5th ACM International Workshop on Challenged Networks (CHANTS 2010)*, Chicago, IL: ACM, 2010, pp. 47–54.
- [7] F. Dressler, P. Handle, and C. Sommer, "Towards a Vehicular Cloud - Using Parked Vehicles as a Temporary Network and Storage Infrastructure," in *15th ACM Int. Symposium on Mobile Ad Hoc Networking and Computing (Mobihoc 2014): ACM International Workshop on Wireless and Mobile Technologies for Smart Cities (WiMobCity 2014)*, Philadelphia: ACM, Aug. 2014, pp. 11–18.
- [8] M. Whaiduzzaman, M. Sookhak, A. Gani, and R. Buyya, "A survey on vehicular cloud computing," *Journal of Network and Computer Applications*, vol. 40, pp. 325–344, 2014.

Inter-Vehicle Communication on the Run

Experiences From Tweaking Veins Runtime Performance

Raphael Riebl, Stefan Neumeier and Christian Facchi
Technische Hochschule Ingolstadt, Research Centre
Email: {raphael.riebl, stefan.neumeier, christian.facchi}@thi.de

Abstract—With increasing complexity of employed simulation models in the field of Inter-Vehicle Communication, the associated computational effort increases as well. Solely adding more hardware resources to simulation computers may be insufficient when the models are not prepared to scale appropriately. Depending on the use case, there can be upper limits for execution times or it is simply annoying to await the end of long-running simulations. For this reason, various methods for identification of performance bottlenecks are examined and with that knowledge methods aiming at improvement of single-threaded simulation execution time are investigated. Those methods were selected with focus on ease of adaption, so they can be easily applied to custom simulations as well. Selected methods were evaluated by the example of OMNeT++ based Veins simulation scenarios.

Index Terms—Inter-Vehicle Communication (IVC), OMNeT++, simulation performance, Vehicle-2-X (V2X), Veins

I. INTRODUCTION

Probably most users of Veins [1], a framework for vehicular communication simulations, have already observed how one of their Central Processing Unit (CPU) cores is utilised to the full while remaining cores are idling. Therefore, the idea of spreading the simulation on more of those CPU cores comes quite naturally. While it is straightforward to run multiple, mutually independent simulation processes in parallel, real-time requirements are not as easy to satisfy. These requirements emerge when the simulation environment acts as testbed for a hardware device, as depicted e.g. in [2]. It is also favourable to have short execution times during the preparation of experiments when the simulation environment is set up and configured. From our experience this process is accompanied by repeated adjusting and testing of simulation settings. Thus, shorter execution times help to speed up this preparation phase.

In the past, when CPU speed doubled every 12–24 months [3], upgrading to newer hardware has been an approach for solving performance issues. Nowadays, CPU speed improvements per core level off, but multi-core CPUs are available en masse [4], [5]. Making use of these cores in order to speed up single execution times of simulations, however, involves in-depth changes to the simulation model [6]. Although OMNeT++ [7] supports parallelisation through e.g. Message Passing Interface (MPI) [8], it is still necessary to partition the simulation model in a smart way so that parallelisation gain is not exhausted by occurring synchronisation overhead.

Hence, instead of investigating performance improvements through horizontal scaling, i.e. distributing the problem on

multiple cores, this paper focusses on enhancing single-thread performance. In the following sections, we present tools for evaluating simulation performance and how even comparatively small changes regarding compiler settings, memory management, data structure and model level can lead to considerable speed-ups.

II. TOOLS FOR MEASURING PERFORMANCE

There exist several ways to measure execution time on a Linux system. For the measurements in this paper, three tools – *GNU time*, *Valgrind*, and *gprof* – were employed. Each of those tools follows a different approach for measuring execution time and gives different results therefore. For interpretation of their results it is inevitable to have a basic understanding of their operation mode.

GNU time [9] reports three time components regarding the execution time of a process. First of all, it reports the total elapsed time from program start until exit, which is also referred to as wall clock time. Furthermore, *GNU time* unveils user time and system time, which are the durations the monitored process actively consumed CPU resources in user mode or system mode respectively. By default, a process runs in user mode and only switches over to system mode when it engages the operating system to accomplish something for it, e.g. allocating memory. It has to be noted, though, that user and system time do not sum up to the wall clock time, because the wall clock time also includes the lapse of time the monitored process is inactive, e.g. when waiting for the completion of Input/Output (I/O) operations.

Valgrind [10], [11], on the other hand, measures radically different. It executes the program on its synthetic CPU, so a sub-tool of *Valgrind* called *callgrind* provides exact information about the number of CPU instructions every single program function takes, how often and by whom those were called. Aside from counting CPU instructions, *callgrind* is also capable to simulate CPU cache behaviour, which allows to estimate the required CPU cycles per instruction and therefore time. Running the program on a synthetic CPU, however, means also some runtime penalty. It is not uncommon to experience the tenfold program runtime. Consequently, I/O operations seem to finish much faster from the perspective of the now slower running simulation program, i.e. only CPU bottlenecks can be detected, but not I/O bottlenecks.

Whereas the aforementioned tools do not require any changes to the simulation program, *gprof* [12], [13] requires code instrumentation by the compiler, i.e. enabling the *-pg* flag

of GNU GCC. When a *gprof* enabled program is run, number of calls per function are logged and time spent in instrumented functions is sampled. Due to sampling the measured data are of coarser granularity than those by *Valgrind*, but *gprof* has also less impact on the execution time compared to *Valgrind*. In the context of this paper, *gprof* has been used for validating plausibility of measurements by other tools.

Interpretation of measurements is easier with less components affecting simulation runtime. Thus, effects slowing down simulation performance have been excluded as far as possible. Namely, all measurements were conducted using OMNeT++’s command line interface *Cmdenv*, so latencies arising from refreshing Graphical User Interface (GUI) components are omitted. Furthermore, printing output to the command line has been reduced to a minimum. Additionally, turning off recording of simulation data diminishes access to hard disc during runtime. Thus, the OMNeT++ process is not suspended by the operating system kernel due to blocking file operations. Of course, recording this data is usually the primary reason for running a simulation at all, but the amount and frequency of data writes depends vastly on an experiment’s intention, so a general advice can not be given. Furthermore, by shifting write operations to a distinctive thread, OMNeT++ recording becomes non-blocking.

III. METHODS AND RESULTS

For the following measurements OMNeT++ 4.6 and Veins 4a1 have been used. Changes to the original source code are marked where applicable. Source code has been compiled with GNU GCC 4.9 and enabled C++11 support. Runtime measurements are affected by many factors like the CPU architecture with its various caches, the operating system’s scheduling of processes, compiler vendors and versions, simulation settings and workload. Therefore, absolute numbers differ largely from system to system and performance improvements are better described by relative gains.

Data from *callgrind* measurements in Table I and *GNU time* measurements in Table II demonstrate how the selected simulation settings affect runtime and shift perception of performance bottlenecks. Except for the additional transmission of beacons in the *Data + Beacons* configuration, the simulation set-up is identical to the *Data* configuration. With the *Data* configuration only a total of 196 data packets are sent during the whole simulation run, whereas *Data + Beacons* includes periodically sent beacon packets, which sum up to a total of 78206 packets sent in the same simulation time span. Consequently, there is a much higher stress on packet processing functions for *Data + Beacons* as can be seen by the function `BasePhyLayer::handleAirFrame`. In this function’s *Inclusive* costs, i.e. the amount of CPU instructions consumed by this function (*Self* costs) and all functions called by it, `Veins::ObstacleControl::calculateAttenuation` is also included, which stands out because of its comparatively high *Self* costs. With a severely reduced number of transmitted packets, measurements

TABLE I
CPU INSTRUCTIONS (RELATIVE TO ALL INSTRUCTIONS) SPENT PER FUNCTION DEPENDING ON SIMULATION SETTINGS

Function	Configuration		Data		Data + Beacons	
	Inclusive	Self	Inclusive	Self	Inclusive	Self
<code>BasePhyLayer::handleAirFrame</code>	12.86 %	0.00 %	95.74 %	0.01 %		
<code>BaseConnectionManager::updateNicConnections</code>	26.26 %	13.84 %	0.48 %	0.25 %		
<code>malloc</code>	9.76 %	3.77 %	22.34 %	9.04 %		
<code>Veins::ObstacleControl::calculateAttenuation</code>	1.94 %	0.17 %	26.32 %	3.37 %		
<code>Veins::TraCIScenarioManager::processVehicleSubscription</code>	81.37 %	1.73 %	1.49 %	0.03 %		

TABLE II
EXECUTION TIMES MEASURED WITH GNU TIME

Configuration	CPU time	Wall clock	CPU utilisation
Data (Debug)	24.44 s	33.90 s	72 %
Data + Beacons (Debug)	1364.57 s	1379.35 s	98 %
Data (Release)	6.91 s	16.71 s	41 %
Data + Beacons (Release)	190.07 s	200.58 s	94 %

with *Data* configuration highlight the runtime costs of processing incoming TraCI data like updating each vehicle node’s position as part of `processVehicleSubscription`. A big chunk of work is apportioned to `updateNicConnections` due to node position changes.

Running a simulation under the supervision of *Valgrind*, however, does not show the delays when Veins has to wait for SUMO to finish its update cycle and sending its data back over TraCI. Since the simulation process primarily needs to be suspended by the operating system scheduler when waiting for SUMO, this delay is roughly the difference between elapsed wall clock and CPU time (sum of user and system time). This difference is about 10 seconds for both configurations, though *Data + Beacons* runs almost 12 times as long as *Data*, as the *GNU time* measurements in Table II show. The CPU utilisation also reflects this observation: Higher utilisation comes along with less waiting on I/O operations, in our case TraCI network communication I/O. Since the amount of TraCI I/O is the same in both configurations, the waiting time for I/O is roughly equal though their total execution times differ vastly.

A. Compiler optimisations

Execution times in Table II show the impact of compiler flags (*Debug* and *Release*) on the simulation performance, which are therefore investigated in more detail. Enabling optimisations offered by the C++ compiler are an easy way to improve code execution time. Figure 1 depicts the averaged execution times achieved using various compiler optimisation levels on three different systems with x86_64 CPU designs. For brevity, measured times for binaries including debug

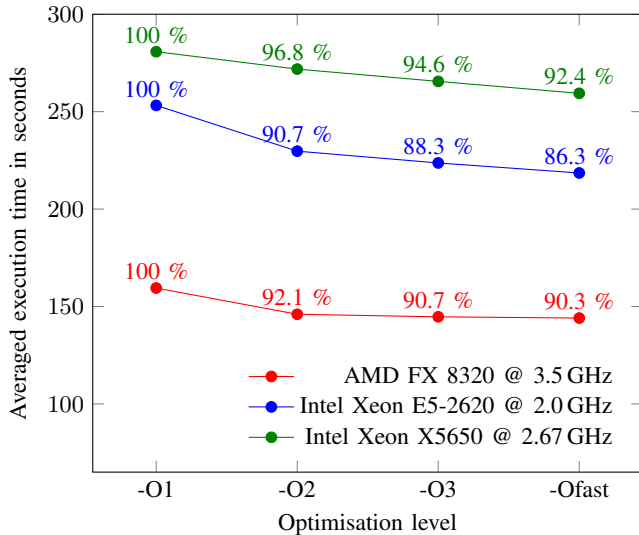


Fig. 1. Influence of compiler optimisation levels on execution time, averaged over 15 measurements for Data + Beacons configuration

symbols, i.e. compiled with `-g` flag set, are omitted, because measurements showed no significant difference to those compiled without debug symbols. Execution time differences between simulations built in *Debug* and *Release* mode have to be attributed to the optimisation level instead. Since execution times increase fourfold and more without any enabled optimisations compared to basic optimisation level `-O1`, those measurements were left out for time reasons.

It can be concluded from these measurements that more modern CPU generations can still out-perform higher CPU clock rates. For similar CPU generations, however, a higher clock rate is advantageous. On each of the investigated systems, the achievable gain levels off with more aggressive optimisation levels. Furthermore, a trade-off between runtime and compilation-time has to be made: Compared to compilation of Veins with all optimisations disabled, `-O2` requires over 25% more time and `-Ofast` even 30%.

B. Memory management

Table I reveals a strong influence of dynamic memory allocation through `malloc` on the overall simulation runtime. This influence is emphasised when a plethora of packets are transmitted, as each transmission involves at least one memory allocation. Because of the way packets are handled in OMNeT++, it is almost unfeasible to reduce the number of `malloc` calls. Instead, it is more promising to reduce the costs per call of `malloc`. This can be achieved by replacing the dynamic memory management functions `malloc` and `free` with an implementation performing better than the default library, in this case `glibc`. In this paper, `TCMalloc` of Google’s `perfutils` [14] has been used as replacement. After mere linking of the `TCMalloc` library to the simulation executable, dynamic memory is handled by `TCMalloc`. No code changes are necessary. Achieved speed-ups are presented in Table III, which reveals improvements up to 24.7% over `glibc`.

TABLE III
GNU TIME MEASUREMENTS REGARDING DYNAMIC MEMORY MANAGEMENT (AVERAGE OVER 15 RUNS)

Configuration System	Data		Data + Beacons	
	<code>glibc</code>	<code>TCMalloc</code>	<code>glibc</code>	<code>TCMalloc</code>
AMD FX 8320	20.4 s	19.6 s (-3.9%)	146.0 s	110.0 s (-24.7%)
Intel Xeon E5-2620	31.2 s	31.7 s (+1.6%)	229.7 s	188.9 s (-17.8%)
Intel Xeon X5650	35.8 s	34.9 s (-2.5%)	271.8 s	216.5 s (-20.3%)

C. Data structure

Analysing the `callgrind` measurements for the *Data + Beacons* configuration, 4.52% of all instructions are a result of rebalancing operations of C++ `std::map`’s the internal tree structure. Most rebalancing operations are due to `ObstacleControl::calculateAttenuation` modifying its `cacheEntries` map, accounting for 2.68% of all instructions. Since there is no reason for keeping the order of these cached entries – only single values are ever stored and retrieved – using `std::unordered_map` instead of `std::map` might improve the runtime performance. However, measured performance gain ranges between 2.2% slower and 2.7% faster execution times on average, depending on the employed computer system. Instead of tree rebalancing there are now costs for calculating hash values for the cache entries. While the implementation effort is acceptable¹, the gain is neither particularly exciting nor present on every system.

D. Simulation model

The level of detail of the employed simulation model has significant influence on the required execution time. By omitting obstacle shadowing effects for radio propagation in Veins, execution time could be reduced by over 30%. If it is possible to neglect obstacle shadowing for an experiment, then this is an obvious runtime optimisation. However, comparing different implementations of a model might reveal performance gains without changing the level of detail covered by the simulation. Thus, we compare standard Veins with a variant where the IEEE 802.11 implementation and radio channel model of INET² are used. Because Medium Access Control (MAC) layer interfaces of Veins and INET are not exactly the same, some adjustments³ had to be done. From a simulation modelling perspective, both variants should cover the same level of detail: They employ the identical Inter-Vehicle Communication (IVC) application, same road map and traffic flow, as well as equivalent radio settings like carrier frequency, transmission power, bitrate and path loss model.

At first, a slow-down of more than 30% compared to the original Veins was disappointing, as listed in Table IV.

¹ Available at <https://gist.github.com/riebl/e98b6ae98a63c5e75e65>

² Tested with commit `bf56c898a0d` from INET integration branch

³ Available at <https://github.com/riebl/veins/tree/fg-ivc-2015>

TABLE IV
EXECUTION TIMES OF MODEL VARIANTS FOR DATA + BEACONS
CONFIGURATION (AVERAGE OVER 15 RUNS)

Model implementation	Execution time	
	seconds	relative
Veins	229.73 s	100 %
INET (with leak)	300.03 s	130.6 %
INET	222.15 s	96.7 %




Fig. 2. Average of 15 measured execution times for combined methods

Further analysis, though, revealed over-proportional memory consumption by the INET variant. After elimination of this memory leak⁴, the INET variant operated slightly faster than the unmodified Veins. It has to be noted, though, that the effective communication range is larger with INET and thus on average packets are received by more vehicles than with Veins, i.e. these implementations behave differently although configuration was careful adjusted to be as similar as possible.

E. Combination of Methods

Summarising the results depicted in Figure 2, the overall improvement by combining all aforementioned methods is higher than by any other single method. For this evaluation the system with the Intel Xeon X5650 CPU has been used, because it exposed mid-level gains for each method.

IV. CONCLUSION

Donald Knuth once warned of misguided code optimisations striving for performance gains: “Premature optimization is the root of all evil” [15]. Considering this advice and seeing the striking dependence of runtime performance on the particular simulation settings and level of detail, we had only a look at optimisation methods general enough to be applicable to a broader range of settings and can be integrated with little effort. But it can be pointed out that even simple means can reduce the execution time by one quarter.

A performance analysis of Veins was conducted using the tools *Valgrind*, *GNU time* and *gprof* to gain insight of present bottlenecks, whereby each tool represents a different point of view. Subsequently, four different countermeasures

have been investigated: compiler optimisation levels, memory management, data structure and simulation model changes. All of these approaches have shown to be appropriate to enhance single-threaded execution time at least in some constellations. Individual changes vary from almost 4% performance loss up to impressive 25% gain per optimisation method. Each approach is applicable without enormous code refactoring or change of a simulation model’s level of detail, which has been an important objective when we selected those approaches.

Despite this success improving single thread speed of the examined OMNeT++ based Veins simulation, there still remains a desire to make use of the nowadays ubiquitously available multi-core computing power. There are already steps towards parallelisation of radio channel calculations in INET⁵ and event processing in OMNeT++ [16]. Future steps will need to go into this direction.

REFERENCES

- [1] C. Sommer, R. German, and F. Dressler, “Bidirectionally Coupled Network and Road Traffic Simulation for Improved IVC Analysis,” *IEEE Trans. Mob. Comput.*, vol. 10, no. 1, pp. 3–15, Jan. 2011.
- [2] R. Riebl and C. Facchi, “Implementation of day one ITS-G5 systems for testing purposes,” in *Proceedings of the 2nd GIITG KuVS Fachgespräch Inter-Vehicle Communication (FG-IVC 2014)*, Feb. 2014, pp. 33–36.
- [3] S. E. Thompson and S. Parthasarathy, “Moore’s law: the future of Si microelectronics,” *Materials Today*, vol. 9, no. 6, pp. 20–25, 2006.
- [4] H. Sutter, “The free lunch is over: a fundamental turn toward concurrency in software,” *Dr. Dobbs’s journal*, vol. 30, no. 3, pp. 202–210, 2005.
- [5] P. Rogers, *The programmer’s guide to the APU galaxy*, AMD Fusion Developer Summit, Jun. 2011.
- [6] M. Stoffers, R. Bettermann, J. Gross, and K. Wehrle, “Enabling distributed simulation of OMNeT++ INET models,” in *Proceedings of 1st OMNeT++ Community Summit*, Sep. 2014.
- [7] 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*, ser. Simutools ’08, ICST, Brussels, Belgium, 2008, 60:1–60:10.
- [8] A. Varga and A. Y. Şekercioğlu, “Parallel simulation made easy with OMNeT++,” 2003. [Online]. Available: <http://www.omnetpp.org/doc/publications/vargasekerici09parsim.pdf>.
- [9] D. C. Niemi *et al.* (Feb. 10, 2015). time - GNU Project - Free Software Foundation. [Online]. Available: <https://www.gnu.org/software/time/>.
- [10] N. Nethercote and J. Seward, “Valgrind: a framework for heavyweight dynamic binary instrumentation,” in *Proceedings of the 2007 ACM SIGPLAN Conference on Programming Language Design and Implementation*, ser. PLDI ’07, 2007, pp. 89–100.
- [11] J. Seward, N. Nethercote, T. Hughes, J. Fitzhardinge, J. Weidendorfer, *et al.* (Feb. 10, 2015). Valgrind. [Online]. Available: <http://valgrind.org/docs/manual/index.html>.
- [12] S. L. Graham, P. B. Kessler, and M. K. Mckusick, “Gprof: a call graph execution profiler,” in *Proceedings of the 1982 SIGPLAN Symposium on Compiler Construction*, ser. SIGPLAN ’82, 1982, pp. 120–126.
- [13] J. Fenlason. (Feb. 10, 2015). GNU gprof. [Online]. Available: <https://sourceware.org/binutils/docs/gprof/>.
- [14] S. Ghemawat. (Feb. 11, 2015). TCMalloc: Thread-Caching Malloc. [Online]. Available: <http://google-perftools.googlecode.com/svn/trunk/doc/tcmalloc.html>.
- [15] D. E. Knuth, “Structured programming with go to statements,” *ACM Comput. Surv.*, vol. 6, no. 4, pp. 261–301, Dec. 1974.
- [16] G. Kunz, M. Stoffers, J. Gross, and K. Wehrle, “Runtime efficient event scheduling in multi-threaded network simulation,” in *Proceedings of the 4th International ICST Conference on Simulation Tools and Techniques*, ser. SIMUTools ’11, 2001, pp. 359–366.

⁴See <https://github.com/inet-framework/inet/pull/44>

⁵Available at <https://github.com/inet-framework/inet/tree/topic/radio-cuda>

Discussing Different Levels of Privacy Protection in Vehicular Ad-Hoc Networks

David Förster
Robert Bosch GmbH
david.foerster@de.bosch.com

Abstract—This paper gives an overview of privacy issues in vehicular ad-hoc networks and connected cars. It illustrates the need for strong privacy by practical examples. Technical solutions that provide different levels of privacy are presented and evaluated against a set of partially contradicting requirements. The available of privacy-friendly solutions shows that strong privacy can in fact be achieved in practice.

I. INTRODUCTION

Vehicular ad-hoc networks based on Car2X (C2X) communication have been subject to research for quite some time. With the imminent deployment of the technology the public is becoming more and more aware of it and its potential privacy implications. Many critical voices have been raised in the media on whether connected cars pose a threat to user privacy [1]–[3]. The benefit expected from C2X applications will only be realized when a certain market penetration has been reached. In order to ensure rapid and smooth adoption of the technology it is imperative to address consumers’ concerns.

BMW has made public the increasing number of requests from third party companies to make use of the vast amount of data collected by modern automobiles [4]. Legal aspects of connectivity and ownership of collected data have yet to be established. Many of these problems can be avoided altogether by using privacy-friendly systems that prevent data collection in the first place. On the other hand data minimization may conflict with (legitimate) use of collected data.

In this paper we discuss real-world privacy requirements, how well they are addressed by proposed privacy-friendly systems, and which other, possibly contradicting requirements exist. Informally, we want to examine how much privacy we can achieve without affecting the functionality of C2X systems.

II. MOTIVATION AND REQUIREMENTS

The parties towards one requires privacy may be the government, official authorities, commercial entities, colleagues, friends, family, or partners. From a lack of privacy one might experience severe (subjective) disadvantages such as *a*) discrimination based on political or other preferences inferred from locations visited, *b*) surveillance or stalking, *c*) unfavorable insurance conditions based on driving behavior, *d*) more efficient prosecution of traffic violation, or other.

Research shows that the availability of location traces alone poses a threat to user privacy as drivers’ homes [5] and even their identities [6] can be inferred.

Schaub et al. define a set of formal privacy requirements for C2X systems [7]. We omit *minimum disclosure*, which is always a requirement for privacy-friendly systems.

Anonymity Messages sent by a participant do not allow to infer his identity. It can be split into anonymity towards other participants and anonymity towards backend providers.

Unlinkability Two “items of interest” related to one participant cannot be linked, e.g. C2X messages sent at two different locations by the same user cannot be linked.

Perfect forward privacy Resolution or revocation of one credential or message does not affect unlinkability of any of the user’s other credentials or messages.

They also define a set of *security* requirements they deem necessary for the undisturbed operation of the C2X system.

Authentication Establish that the sender of a message is a valid participant of the C2X system.

Accountability The authorities can hold users accountable for their action, i.e. for the messages they send.

Restricted credential usage Limit the number of credentials a user can use at the same time to prevent “sybil” attacks.

Credential revocation Revoke a user’s credentials.

The publication also describes that there is a contradiction between some of the privacy and some of the security requirements [7]. Intuitively, authentication is detrimental to anonymity. Also, accountability contradicts anonymity and unlinkability, at least with respect to the party that can hold users accountable.

There is one additional *functional* requirement.

Local linkability Messages sent must be linkable on a local level in order to allow vehicles to maintain a “local dynamic map” of nearby vehicles. Insufficient local linkability can negatively affect C2X-based safety functions [8].

III. PRIVACY-FRIENDLY SYSTEMS

Privacy-friendly authentication is usually achieved by using “pseudonym certificates”, which are issued to participants of a C2X system [9], [10]. As these pseudonym certificates do not contain their owner’s identity, they can be used for signing C2X messages without violating the requirement of anonymity. In the “basic scheme” pseudonym certificates are usually obtained from backend providers in exchange for presenting a “long-term certificate” that does contain the owner’s identity. Hence, this approach does not provide anonymity towards the backend providers as they can record which user they issued a certain pseudonym certificate to. That way accountability

	Anonymity among participants	Anonymity towards backend providers	Unlinkability of messages	Perfect forward privacy	Authentication	Accountability	Restricted credential usage	Credential Revocation	Local linkability
Basic pseudonym scheme	✓	✗ ¹	✓ ²	✗ ¹	✓	✓	✓	✓	✓
PUCA	✓	✓	✓ ²	✓	✓	✗	✓	✓ ³	✓
Credential based approaches	✓	✗ ⁴	✓	✗	✓	✓	✗	✓	✗

Table I. COMPARISON OF MESSAGE AUTHENTICATION SCHEMES WITH RESPECT TO THE REQUIREMENTS.

can be achieved. Privacy protection against backend providers can be implemented by separation of duties, however, this protection is organizational only [11].

Full anonymity is implemented by the “PUCA” scheme, that uses anonymous credentials for privacy-friendly authentication with the backend providers when requesting pseudonym certificates [12]. As a consequence it does not offer accountability.

There are several suggestions for using anonymous credentials [13] or group signatures [14] directly for C2X message authentication. We call this the “credential based approaches”. Their drawback is questionable performance on the hand and complete unlinkability of messages sent on the other hand, which violates the requirement of local linkability.

Table III shows a comparison of the three approaches. None of the schemes fulfills all requirements. The credential based approaches are impractical due to their insufficient performance negative effect on C2X core functionality. Choosing between the basic pseudonym scheme and the advanced PUCA system is a trade-off between privacy and control. It has yet to be decided whether accountability, as a regulatory requirement, is really needed. In certain scenarios, e.g. in countries where civil rights are not firmly established, a fully anonymous system might be preferable. Certain revocation scenarios rely on traceability (resolving a sender’s identity from a message; implied by accountability). Revoking credentials from a malicious attacker always requires traceability. Excluding a malfunctioning vehicle may be possible using a cooperative approach, where other vehicles that detect the malfunction request the vehicle to revoke its own credentials. We did not examine requirements for misbehavior detection, which may present similar challenges as revocation.

C2X communication is not the only area where privacy is important. For electric vehicles charging may have privacy implications. Höfer et al. developed POPCORN [15] a system for privacy-friendly vehicle charging. A serious threat to location privacy are mobile phones as current protocols expose user locations to the providers. Angermeier et al. propose PAL-privacy augmented LTE [16] as a privacy-friendly replacement of the standard LTE protocol.

¹Optional protection by organizational separation of entities ²After pseudonym change ³Only with the driver’s consent ⁴Technically possible but not implemented by the proposed schemes

IV. CONCLUSION

Vehicular ad-hoc networks and connected cars bring new challenges with regard to user privacy. Data minimization becomes more important as collected data could potentially be forwarded among vehicles and uploaded to backend systems. Strong privacy is not trivial to achieve as it sometimes collides with security and functional requirements. Yet, a number of privacy-friendly schemes have been proposed in different areas such as C2X message authentication, electronic vehicle charging, and privacy augmented LTE. This shows that strong privacy is in fact possible but at the same time highlights the need for further research.

REFERENCES

- [1] H. Foy, “Connected cars raise privacy and safety worries,” *Financial Times*, 20 November 2014. [Online]. Available: <http://www.ft.com/cms/s/0/e663c6fa-643b-11e4-bac8-00144feabdc0.html>
- [2] M. Wall, “Is your connected car spying on you?” *BBC*, 5 November 2014. [Online]. Available: <http://www.bbc.com/news/business-29566764>
- [3] M. Ramsey and N. Boudette, “Auto makers face questions over privacy in connected cars,” *The Wall Street Journal*, 16 January 2014. [Online]. Available: <http://www.wsj.com/articles/SB1000142405270230414940457932454223762862>
- [4] A. Sharman, “BMW sounds alarm over tech companies seeking connected car data,” *Financial Times*, 14 January 2015. [Online]. Available: <http://www.ft.com/cms/s/0/685fe610-9ba6-11e4-950f-00144feabdc0.html>
- [5] B. Hoh, M. Gruteser, H. Xiong, and A. Alrabady, “Enhancing security and privacy in traffic-monitoring systems,” *Pervasive Computing, IEEE*, vol. 5, no. 4, pp. 38–46, 2006.
- [6] J. Krumm, “Inference attacks on location tracks,” in *Pervasive Computing*. Springer, 2007, pp. 127–143.
- [7] F. Schaub, Z. Ma, and F. Kargl, “Privacy requirements in vehicular communication systems,” in *Computational Science and Engineering, 2009. CSE '09. International Conference on*. IEEE, 2009, pp. 139–145.
- [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] P. Papadimitratos, L. Buttyan, T. Holzer, E. Schoch, J. Freudiger, M. Raya, Z. Ma, F. Kargl, A. Kung, and J.-P. Hubaux, “Secure vehicular communication systems: design and architecture,” *Communications Magazine, IEEE*, vol. 46, no. 11, pp. 100–109, 2008.
- [10] J. Petit, F. Schaub, M. Feiri, and F. Kargl, “Pseudonym schemes in vehicular networks: A survey,” *Communications Surveys Tutorials, IEEE*, vol. PP, no. 99, p. 1, August 2014.
- [11] N. Bissmeyer, H. Stubing, E. Schoch, S. Gotz, J. P. Stotz, and B. Lonc, “A generic public key infrastructure for securing car-to-x communication,” in *18th ITS World Congress*, Orlando, USA, 2011.
- [12] D. Förster, F. Kargl, and H. Löhr, “PUCA: A pseudonym scheme with user-controlled anonymity for vehicular ad-hoc networks (VANET),” in *Vehicular Networking Conference (VNC)*. IEEE, 2014, pp. 25–32.
- [13] A. Singh, “Restricted usage of anonymous credentials in VANET for misbehavior detection,” Master’s thesis, University of Applied Sciences, Frankfurt am Main, Germany, Jun. 2012.
- [14] J. Guo, J. P. Baugh, and S. Wang, “A group signature based secure and privacy-preserving vehicular communication framework,” in *2007 Mobile Networking for Vehicular Environments*. IEEE, 2007, pp. 103–108.
- [15] C. Höfer, J. Petit, R. Schmidt, and F. Kargl, “Popcorn: Privacy-preserving charging for mobility,” in *Proceedings of the 2013 ACM workshop on Security, privacy & dependability for cyber vehicles*. ACM, 2013, pp. 37–48.
- [16] D. Angermeier, A. Kiening, and F. Stumpf, “PAL - privacy augmented LTE: A privacy-preserving scheme for vehicular LTE communication,” in *Proceeding of the tenth ACM international workshop on Vehicular inter-networking, systems, and applications*. ACM, 2013, pp. 1–10.

The case for announcing pseudonym changes

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

Jonathan Petit
Department of Computer Science
University College Cork
Email: j.petit@cs.ucc.ie

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

Abstract—The use of pseudonyms is expected to protect vehicular networks against location tracking. However, a lack of recommendations for pseudonym change strategies is a critical missing component in vehicular communication standards. Through a rational discussion of related work we identify the achievable effectiveness against different observer classes as a major determining factor for the selection of safe and effective pseudonym changes strategies. We argue that announcing pseudonym changes serves to make pseudonym changes safe and justifiable while a vehicle is on the road. The achievable level of unlinkability of vehicular communication still serves to increase uncertainty for medium sized observers and helps to increase the cost of surveillance for a potential attacker.

I. INTRODUCTION

The exchange of vehicle data through vehicular communication is expected to enable new driver assistance services that can enhance safety-of-life services in vehicles. The potential for such enhancements is a major motivation for efforts to deploy vehicular communication infrastructure. However, the continuous exchange of vehicle data enables efficient mass surveillance of vehicle movements and indirectly enables tracking of vehicle drives and passengers.

Including the protection of privacy in the design of vehicle to vehicle (V2V) communication solutions is a worthwhile goal for user acceptance and in some jurisdictions might be a fundamental requirement for deployment. Academic research and major contemporary standardization have adopted the protection of privacy as a goal in vehicular communication. This pertains primarily the attribute of unlinkability of V2V communications for the purpose of making location tracking impossible or at least more difficult. Additional relevant attributes can be the general minimization of information disclosure and offering protected ways to revoke privacy and authorizations in cases of abuse.

The security and privacy features proposed in recent versions of the IEEE 1609 and ETSI C-ITS families of standards [1], [2] foresee the use of digital signatures for broadcast authentication of V2V messages. These signatures are based on asymmetric key cryptography and authenticated through certificates issued by trusted authorities. Such a solution offers integrity, authenticity, and non-repudiation with acceptable levels of communication and computation overhead. To protect privacy, in the sense of avoiding linkability, both standards allow sets of multiple valid certificates to be issued for vehicles. These certificates serve as pseudonymous identities and can allow a vehicle to break linkability by changing a short term pseudonymous identity. Unique long-term identities are issued to vehicles in addition to short term pseudonymous

identities to enable privileged operations such as pseudonym refill procedures.

Depending on the characteristics of the certificate issuance method, this kind of pseudonymous authentication scheme can offer revocable privacy [3] or even reasonable levels of true anonymity [4]. However, this implies that the use of pseudonyms is in fact effective to break linkability. This assumption highly depends on the pseudonym change strategy that is used by the vehicles and additionally depends on the coverage level of an eavesdropping attacker.

II. RELATED WORK

Pseudonymous authentication with digital signatures is widely recognized as a suitable solution for privacy preserving broadcast communication in vehicular communication. Yet, specific guidelines on reasonable pseudonym change strategies are lacking. Major standardization efforts offer all the required infrastructure for pseudonymous authentication, but avoid specific recommendations regarding pseudonym change strategies.

In field-operational-tests (FOTs) and publications surrounding the relevant standardization efforts we commonly find assumptions of periodic pseudonym change strategies. In the context of early solutions in the context of IEEE 1609.2 we find estimates of pseudonym changes periods of around 5 minutes [5], [6].

A PKI model proposed within the CAR 2 CAR Communication Consortium has influenced recent FOTs by generally assuming the availability of certificates with multiple overlapping validity periods. This allows flexible change strategies [7]. Yet the proposal avoids specific recommendations, instead calling for standardization of boundaries without providing further suggestions:

The pseudonym change strategy and frequency is out of scope of this work, since we consider it as a feature specific to manufacturers. For security and effectiveness reasons, we only advocate to standardize boundaries of maximum and minimum frequency.

In academic literature we find recommendations of periodic pseudonym change strategies with time periods between 1 minute [8] and 10 minutes [9]. A recent survey of pseudonymity schemes for vehicular networks [10] covers a multitude of strategies, classifying them into 6 categories

- 1) Fixed time change (periodic)
- 2) Random change

- 3) Silent period between change
- 4) Vehicle-centric
- 5) Density-based
- 6) Collaborative (synchronous) change

A ranking of performance characteristics is not included in the aforementioned survey due to lack of universal privacy impacts metrics and lack of suitable quantifications of side-effects on safety and scalability. These problems are identified as future work.

III. DISCUSSION

It is useful to narrow the solution space by identifying what can realistically be achieved in practice. The primary boundary is defined by the attacker model that informs the effectiveness of pseudonym changes. Any pseudonym change strategy is ineffective if an attacker can link different pseudonyms through simple observation. The main design criteria for effective pseudonym change strategies are assumptions about the attacker coverage and the consequences thereof.

A second boundary is imposed on pseudonym change strategies by the fact that side-effects on safety and scalability of V2V communication should not affect safety-of-life services in negative ways. We assume that concerns about service quality for safety-of-life applications in vehicular communication networks will take precedence over privacy considerations.

Investigating achievable goals against attackers and considering negative impact on service quality as unacceptable, provides tight bounds for the solution space. Additionally some techniques that appear detrimental to unlinkability become acceptable, once fundamental constraints of attacker uncertainty and application service quality become apparent.

A. Local observer

Intuitively it appears desirable to change pseudonyms frequently and unlinkably. Nevertheless, studies of data plausibility checks [11] have demonstrated that - even under the assumption of perfect unlinkability - protection against tracking by entities in local communication range is futile. This is due to the fact that vehicles continuously broadcast their precise position and trajectories. These announcements are known as Basic Safety Messages (BSM) or Cooperative awareness Messages (CAM) and represent a core feature of vehicular communication. In fact, it would be counterproductive to aim for location privacy against vehicles in local communication range, because achieving local tracking is the fundamental goal of these broadcasts. It is a feature of vehicular communication to create authenticated linkability for local entities.

B. Global observer

Protection against a global all-seeing attacker is practically impossible due to the same considerations as above. An attacker with universal coverage can create linkability through observations of all local BSM and CAM messages. Even under the absence of any other identifiers, interpretation of the positions and trajectories will enable effective tracking for global observers [12].

Silent periods and mix zones are effective techniques to create uncertainty even for a global all-seeing observer.

However, the use of silent periods is not acceptable while a vehicle is participating in traffic. The introduction of silent periods can degrade the quality of service of important safety-of-life applications, such as intersection collision avoidance applications [13].

Proposals exist to introduce cryptographic silent zones [14], which can protect against passive global observers. Nevertheless, there will always be a degradation of service while enrolling newly arriving vehicles into cryptographic silent zones. All active entities in a cryptographic silent zone need to be supplied with valid cryptographic key material to be able to decrypt position beacons of neighboring vehicles. Furthermore, an attacker can participate actively in a mix zone to receive relevant key material. Additional assumptions about the availability of supporting infrastructure, such as road side units (RSU), limit the practical applicability of cryptographic mix zones.

C. Medium observer

The previous sections indicate that pseudonym change strategies can not and should not be effective against local attackers. Furthermore, techniques to provide unlinkability against global observers, such as silent zones, would have a negative impact on service quality. This negative impact might be small, but will be unacceptable in the context of safety-of-life applications.

This leaves attackers with gaps in coverage as the only model that pseudonym changes can reasonably be effective against. Such protection against medium sized observers is still useful. We expect a large class of potential attackers to have stationary listening stations with non-perfect coverage. The goal is to maximize uncertainty for the attacker and maximize the cost of effective location tracking. A comparable model is the protection of metadata through onion routing, which can effectively only increase the resources required for successful surveillance while not guaranteeing perfect protection of metadata against attackers with very broad or strategic observation capabilities over a network.

D. Safe and effective pseudonym changes

In the discussion above we see that the concern to avoid degradation of service quality in safety-of-life applications is a strong boundary on effective pseudonym change strategies. Essentially the only time when it is perfectly safe to change a pseudonym is when a vehicle is parked. Consequently, the C2C-CC PKI model [7] supports and suggests "per-trip" pseudonym changes. It appears likely that vehicle manufacturers will utilize this strategy to prioritize safety over privacy. However we expect that this pseudonym change strategy is practically ineffective against all but the most primitive stationary observers.

On the other hand, the fact that it is impractical and even counterproductive to protect against tracking by local observers is not an exclusively negative result. This insight makes it possible to consider techniques that would otherwise intuitively appear unjustifiable. It is not necessary to locally hide the fact that a pseudonym change is, was, or will be performed. Instead, it is justifiable and useful to pre-announce pseudonym changes. Pre-announcements help to ensure the continuity of

local cooperative awareness and can support the elimination of cryptographic packet loss [15]. Such pseudonym change announcements can further be used as an enabler for the category of collaborative pseudonym change schemes, as classified in [10]. The synchronization of upcoming pseudonym changes among multiple vehicles serves to increase uncertainty in medium sized observers. Otherwise, without any mix partners between two observation points, it would become easy for a medium sized observer to identify which vehicle performed a pseudonym change. This generally represents the approach of a mix-context pseudonym change strategy [16].

IV. CONCLUSION

The lack of recommendations for pseudonym change strategies is a critical missing component in current plans for vehicular communication standards. Through a rational discussion of related work we identify that ineffectiveness against certain observer classes is a major boundary for the selection of pseudonym changes strategies. Nevertheless, unlinkability of V2V communication through pseudonyms remains a realistically achievable goal against medium sized observers.

We identify the potential for degradation of service quality in safety-of-life applications as a major limitation on pseudonym change strategies. Our considerations lead us to dismiss silent periods due to concerns about degradations of vehicle safety. Furthermore, we point out that pseudonym changes do not need to be perfectly unlinkable against local observers. In fact, local linkability is a goal of cooperative awareness. Thus, local pre-announcements of pseudonym changes can be justifiable. This technique can serve to make pseudonym changes safe while a vehicle is on the road by avoiding cryptographic packet loss and supporting continuous authenticated awareness.

REFERENCES

- [1] IEEE, "IEEE 1609.2v2 - Standard for Wireless Access in Vehicular Environments (WAVE) - Security Services for Applications and Management Messages," 2011.
- [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] F. Schaub, F. Kargl, Z. Ma, and M. Weber, "V-tokens for conditional pseudonymity in vanets," in *Wireless Communications and Networking Conference (WCNC), 2010 IEEE*. IEEE, 2010, pp. 1–6.
- [4] D. Forster, F. Kargl, and H. Lohr, "Puca: A pseudonym scheme with user-controlled anonymity for vehicular ad-hoc networks (vanet)," in *Vehicular Networking Conference (VNC), 2014 IEEE*. IEEE, 2014, pp. 25–32.
- [5] 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.
- [6] W. Whyte, A. Weimerskirch, V. Kumar, and T. Hehn, "A security credential management system for v2v communications," in *Vehicular Networking Conference (VNC), 2013 IEEE*. IEEE, 2013, pp. 1–8.
- [7] N. Bismeyer, H. Stubing, E. Schoch, S. Gotz, J. P. Stotz, and B. Lonc, "A generic public key infrastructure for securing car-to-x communication," in *18th ITS World Congress*, 2011.
- [8] M. Raya, P. Papadimitratos, and J.-P. Hubaux, "Securing vehicular communications," *IEEE Wireless Communications Magazine, Special Issue on Inter-Vehicular Communications*, vol. 13, no. LCA-ARTICLE-2006-015, pp. 8–15, 2006.
- [9] J. J. Haas, Y.-C. Hu, and K. P. Laberteaux, "Design and analysis of a lightweight certificate revocation mechanism for vanet," in *Proceedings of the sixth ACM international workshop on Vehicular InterNetworking*. ACM, 2009, pp. 89–98.
- [10] J. Petit, F. Schaub, M. Feiri, and F. Kargl, "Pseudonym schemes in vehicular networks: a survey," to appear.
- [11] 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.
- [12] 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.
- [13] 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.
- [14] 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.
- [15] M. Feiri, R. Pielage, J. Petit, N. Zannone, and F. Kargl, "Pre-distribution of certificates for pseudonymous broadcast authentication in vanet," in *Vehicular Technology Conference, 2015. VTC2015-Spring. IEEE 81th. IEEE*, to appear.
- [16] M. Gerlach and F. Guttler, "Privacy in vanets using changing pseudonyms-ideal and real," in *Vehicular Technology Conference, 2007. VTC2007-Spring. IEEE 65th. IEEE*, 2007, pp. 2521–2525.

Assessing Artificially Caused Congestion on Urban Scenarios: A Case Study on Luxembourg Sumo Traffic

Lautaro Dolberg*, Jérôme François*[†], Thomas Engel*

*SnT (Interdisciplinary Centre for Security, Reliability and Trust), University of Luxembourg, Luxembourg
Email: lautaro.dolberg@uni.lu, thomas.engel@uni.lu

[†]INRIA Nancy Grand Est, France, email: jerome.francois@inria.fr

Abstract—Position information is essential for many vehicular applications like traffic information and route planning but also for enabling vital network routing services. However, the accuracy of the latter is highly dependent of a precise location service which might be altered by attackers forging falsified position data. In this paper, we assess the impact of congestion caused by both circumstantial conditions and an attacker injecting fake coordinates and propose a methodology for tracking and removing location spoofing anomalies. We perform this analysis under the most general conditions, hence we don't control of the communication infrastructure or additional hardware nor software at the client side. Our system uses aggregation of location information to compute relevant stability metrics which may reveal anomalous events.

I. INTRODUCTION

Positioning systems are considered a crucial service in vehicular navigation and safety applications. Typically, specific routing protocols may rely on vehicle positions to infer the forwarding path of packets [1]–[3], services are qualified as Vehicle-to-Vehicle and Vehicle-to-Infrastructure. In some cases, it is possible to rely on the vehicle movements for delivering packets, according to authors of [4]. Simultaneously, a fast paced growth has occurred on the application market for improving the driving experience. These applications offer services such as trip planning considering road conditions (e.g. LuxTraffic [5]), alert broadcasting and car-sharing services based on commuters patterns. Most of these applications rely on a combination of human interaction and technological improvements. In particular, the high penetration of mobile devices equipped with positioning services such as GPS has contributed enormously to the popular adoption of such applications. Therefore, vehicular-based applications can be marketed and deployed without specific investment on infrastructure like roadside units or embedded on-board units.

The infrastructure leverages the communication channels between the users and the service, which can be protected by standard security protocols. However, such mechanisms do not prevent attack in crowd-sourced applications since the access is usually open and it is rather simple to create a large

number of fake accounts, also affecting negatively blacklisting as a counter measure. Injecting a large number of vehicles leads to have some of them at plausible positions which limits the detection on such observations.

The services mentioned above require a very reliable database of vehicle positions. Those databases can be poisoned with forged positions reported by attackers (location spoofing attack). Thus, an attacker can fake traffic congestion or alter traffic routing systems [6]. While most of popular vehicle based applications like traffic congestion avoidance or event network level routing [7] consider aggregated data, this kind of attacks are still valid by reporting a large number of spoofed or faked vehicle positions.

In this paper, we show an application of our methodology for tackling position spoofing to the urban scenario of Luxembourg city. It is inspired from works relying on movement plausibility verifications [8]–[10] but enhances the scalability by leveraging aggregation which so avoids the costly verification of the movements of each individual vehicle. In fact, vehicles are grouped within different areas of various sizes of a map depending on the traffic density using Multidimensional Aggregation Monitoring (MaM) [11], [12]. As a contribution, we study dedicated metrics that are applied to such aggregated representations to verify the plausibility of global vehicle repartition on a map and the associated movement dynamic. Unlike standard approaches relying on individual car movements, our method focuses large datasets such as those which can be collected from navigation systems and crowd-sourced applications.

This paper is structured as follows. Section II describes our location spoofing detection and recovery method. Section III details the datasets used in the evaluation in Section IV. Finally, Section V concludes the paper.

II. AGGREGATION BASED DETECTION

A. MaM for vehicle locations

In this work, we decided to model an area as bounded intervals. In fact, an area is defined as a bi dimensional space using

Cartesian coordinates to represent the longitude and latitude (altitude is not considered). Hence, the areas are represented as rectangles defined by two points: $((X_1, Y_1) : (X_2, Y_2))$. The most specific representation is an interval of a single point: $X_1 = X_2$ and $Y_1 = Y_2$.

Each node represents a specific rectangular area defined as:

- 1) X dimension range: (X_1, X_2) where $X_1 \in \mathfrak{R}$ and $X_2 \in \mathfrak{R}$
- 2) Y dimension range: (Y_1, Y_2) where $Y_1 \in \mathfrak{R}$ and $Y_2 \in \mathfrak{R}$
- 3) Percentage of vehicles in the corresponding area excluding its children: $vol \in \mathfrak{R}$
- 4) Cumulated percentage of vehicles in the corresponding area (including its children): $acc_vol \in \mathfrak{R}$

The corresponding area of a node is actually embedded into the area associated to its parent node.

As highlighted before, a hierarchical model has to be built over such dimensions. In order to keep the advantages of *MaM*, the areas are not fixed in advance as we could have done using a grid-based approach. Rectangles are built regarding the percentage of activity which, in this case, represents the percentage of vehicles. To do so, the areas are created on the fly by assembling points together. For example, assuming two individual points as $((X_1, Y_1) : (X_1, Y_1))$ and $((X_2, Y_2) : (X_2, Y_2))$, this will entail the creation of the area: $((mix(X_1, X_2), mix(Y_1, Y_2)) : (max(X_1, X_2), max(Y_1, Y_2)))$. When the third point has to be added, either it falls into this created area which implies the creation of an embedded area, or it falls outside creating a parent area. This process continues until no new data points have to be added. Hence, this automatically creates a hierarchy between areas which is not known a priori unlike IP subnetworks. In the meantime, each area maintains a counter about the number of vehicles it contains. Aggregation is then performed from the leaves to the root. For each node, if $vol > \alpha$, the node is kept, otherwise it is discarded and its vol value is added to the one of its parent node. This aggregation process is the standard one of *MaM* which is fully described in [11].

B. Metrics

We propose a method for analyzing series of aggregated trees based on stability over time. The intuitive idea behind is similar to plausibility check by considering that vehicles movements from one area to another are bounded by physical properties but also depending on the traffic conditions. To achieve that, we defined a stability metric that captures the dynamic of the traffic, i.e. the movement of the car, by comparing consecutive vehicle positions. Observing the variation of the stability when no attack occurs, we are then capable of detecting future abnormal variations. The aggregation helps in considering the global road traffic dynamic. In addition, this avoids to detect isolated deviant behaviors which are not necessary malicious and, in fact, does not have an impact on the large scale location-based service that relies on mass information. From a general point of view, stability based




Fig. 1: Area comparison

detection is equivalent to plausibility checks using aggregated views.

As explained, the stability captures the dynamic of the car. It is a bounded value between 0 and 1. Assuming cars which are blocked in a traffic jam, the stability will be very high. On the contrary, cars moving very fast entails a low stability. Therefore, the goal is to evaluate the stability of the traffic (and so vehicles positions) between consecutive time windows. To do so, it is firstly required to define the stability between two nodes, $n1$ and $n2$. It is higher if the distance between the associated areas ($Area_Distance(n1, n2)$) are close, if the overlap between them is high ($Area_Com(n1, n2)$) and if the number of vehicles is similar:

$$\begin{aligned}
 Stability(n1, n2) = & \eta \times Area_Com(n1, n2) + \\
 & \psi \times Area_Distance(n1, n2) + \\
 & \gamma \times (1 - |n1.vol - n2.vol|)
 \end{aligned} \quad (1)$$

The common space is evaluated regarding the ratio between the intersected rectangular area and the merged rectangular area: $Area_Com(n1, n2) = \frac{\cap(n1, n2)}{\cup(n1, n2)}$

The distance is evaluated regarding the centroid of the areas and normalized regarding the diagonal distance of the monitored map: $Area_Distance(n1, n2) = 1 - \frac{|n1.centroid - n2.centroid|}{DIAG_DISTANCE}$

The common area and centroids can be easily retrieved from the coordinates and Figure 1 illustrates them. Evidently, areas are not always intersected and in such a case, $Area_Com(n1, n2) = 0$. Without specific knowledge, $\eta = \psi = \gamma = 1/3$ which weights equally each factor.

The objective is to compare the stability between t_i and a previous tree t_j . To achieve that, the individual stability of each node $n_i \in t_i$ is calculated. This is calculated by comparing with the most similar node in t_j which is $mostsim(n_i, t_j) \in t_j$. $\forall n_j \in t_j, Stability(n_i, mostsim(n_i, t_j)) \geq Stability(n_i, n_j)$.

Finally, to globally assess the abnormality of a tree t_i , the average stability over all nodes is considered:

$$Avg_stab(t_i, t_j) = \frac{\sum_{n_i \in t_i} Stability(n_i, mostsim(n_i, t_j))}{|t_i|} \quad (2)$$

III. DATA SET

A. Traffic Simulator

In order to perform realistic experiments, the simulator SUMO (Simulation of Urban MObility) [13] was used in order to generate vehicle traffic flows in an urban environment. In particular, we used as main scenario Luxembourg City, a detailed map including bus stops, exclusive lanes and realistic traffic conditions^{1 2}.

B. Luxembourg SUMO Traffic Scenario

Our evaluation is based on the Luxembourg SUMO Traffic Scenario (LuST) as a main scenario for simulating traffic. Our choice is based on the supported capabilities offered by this scenario. These capabilities include: support for free flow and congested lanes, multi-modal traffic and avoiding teleportations as much as possible.

C. Traffic Simulation

For simulating traffic we choose three versions of LuST, varying the mobility patterns. As a baseline, we use the original traffic pattern included in the scenario. It consists in three main peaks of traffic defined by an activity generation included in SUMO. The three peaks include one large peak in the morning, and two minor peaks at noon and in the evening.

From this scenario we derived two another scenarios. The first one, consisted in adding duplicated trips of already existing trips in the scenario original configuration. Hence, congestion can be increased gradually. The second, consists in injecting random vehicles as follows. For each vehicle to be injected, a random departure and an arrival street point. In addition, when the car arrives at its destination, its location information is not provided anymore to mimic a real vehicle. A total of 188260 (50000 more than in the baseline) vehicles were injected along the simulation. For being realistic, the cars are not injected simultaneously but also at a random number, every seconds. Around 20833 new cars are injected every hour as highlighted in Table I, which gives also some other characteristics of the datasets. Moreover, each vehicle reports its position every second.

	Baseline	Gradual Congestion	Injection
Total Cars	138260	188260	188260
Mean Speed	11.17	9	6
Mean Car Trip Time	11 min	18 min	22 min

TABLE I: Simulation Performance in numbers

D. Malicious Traffic Data Generation & Injection

Malicious traffic was injected in order to extend each original dataset during our experiments. The main principle consists into injecting faked reported positions of vehicles in a predefined area which thus mimics spoofed vehicles in this area. Therefore, the following parameters are used:

¹<https://github.com/lcodeca/LuSTScenario>

²http://www.vehicularlab.uni.lu/?page_id=264

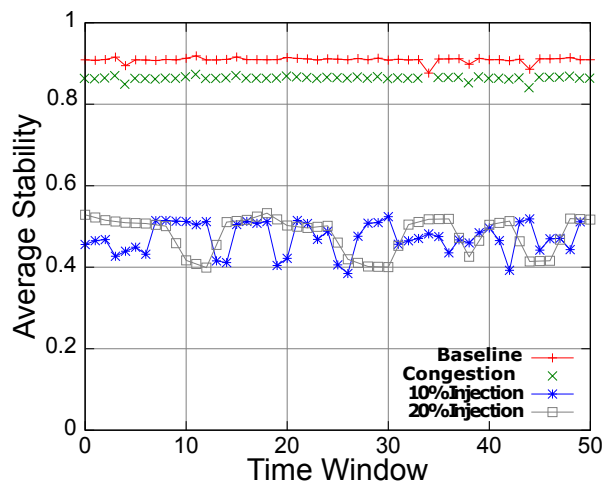


Fig. 2: Average Stability with various attack level, $\beta = 600$ seconds

- Time: the time interval to perform the data generation.
- Center: the center of the area where faked vehicles are injected
- Radius: the radius size of the area where faked vehicles are injected
- Volume of cars: overall amount of injected vehicles.
- Frequency: proportion of vehicles per time unit to be injected regarding the number of normal cars.

The injection takes in consideration the map topology for creating vehicles at valid locations. Besides, the traffic injection models may generate spoofed location attacks against multiple areas. To achieve that, several locations (center and radius) can be defined. In such a case, the volume of cars and the frequency of data generation are global to all targeted locations. In our experiments, the radius is fixed to 500 meters the number of targeted locations is 2 for the Luxembourg scenarios near the highway ramps. We deliberately choose small values in order to strengthen our approach by considering targeted attacks unlike global attacks that could impact all the map and so be highly visible.

IV. EXPERIMENTAL RESULTS

In this section, experimental results are presented. The first experiment was conducted to evaluate the impact of the attacks on the stability and so to show that a threshold-based technique (θ) is viable. The second experiment assesses the impact of injection within different topologies. The third experiment is dedicated to the localization of the attack and its recovery. Then, this section also highlights the gain of using the aggregated trees compared to individual vehicle positions.

Preliminary experiments help in identifying good parameter values for the aggregation threshold $\alpha = 2\%$ and the sliding window size $S = 5$ for computing stability.

A. Attack Identification

In this experiment, $\beta = 600$ seconds and the objective is to assess the impact of the attack injection level, between 10%

and 20% of normal vehicles. In Figure 2, both the Luxembourg injection scenarios is considered. The curve on top of each graph represents the average stability as defined in equation (2) when no attack occurs (with incremental congestion and the baseline). In such a case values are between 0.85 and 0.95. Logically, when the attack aggressiveness increases, the average stability drops in higher proportion. However, considering 10% of spoofed vehicles, reaching a limit of disturbance for this experiments. On a tiny scale with congestion, the average stability is lowered and so clearly distinguishable from baseline. Therefore, setting $\theta = 0.15$ is enough to detect stealthy attacks independently of the topology.

The variation threshold θ is evaluated in Figure 3 by calculating the True Positive Rate (TPR) and False Positive Rate (FPR). FP are due to random normal traffic which, at some high attack rates, can be interpreted as an alert. However, most of them occur with excessively high attack rates.

The FP represents the real road traffic that will be discarded and so can have an hight impact on monitoring, undetected traffic jam for example. Such examples can be more discussed to show how FP may impact applications. Where we assess a "single" attack, it corresponds to test every individual one separately, and then compute an average, but all are tested. The FPR and TPR are calculated based on the stability value of the leaf nodes. For instance, assuming $\theta = 0.15$ and a normal stability around 0.9 as highlighted in the first phase of figure ??, this means that we are considering an absolute threshold of 0.75. In such a case, the TPR reaches 89% with less than 27% of false positives. FPR can be drastically reduced using a smaller threshold like 0.4 which leads to a TPR equal 75% and FPR equal 4%.

V. CONCLUSION

This paper describes a example of location spoofing detection and recovery mechanism that can be applied to a large scale database of position data such as those used in many vehicular applications. Additionally, we were able to show it on a practical example of the Luxembourg SUMO Traffic Scenario (LuST). In this work we were able to show that

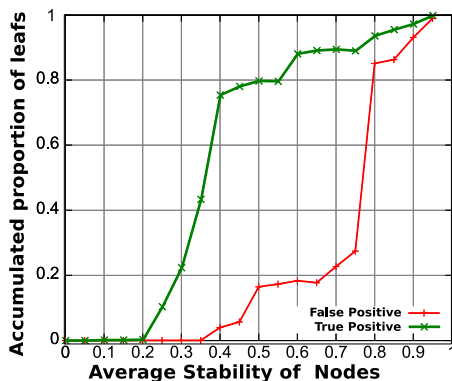


Fig. 3: Threshold detection for simulation containing attacks with up to 10% of increased traffic

average stability metric is a rough indicator of anomaly presence in traffic flows which have been realistically simulated. However, for data collection to be representative we face a scale challenge. Moreover, to re validate our experiments in specific topologies would not add more generality. Different kinds of architectural design of cities are already taken in account for a future work.

Acknowledgement: this work was partially funded by IoT6, a European FP7 funded project under the grant agreement 288445 and MOVE, a CORE project funded by FNR in Luxembourg.

REFERENCES

- [1] L.-D. Chou, J.-Y. Yang, Y.-C. Hsieh, D.-C. Chang, and C.-F. Tung, "Intersection-based routing protocol for vanets," *Wirel. Pers. Commun.*, vol. 60, no. 1, Sep. 2011.
- [2] T. Li, S. Hazra, and W. Seah, "A position-based routing protocol for metropolitan bus networks," in *Vehicular Technology Conference, 2005. VTC 2005-Spring. 2005 IEEE 61st*, 2005.
- [3] C. Lochert, M. Mauve, H. Fussler, and H. Hartenstein, "Geographic routing in city scenarios," *SIGMOBILE Mob. Comput. Commun. Rev.*, vol. 9, no. 1, pp. 69–72, Jan. 2005.
- [4] J. Zhao and G. Cao, "Vadd: Vehicle-assisted data delivery in vehicular ad hoc networks," *Vehicular Technology, IEEE Transactions on*, vol. 57, no. 3, pp. 1910–1922, may 2008.
- [5] A. Kovacheva, R. Frank, and T. Engel, "Luxtraffic: A collaborative traffic sensing system," in *Local & Metropolitan Area Networks (LANMAN), 2013 19th IEEE Workshop on*. IEEE, 2013, pp. 1–6.
- [6] G. Samara, W. Al-Salihy, and R. Sures, "Security issues and challenges of vehicular ad hoc networks (vanet)," in *New Trends in Information Science and Service Science (NISS), 2010 4th International Conference on*, may 2010, pp. 393–398.
- [7] H. Saleet, O. Basir, R. Langar, and R. Boutaba, "Region-based location-service-management protocol for vanets," *Vehicular Technology, IEEE Transactions on*, vol. 59, no. 2, pp. 917–931, 2010.
- [8] T. Leinmüller, C. Maihöfer, E. Schoch, and F. Kargl, "Improved security in geographic ad hoc routing through autonomous position verification," in *International workshop on Vehicular ad hoc networks - VANET*. ACM, 2006.
- [9] T. Leinmuller, E. Schoch, and F. Kargl, "Position verification approaches for vehicular ad hoc networks," *Wireless Communications, IEEE*, vol. 13, no. 5, pp. 16–21, 2006.
- [10] N. Bißmeyer, C. Stresing, and K. Bayarou, "Intrusion detection in vanets through verification of vehicle movement data," in *Vehicular Networking Conference (VNC), IEEE*, 2010.
- [11] L. Dolberg, J. François, and T. Engel, "Efficient multidimensional aggregation for large scale monitoring," in *Proceedings of the 26th international conference on Large Installation System Administration: strategies, tools, and techniques*. USENIX Association, 2012.
- [12] L. Dolberg, J. Francis, and T. Engel, "Tracking spoofed locations in crowdsourced vehicular applications," in *Network Operations and Management Symposium (NOMS), 2014 IEEE*. IEEE, 2014, pp. 1–9.
- [13] 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*, 2011, pp. 55–60.

How Electric Vehicles Can Benefit from Vehicular Networking

Sebastian Schellenberg, Rüdiger Berndt, David Eckhoff, and Reinhard German
Computer Networks and Communication Systems

Dept. of Computer Science, University of Erlangen, Germany

{sebastian.schellenberg, ruediger.berndt, david.eckhoff, reinhard.german}@fau.de

Abstract—Wireless exchange of information between vehicles or vehicles and infrastructure nodes is believed to both improve traffic safety and traffic efficiency. Another trend that can currently be observed is a growing market share of electric vehicles. Unfortunately, electric vehicles still suffer from limited driving ranges and long charging times. In this paper, we not only show possible applications to alleviate these problems by the help of wireless communication but also point to models and tools necessary for their evaluation. We especially discuss how – when a vehicle is aware of its surroundings and able to predict the behavior of other vehicles or traffic lights – it can potentially optimize the process of recuperation, that is, the charging of the battery when braking. Furthermore, we show that vehicular communication can be beneficial to optimize charging strategies or job dispatching in vehicle fleets.

Index Terms—Electromobility; Vehicular Communication; Discrete-Event Simulation;

I. INTRODUCTION

The electrification of a vehicle’s powertrain is an important building block towards the CO₂ reduction goal of the European Union for the year 2050. Unfortunately, Electric Vehicles (EVs) come with some non-negligible drawbacks: the reduced driving range due to limited battery capacity (also known as range anxiety), long charging times compared to the refueling of combustion vehicles, or limited accessibility to charging infrastructure, just to name a few. Considering road traffic, another emerging field can be observed: the deployment of Intelligent Transportation Systems (ITSs), commonly based on modern communication technologies, such as cellular or ad-hoc networks. Whether or not a combination of these technological advances might improve the whole system, and which applications are required, is still subject of research, but was shown to be generally possible [1], [2].

This article will shed light on potential future applications considering electrified road traffic. Further to that, this article will introduce a simulation-based approach allowing the evaluation of the effectiveness of such applications.

The paper is structured as follows. In the next section we will discuss potential applications for Inter-Vehicle Communication (IVC) in the context of electric mobility. Then, in Section III, the necessary components for a simulation-based evaluation are explained. Finally, in Section IV, we draw a conclusion and discuss potential future work.

II. APPLICATIONS

In this section, we give an overview on how wireless vehicular communication can help electric vehicles save energy.

A. Utilizing Recuperation

Recuperation is the process of recovering energy by converting kinetic to electric energy when the car is braking or coasting. During recuperation, the vehicle is decelerated and the battery is charged. Such systems are widely used in today’s electric and hybrid cars.

Usually, an EV comes with multiple levels of recuperation which influence the amount of converted energy and thereby the extent by which a vehicle is slowed down. If a vehicle was able to predict when it had to brake and also the position before which it had to be stopped, it could maximize the converted energy by setting the optimal recuperation level.

One application to maximize recuperation energy by utilizing IVC are Green Light Optimal Speed Advisory (GLOSA) systems, where traffic lights communicate with vehicles to announce green and red phases [2], [3].

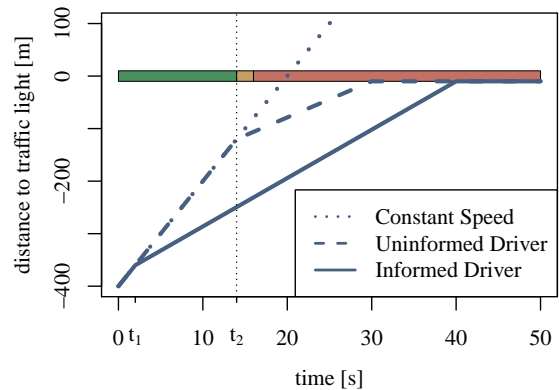


Figure 1. Trajectories of different scenarios at traffic lights where stopping can not be avoided

If a vehicle knows when a traffic light switches from green to red (or vice versa), it can calculate if it can reach the intersection in time to pass it. Otherwise, it can automatically set the highest recuperation level to get the highest energy return flow possible. Another positive side effect is the reduced stress

on the brake discs, which results in less wear and therefore reduced costs.

Figure 1 shows different trajectories for a traffic light scenario for informed and uninformed drivers. Consider a vehicle approaching a traffic light regulated intersection at constant speed: Uninformed driver would only recognize that they are unable to pass the intersection before the next red phase at the moment when they see the lights switch from green to yellow (cf. time t_2 in Figure 1). They consequently have to brake to stop the car in time, leading to a loss of most kinetic energy without converting it to electric energy. An informed driver can calculate the necessary stop beforehand (at time t_1 in Figure 1) and can switch to freewheeling with an optimal recuperation level. This mode recovers the most energy and, ideally, braking can be potentially avoided all together.

While this application only applies to intersections, an automatic adjustment of the recuperation level can also be beneficial in other driving scenarios. Much effort has been put into the development of Emergency Brake Assist (EBA) systems [4]. If a vehicle has to brake suddenly, an EBA system can decelerate a succeeding vehicle automatically if, for example, the preceding car sends a wireless broadcast message when the driver starts braking. This mechanism could be coupled with the recuperation module, and thereby recover energy by supporting the brakes with the engine retarder.

Another possible application utilizing optimal recuperation is an intersection where left yields to right. If a vehicle knows that another car is emerging from the right side on such a crossing, even without seeing it visually but by using IVC, it can brake earlier and gain maximum energy. In an idealized scenario with 100% penetration rate (e.g., see Virtual Traffic Lights [5]) a vehicle knows beforehand whether there is traffic to give way to or not, so it could either avoid braking physically and only use recuperation, or not decelerate at all and cross the intersection.

B. Optimized Scheduling

EVs need to be recharged when the battery is empty, or if the State of Charge (SOC) is too low with regard to an upcoming trip. Considering that, two essential facts need to be taken into account: first, the recharging of an EV takes considerably more time than refueling a combustion vehicle and second, the accessibility to charging infrastructure is still limited. The required charging time is determined by the charging rate and the amount of energy. The charging duration, in turn, depends on the available charging plug (varying from 3.6 to 44 kW) and the battery. Assuming a 12 kWh battery and a plug with 12 kW, charging will take 1 h.

There are several possibilities which might help alleviate these technology-induced impacts on the energy management of EVs. This is where communication comes into play.

The following considerations are particularly suited for (commercial) vehicle fleets, assuming a sufficiently dimensioned charging infrastructure. In order to optimize the charging and job schedule, each EV needs to periodically transmit information about the *expected point of return*, the *expected*

distance to travel, and the *current SOC* to the *scheduler* – i.e., a centralized service.

Based on such information, the system computes an optimal charging scheduling strategy and sends the corresponding information, e.g. the location of an appropriated charging station, back to the vehicles. To decrease the size of waiting queues, waiting times, and charging times, respectively, the *scheduler* will calculate the required energy for each EV and reserve a proper charging plug. In addition, the charging rate could be optimized, for example with respect to the vehicles' upcoming trips. If possible, batteries are recharged gently so that heat generation is kept small and lifetime is increased [6].

Having vehicles periodically report their SOC also allows the scheduler to balance utilization among the vehicle fleet to increase their overall lifetime. This knowledge also allows the scheduler to (re-)assign vehicles to jobs with respect to their current SOC, their mileage, battery age, and the trip distance. Since the necessary information exchange is centralized and not time-critical, cellular-based communication would be the preferred choice for this application.

Communication can also improve the charging process for non-commercial vehicles. Based on the SOC, an in-car navigation system can compute different routes to a given destination. If required, the on-board unit could then contact charging stations along the routes and query them whether a charging plug will be available at the expected time of arrival and makes a reservation. Based on the expected waiting and charging time, one particular route is then recommended to the driver. In a more far fetched scenario, this could also work when there is not a distinct destination but a type of destination, e.g., a supermarket: the vehicle could contact the charging infrastructure of different supermarkets to find out about currently available charging plugs.

C. Platooning

Platooning, that is, multiple cars autonomously driving in formation and close distance, is a promising approach to reduce traffic congestion and increase fuel economy due to reduced air resistance. The close distances between vehicles in a platoon require accurate distance sensors and low-latency communication for safe operation [7]. Wireless and reliable communication between vehicles further allows for platoon management, e.g., joining, leaving, or overtaking maneuvers [8]. In the event of predictable braking maneuvers, the lead vehicle could broadcast its chosen recuperation level and deceleration rate to allow succeeding vehicles to set their optimal recuperation level to maximize battery recharging.

The lead vehicle in a platoon, however, will not encounter reduced fuel or energy consumption and will therefore, in the case of electric vehicles, experience a higher battery drain. To ensure fairness in a platoon, wireless communication could be used to periodically trigger a changing of the lead vehicle in order to balance the SOC between all participants. A simple time-based approach is not sufficient because the SOC is not only depending on the time a vehicle leads the platoon, but also on the maximum capacity of the battery, the average drain




Figure 2. Convoy of vehicles driving in a platoon

(based on engine, vehicle weight, etc.), and the age of the battery [9].

Figure 2 shows an example scenario, where five vehicles drive in a platoon. The leading one obviously consumes the most energy because of the highest air resistance. If we assume the second vehicle to have the currently highest SOC, by appropriate distribution of the individual battery states for instance via WiFi, the leading and the second vehicle could recognize the gap and change the position to balance the battery drains.

III. SIMULATION

It has been shown that simulations are a powerful tool when it comes to performance evaluation of vehicular networking applications [1]. In this section, we want to discuss the requirements and approaches for the simulation of wirelessly communicating electric vehicles.

A. Network Models

It is apparent that to study potential benefits of wireless communication on the battery management of vehicles a model for these networks is required. However, the choice of detail depends on the investigated scenario: for example, if the battery management application can be (strongly) influenced by packet loss, then the used radio and channel model should consider all effects that can decrease the Signal to Noise + Interference Ratio (SNIR) and therefore lead to undecodable. These include but are not limited to path loss, cancellation through multipath propagation [10], and interference from other packets or systems [11]. If low latency is a requirement, the underlying network model must account for MAC and PHY properties that may increase message latency, such as delays for channel reservations, scheduling, CSMA functionality, back-off times and so on. However, in scenarios where these features do not play a significant role, for example when vehicles contact centralized traffic information centers in a non time-critical manner, it can be sufficient to use a very abstract network model that simply delivers packets regardless of delays or throughput.

Many battery management applications require a detailed network model because of their time criticality. If a preceding vehicle informs its succeeding vehicles about a braking event, the latency of this message can have a considerable impact on the effectiveness of the battery management of the preceding vehicle. It was shown that it is necessary to use detailed and exact models to capture these effects accordingly [12].

We therefore recommend the use of widely employed, peer-reviewed, open-source models [13], [14].

B. Mobility and Battery Model

In [15], we presented an evaluated battery model for electric vehicles to be used in micro-simulation. We showed, how the energy consumption and the regarding impact on the SOC of the battery can be derived from the vehicle’s velocity. The power flows in or out of the battery are calculated based on the kinetic power flows, e.g. the power to overcome the air and roll resistance or the power to accelerate the car. The validation of our model using a reference car shows that we are able to reproduce realistic consumptions.

Furthermore, our model comes with two additional modules. The first one generates the energy that is fed back to the battery by an additional range extender, if applicable. Such additional engines are widely used in EVs in order to increase the total range. The second module deals with the calculation of the recuperation energy based on the strength of deceleration and the set recuperation level. Both modules are also validated with the data collected by the reference car [15].

In order to simulate a realistic scenario, an adequate modeling of the scenario, e.g., the traffic, is important. There exist several approaches to increase the realism of a given scenario (assume a city [16]), for example by incorporating information on traffic lights, parking facilities, residential and industrial areas, among others. The used road topology should therefore be based on real map data, for example, obtained from OpenStreetMap (OSM). Start and destination points as well as the number of participating vehicles can then be chosen according to the scenario.

In order to simulate charging procedures, a charging infrastructure has to be included in the scenario, such as the one we presented in [17]. A charging station consists of one or more charging plugs (possibly with different charging rates) that can be placed anywhere on the map. The maximum charging rate also depends on the maximum rate the vehicle’s battery can handle, which is defined in the vehicle model. If all slots are occupied, incoming new vehicles are added to a waiting queue. The number of free slots or waiting vehicles, respectively, can be distributed by the charging stations using cell-based communication in order to balance the load. Lastly, to simulate meaningful scenarios, we create vehicle trips based on input distribution from real commercial vehicle fleets [17].

C. Coupling

In the context of vehicular networking applications it is necessary that both the network and the road traffic simulator are coupled and able to exchange information [14]. In our case, we connect the discrete-event simulator OMNeT++ [18] and the road traffic simulator SUMO [19]. The mobility of the vehicles affects the network topology, hence, the network simulator must move the network nodes according to the vehicles' position in the traffic simulator. But vehicles must also be controllable from the network simulator because, for example, receiving a message from another car or a traffic light can potentially affect the mobility of a vehicle if (based on the information in that message) a certain recuperation level is set or another route should be taken.

IV. CONCLUSION AND FUTURE WORK

In this article, we discussed potential applications for electric vehicles to improve the battery management using wireless vehicular communication. We identified the requirements needed for a simulative study of these scenarios, mainly consisting of accurate battery models, network models, and realistic mobility patterns.

One of the major possibilities to reduce battery drain is the intelligent use of the recuperation level based on the vehicle's surroundings, such as neighboring vehicles, traffic lights, or even road topology. Additionally, vehicle platoons could use inter-vehicle communication to introduce additional benefits to, e.g., balance the State of Charge of all participating vehicles. Lastly, information exchange between vehicles or the infrastructure can be used not only to reduce waiting times and queue lengths at charging stations but also to optimize the job schedule in vehicle fleets.

Future work will mainly focus on detailed evaluation and simulation of listed applications to determine their actual benefit for the battery management of electric vehicles.

V. ACKNOWLEDGMENT

This work is funded by the German national program *Schaufenster Elektromobilität, e-NUE*. We would like to thank *AUDI AG, Ingolstadt, Germany* and *N-ERGIE Aktiengesellschaft, Nürnberg, Germany* for providing the necessary data to conduct this research.

REFERENCES

- [1] 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.
- [2] 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, USA: IEEE, December 2013, pp. 103–110.
- [3] C. Raubitschek, N. Schütze, E. Kozlov, and B. Baker, "Predictive Driving Strategies Under Urban Conditions for Reducing Fuel Consumption Based on Vehicle Environment Information," in *Forum on Integrated and Sustainable Transportation System (FISTS)*. Vienna, Austria: IEEE, July 2011, pp. 13–19.
- [4] N. Gkikas, J. Hill, and J. Richardson, "Towards a driver-centred brake assist system," in *The Proceedings of Braking Conference (Braking 2009)*, The Institute of Mechanical Engineers, St Williams College, York, UK, June 2009, pp. 85–92.
- [5] M. Ferreira and P. d'Orey, "On the Impact of Virtual Traffic Lights on Carbon Emissions Mitigation," *IEEE Transactions on Intelligent Transportation Systems*, vol. 13, no. 1, pp. 284–295, February 2012.
- [6] I. Fernández, C. Calvillo, A. Sánchez-Mirallas, and J. Boal, "Capacity fade and aging models for electric batteries and optimal charging strategy for electric vehicles," *Elsevier Energy*, vol. 60, pp. 35–43, October 2013.
- [7] M. Segata, B. Bloessl, S. Joerer, C. Sommer, M. Gerla, R. Lo Cigno, and F. Dressler, "Towards Inter-Vehicle Communication Strategies for Platooning Support," in *7th IFIP/IEEE International Workshop on Communication Technologies for Vehicles (Nets4Cars 2014-Fall)*. Saint-Petersburg, Russia: IEEE, October 2014, pp. 1–6.
- [8] M. Segata, B. Bloessl, S. Joerer, F. Dressler, and R. Lo Cigno, "Supporting Platooning Maneuvers through IVC: An Initial Protocol Analysis for the Join Maneuver," in *11th IEEE/IFIP Conference on Wireless On demand Network Systems and Services (WONS 2014)*. Obergurgl, Austria: IEEE, April 2014, pp. 130–137.
- [9] L. Serrao, Z. Chehab, Y. Guezennec, and G. Rizzoni, "An Aging Model of Ni-MH Batteries for Hybrid Electric Vehicles," in *IEEE Vehicle Power and Propulsion Conference (VPPC 2005)*. Chicago, IL, USA: IEEE, September 2005, pp. 78–85.
- [10] C. Sommer, D. Eckhoff, R. German, and F. Dressler, "A Computationally Inexpensive Empirical Model of IEEE 802.11p Radio Shadowing in Urban Environments," in *8th IEEE/IFIP Conference on Wireless On demand Network Systems and Services (WONS 2011)*, Bardonecchia, Italy, January 2011, pp. 84–90.
- [11] K. Wessel, M. Swigulski, A. Köpke, and D. Willkomm, "MiXiM – The Physical Layer: An Architecture Overview," in *2nd ACM/ICST International Conference on Simulation Tools and Techniques for Communications, Networks and Systems (SIMUTools 2009): 2nd ACM/ICST International Workshop on OMNeT++ (OMNeT++ 2009)*. Rome, Italy: ACM, March 2009, Article No. 78.
- [12] D. Eckhoff, C. Sommer, and F. Dressler, "On the Necessity of Accurate IEEE 802.11p Models for IVC Protocol Simulation," in *75th IEEE Vehicular Technology Conference (VTC2012-Spring)*, Yokohama, Japan, May 2012, pp. 1–5.
- [13] F. Hagenauer, F. Dressler, and C. Sommer, "A Simulator for Heterogeneous Vehicular Networks," in *6th IEEE Vehicular Networking Conference (VNC 2014), Poster Session*. Paderborn, Germany: IEEE, December 2014, pp. 485–486.
- [14] 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.
- [15] S. Schellenberg, R. Berndt, D. Eckhoff, and R. German, "A Computationally Inexpensive Battery Model for the Microscopic Simulation of Electric Vehicles," in *IEEE 80th Vehicular Technology Conference (VTC2014-Fall)*. Vancouver, BC, Canada: IEEE, September 2014, pp. 1–6.
- [16] S. Uppoor and M. Fiore, "Large-scale Urban Vehicular Mobility for Networking Research," in *3rd IEEE Vehicular Networking Conference (VNC 2011)*. Amsterdam, Netherlands: IEEE, November 2011, pp. 62–69.
- [17] S. Schellenberg, R. Berndt, R. German, and D. Eckhoff, "Evaluating the Electrification of Vehicle Fleets Using the Veins Framework," in *1st OMNeT++ Community Summit 2014*, Hamburg, Germany, September 2014, arXiv:1409.1003.
- [18] 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, Article No. 60.
- [19] 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.

Liste der bisher erschienenen Ulmer Informatik-Berichte

Einige davon sind per FTP von `ftp.informatik.uni-ulm.de` erhältlich

Die mit * markierten Berichte sind vergriffen

List of technical reports published by the University of Ulm

Some of them are available by FTP from `ftp.informatik.uni-ulm.de`

Reports marked with * are out of print

- 91-01 *Ker-I Ko, P. Orponen, U. Schöning, O. Watanabe*
Instance Complexity
- 91-02* *K. Gladitz, H. Fassbender, H. Vogler*
Compiler-Based Implementation of Syntax-Directed Functional Programming
- 91-03* *Alfons Geser*
Relative Termination
- 91-04* *J. Köbler, U. Schöning, J. Toran*
Graph Isomorphism is low for PP
- 91-05 *Johannes Köbler, Thomas Thierauf*
Complexity Restricted Advice Functions
- 91-06* *Uwe Schöning*
Recent Highlights in Structural Complexity Theory
- 91-07* *F. Green, J. Köbler, J. Toran*
The Power of Middle Bit
- 91-08* *V.Arvind, Y. Han, L. Hamachandra, J. Köbler, A. Lozano, M. Mundhenk, A. Ogiwara,
U. Schöning, R. Silvestri, T. Thierauf*
Reductions for Sets of Low Information Content
- 92-01* *Vikraman Arvind, Johannes Köbler, Martin Mundhenk*
On Bounded Truth-Table and Conjunctive Reductions to Sparse and Tally Sets
- 92-02* *Thomas Noll, Heiko Vogler*
Top-down Parsing with Simultaneous Evaluation of Noncircular Attribute Grammars
- 92-03 *Fakultät für Informatik*
17. Workshop über Komplexitätstheorie, effiziente Algorithmen und Datenstrukturen
- 92-04* *V. Arvind, J. Köbler, M. Mundhenk*
Lowness and the Complexity of Sparse and Tally Descriptions
- 92-05* *Johannes Köbler*
Locating P/poly Optimally in the Extended Low Hierarchy
- 92-06* *Armin Kühnemann, Heiko Vogler*
Synthesized and inherited functions -a new computational model for syntax-directed semantics
- 92-07* *Heinz Fassbender, Heiko Vogler*
A Universal Unification Algorithm Based on Unification-Driven Leftmost Outermost Narrowing

- 92-08* *Uwe Schöning*
On Random Reductions from Sparse Sets to Tally Sets
- 92-09* *Hermann von Hasseln, Laura Martignon*
Consistency in Stochastic Network
- 92-10 *Michael Schmitt*
A Slightly Improved Upper Bound on the Size of Weights Sufficient to Represent Any Linearly Separable Boolean Function
- 92-11 *Johannes Köbler, Seinosuke Toda*
On the Power of Generalized MOD-Classes
- 92-12 *V. Arvind, J. Köbler, M. Mundhenk*
Reliable Reductions, High Sets and Low Sets
- 92-13 *Alfons Geser*
On a monotonic semantic path ordering
- 92-14* *Joost Engelfriet, Heiko Vogler*
The Translation Power of Top-Down Tree-To-Graph Transducers
- 93-01 *Alfred Lupper, Konrad Froitzheim*
AppleTalk Link Access Protocol basierend auf dem Abstract Personal Communications Manager
- 93-02 *M.H. Scholl, C. Laasch, C. Rich, H.-J. Schek, M. Tresch*
The COCOON Object Model
- 93-03 *Thomas Thierauf, Seinosuke Toda, Osamu Watanabe*
On Sets Bounded Truth-Table Reducible to P-selective Sets
- 93-04 *Jin-Yi Cai, Frederic Green, Thomas Thierauf*
On the Correlation of Symmetric Functions
- 93-05 *K.Kuhn, M.Reichert, M. Nathe, T. Beuter, C. Heinlein, P. Dadam*
A Conceptual Approach to an Open Hospital Information System
- 93-06 *Klaus Gaßner*
Rechnerunterstützung für die konzeptuelle Modellierung
- 93-07 *Ullrich Keßler, Peter Dadam*
Towards Customizable, Flexible Storage Structures for Complex Objects
- 94-01 *Michael Schmitt*
On the Complexity of Consistency Problems for Neurons with Binary Weights
- 94-02 *Armin Kühnemann, Heiko Vogler*
A Pumping Lemma for Output Languages of Attributed Tree Transducers
- 94-03 *Harry Buhrman, Jim Kadin, Thomas Thierauf*
On Functions Computable with Nonadaptive Queries to NP
- 94-04 *Heinz Faßbender, Heiko Vogler, Andrea Wedel*
Implementation of a Deterministic Partial E-Unification Algorithm for Macro Tree Transducers

- 94-05 *V. Arvind, J. Köbler, R. Schuler*
On Helping and Interactive Proof Systems
- 94-06 *Christian Kalus, Peter Dadam*
Incorporating record subtyping into a relational data model
- 94-07 *Markus Tresch, Marc H. Scholl*
A Classification of Multi-Database Languages
- 94-08 *Friedrich von Henke, Harald Rueß*
Arbeitstreffen Typtheorie: Zusammenfassung der Beiträge
- 94-09 *F.W. von Henke, A. Dold, H. Rueß, D. Schwier, M. Strecker*
Construction and Deduction Methods for the Formal Development of Software
- 94-10 *Axel Dold*
Formalisierung schematischer Algorithmen
- 94-11 *Johannes Köbler, Osamu Watanabe*
New Collapse Consequences of NP Having Small Circuits
- 94-12 *Rainer Schuler*
On Average Polynomial Time
- 94-13 *Rainer Schuler, Osamu Watanabe*
Towards Average-Case Complexity Analysis of NP Optimization Problems
- 94-14 *Wolfram Schulte, Ton Vullings*
Linking Reactive Software to the X-Window System
- 94-15 *Alfred Lupper*
Namensverwaltung und Adressierung in Distributed Shared Memory-Systemen
- 94-16 *Robert Regn*
Verteilte Unix-Betriebssysteme
- 94-17 *Helmuth Partsch*
Again on Recognition and Parsing of Context-Free Grammars:
Two Exercises in Transformational Programming
- 94-18 *Helmuth Partsch*
Transformational Development of Data-Parallel Algorithms: an Example
- 95-01 *Oleg Verbitsky*
On the Largest Common Subgraph Problem
- 95-02 *Uwe Schöning*
Complexity of Presburger Arithmetic with Fixed Quantifier Dimension
- 95-03 *Harry Buhrman, Thomas Thierauf*
The Complexity of Generating and Checking Proofs of Membership
- 95-04 *Rainer Schuler, Tomoyuki Yamakami*
Structural Average Case Complexity
- 95-05 *Klaus Achatz, Wolfram Schulte*
Architecture Independent Massive Parallelization of Divide-And-Conquer Algorithms

- 95-06 *Christoph Karg, Rainer Schuler*
Structure in Average Case Complexity
- 95-07 *P. Dadam, K. Kuhn, M. Reichert, T. Beuter, M. Nathe*
ADEPT: Ein integrierender Ansatz zur Entwicklung flexibler, zuverlässiger kooperierender Assistenzsysteme in klinischen Anwendungsumgebungen
- 95-08 *Jürgen Kehrer, Peter Schulthess*
Aufbereitung von gescannten Röntgenbildern zur filmlosen Diagnostik
- 95-09 *Hans-Jörg Burtschick, Wolfgang Lindner*
On Sets Turing Reducible to P-Selective Sets
- 95-10 *Boris Hartmann*
Berücksichtigung lokaler Randbedingung bei globaler Zielloptimierung mit neuronalen Netzen am Beispiel Truck Backer-Upper
- 95-11 *Thomas Beuter, Peter Dadam:*
Prinzipien der Replikationskontrolle in verteilten Systemen
- 95-12 *Klaus Achatz, Wolfram Schulte*
Massive Parallelization of Divide-and-Conquer Algorithms over Powerlists
- 95-13 *Andrea Mößle, Heiko Vogler*
Efficient Call-by-value Evaluation Strategy of Primitive Recursive Program Schemes
- 95-14 *Axel Dold, Friedrich W. von Henke, Holger Pfeifer, Harald Rueß*
A Generic Specification for Verifying Peephole Optimizations
- 96-01 *Ercüment Canver, Jan-Tecker Gayen, Adam Moik*
Formale Entwicklung der Steuerungssoftware für eine elektrisch ortsbediente Weiche mit VSE
- 96-02 *Bernhard Nebel*
Solving Hard Qualitative Temporal Reasoning Problems: Evaluating the Efficiency of Using the ORD-Horn Class
- 96-03 *Ton Vullingsh, Wolfram Schulte, Thilo Schwinn*
An Introduction to TkGofer
- 96-04 *Thomas Beuter, Peter Dadam*
Anwendungsspezifische Anforderungen an Workflow-Mangement-Systeme am Beispiel der Domäne Concurrent-Engineering
- 96-05 *Gerhard Schellhorn, Wolfgang Ahrendt*
Verification of a Prolog Compiler - First Steps with KIV
- 96-06 *Manindra Agrawal, Thomas Thierauf*
Satisfiability Problems
- 96-07 *Vikraman Arvind, Jacobo Torán*
A nonadaptive NC Checker for Permutation Group Intersection
- 96-08 *David Cyrluk, Oliver Möller, Harald Rueß*
An Efficient Decision Procedure for a Theory of Fix-Sized Bitvectors with Composition and Extraction

- 96-09 *Bernd Biechele, Dietmar Ernst, Frank Houdek, Joachim Schmid, Wolfram Schulte*
Erfahrungen bei der Modellierung eingebetteter Systeme mit verschiedenen SA/RT-
Ansätzen
- 96-10 *Falk Bartels, Axel Dold, Friedrich W. von Henke, Holger Pfeifer, Harald Rueß*
Formalizing Fixed-Point Theory in PVS
- 96-11 *Axel Dold, Friedrich W. von Henke, Holger Pfeifer, Harald Rueß*
Mechanized Semantics of Simple Imperative Programming Constructs
- 96-12 *Axel Dold, Friedrich W. von Henke, Holger Pfeifer, Harald Rueß*
Generic Compilation Schemes for Simple Programming Constructs
- 96-13 *Klaus Achatz, Helmuth Partsch*
From Descriptive Specifications to Operational ones: A Powerful Transformation
Rule, its Applications and Variants
- 97-01 *Jochen Messner*
Pattern Matching in Trace Monoids
- 97-02 *Wolfgang Lindner, Rainer Schuler*
A Small Span Theorem within P
- 97-03 *Thomas Bauer, Peter Dadam*
A Distributed Execution Environment for Large-Scale Workflow Management
Systems with Subnets and Server Migration
- 97-04 *Christian Heinlein, Peter Dadam*
Interaction Expressions - A Powerful Formalism for Describing Inter-Workflow
Dependencies
- 97-05 *Vikraman Arvind, Johannes Köbler*
On Pseudorandomness and Resource-Bounded Measure
- 97-06 *Gerhard Partsch*
Punkt-zu-Punkt- und Mehrpunkt-basierende LAN-Integrationsstrategien für den
digitalen Mobilfunkstandard DECT
- 97-07 *Manfred Reichert, Peter Dadam*
 $ADEPT_{flex}$ - Supporting Dynamic Changes of Workflows Without Loosing Control
- 97-08 *Hans Braxmeier, Dietmar Ernst, Andrea Mößle, Heiko Vogler*
The Project NoName - A functional programming language with its development
environment
- 97-09 *Christian Heinlein*
Grundlagen von Interaktionsausdrücken
- 97-10 *Christian Heinlein*
Graphische Repräsentation von Interaktionsausdrücken
- 97-11 *Christian Heinlein*
Sprachtheoretische Semantik von Interaktionsausdrücken

- 97-12 *Gerhard Schellhorn, Wolfgang Reif*
Proving Properties of Finite Enumerations: A Problem Set for Automated Theorem Provers
- 97-13 *Dietmar Ernst, Frank Houdek, Wolfram Schulte, Thilo Schwinn*
Experimenteller Vergleich statischer und dynamischer Softwareprüfung für eingebettete Systeme
- 97-14 *Wolfgang Reif, Gerhard Schellhorn*
Theorem Proving in Large Theories
- 97-15 *Thomas Wennekers*
Asymptotik rekurrenter neuronaler Netze mit zufälligen Kopplungen
- 97-16 *Peter Dadam, Klaus Kuhn, Manfred Reichert*
Clinical Workflows - The Killer Application for Process-oriented Information Systems?
- 97-17 *Mohammad Ali Livani, Jörg Kaiser*
EDF Consensus on CAN Bus Access in Dynamic Real-Time Applications
- 97-18 *Johannes Köbler, Rainer Schuler*
Using Efficient Average-Case Algorithms to Collapse Worst-Case Complexity Classes
- 98-01 *Daniela Damm, Lutz Claes, Friedrich W. von Henke, Alexander Seitz, Adelinde Uhrmacher, Steffen Wolf*
Ein fallbasiertes System für die Interpretation von Literatur zur Knochenheilung
- 98-02 *Thomas Bauer, Peter Dadam*
Architekturen für skalierbare Workflow-Management-Systeme - Klassifikation und Analyse
- 98-03 *Marko Luther, Martin Strecker*
A guided tour through *Typelab*
- 98-04 *Heiko Neumann, Luiz Pessoa*
Visual Filling-in and Surface Property Reconstruction
- 98-05 *Ercüment Canver*
Formal Verification of a Coordinated Atomic Action Based Design
- 98-06 *Andreas Küchler*
On the Correspondence between Neural Folding Architectures and Tree Automata
- 98-07 *Heiko Neumann, Thorsten Hansen, Luiz Pessoa*
Interaction of ON and OFF Pathways for Visual Contrast Measurement
- 98-08 *Thomas Wennekers*
Synfire Graphs: From Spike Patterns to Automata of Spiking Neurons
- 98-09 *Thomas Bauer, Peter Dadam*
Variable Migration von Workflows in *ADEPT*
- 98-10 *Heiko Neumann, Wolfgang Sepp*
Recurrent V1 – V2 Interaction in Early Visual Boundary Processing

- 98-11 *Frank Houdek, Dietmar Ernst, Thilo Schwinn*
Prüfen von C-Code und Statmate/Matlab-Spezifikationen: Ein Experiment
- 98-12 *Gerhard Schellhorn*
Proving Properties of Directed Graphs: A Problem Set for Automated Theorem Provers
- 98-13 *Gerhard Schellhorn, Wolfgang Reif*
Theorems from Compiler Verification: A Problem Set for Automated Theorem Provers
- 98-14 *Mohammad Ali Livani*
SHARE: A Transparent Mechanism for Reliable Broadcast Delivery in CAN
- 98-15 *Mohammad Ali Livani, Jörg Kaiser*
Predictable Atomic Multicast in the Controller Area Network (CAN)
- 99-01 *Susanne Boll, Wolfgang Klas, Utz Westermann*
A Comparison of Multimedia Document Models Concerning Advanced Requirements
- 99-02 *Thomas Bauer, Peter Dadam*
Verteilungsmodelle für Workflow-Management-Systeme - Klassifikation und Simulation
- 99-03 *Uwe Schöning*
On the Complexity of Constraint Satisfaction
- 99-04 *Ercument Canver*
Model-Checking zur Analyse von Message Sequence Charts über Statecharts
- 99-05 *Johannes Köbler, Wolfgang Lindner, Rainer Schuler*
Derandomizing RP if Boolean Circuits are not Learnable
- 99-06 *Utz Westermann, Wolfgang Klas*
Architecture of a DataBlade Module for the Integrated Management of Multimedia Assets
- 99-07 *Peter Dadam, Manfred Reichert*
Enterprise-wide and Cross-enterprise Workflow Management: Concepts, Systems, Applications. Paderborn, Germany, October 6, 1999, GI-Workshop Proceedings, Informatik '99
- 99-08 *Vikraman Arvind, Johannes Köbler*
Graph Isomorphism is Low for ZPP^{NP} and other Lowness results
- 99-09 *Thomas Bauer, Peter Dadam*
Efficient Distributed Workflow Management Based on Variable Server Assignments
- 2000-02 *Thomas Bauer, Peter Dadam*
Variable Serverzuordnungen und komplexe Bearbeiterzuordnungen im Workflow-Management-System ADEPT
- 2000-03 *Gregory Baratoff, Christian Toepfer, Heiko Neumann*
Combined space-variant maps for optical flow based navigation

- 2000-04 *Wolfgang Gehring*
Ein Rahmenwerk zur Einführung von Leistungspunktsystemen
- 2000-05 *Susanne Boll, Christian Heinlein, Wolfgang Klas, Jochen Wandel*
Intelligent Prefetching and Buffering for Interactive Streaming of MPEG Videos
- 2000-06 *Wolfgang Reif, Gerhard Schellhorn, Andreas Thums*
Fehlersuche in Formalen Spezifikationen
- 2000-07 *Gerhard Schellhorn, Wolfgang Reif (eds.)*
FM-Tools 2000: The 4th Workshop on Tools for System Design and Verification
- 2000-08 *Thomas Bauer, Manfred Reichert, Peter Dadam*
Effiziente Durchführung von Prozessmigrationen in verteilten Workflow-Management-Systemen
- 2000-09 *Thomas Bauer, Peter Dadam*
Vermeidung von Überlastsituationen durch Replikation von Workflow-Servern in ADEPT
- 2000-10 *Thomas Bauer, Manfred Reichert, Peter Dadam*
Adaptives und verteiltes Workflow-Management
- 2000-11 *Christian Heinlein*
Workflow and Process Synchronization with Interaction Expressions and Graphs
- 2001-01 *Hubert Hug, Rainer Schuler*
DNA-based parallel computation of simple arithmetic
- 2001-02 *Friedhelm Schwenker, Hans A. Kestler, Günther Palm*
3-D Visual Object Classification with Hierarchical Radial Basis Function Networks
- 2001-03 *Hans A. Kestler, Friedhelm Schwenker, Günther Palm*
RBF network classification of ECGs as a potential marker for sudden cardiac death
- 2001-04 *Christian Dietrich, Friedhelm Schwenker, Klaus Riede, Günther Palm*
Classification of Bioacoustic Time Series Utilizing Pulse Detection, Time and Frequency Features and Data Fusion
- 2002-01 *Stefanie Rinderle, Manfred Reichert, Peter Dadam*
Effiziente Verträglichkeitsprüfung und automatische Migration von Workflow-Instanzen bei der Evolution von Workflow-Schemata
- 2002-02 *Walter Guttmann*
Deriving an Applicative Heapsort Algorithm
- 2002-03 *Axel Dold, Friedrich W. von Henke, Vincent Vialard, Wolfgang Goerigk*
A Mechanically Verified Compiling Specification for a Realistic Compiler
- 2003-01 *Manfred Reichert, Stefanie Rinderle, Peter Dadam*
A Formal Framework for Workflow Type and Instance Changes Under Correctness Checks
- 2003-02 *Stefanie Rinderle, Manfred Reichert, Peter Dadam*
Supporting Workflow Schema Evolution By Efficient Compliance Checks

- 2003-03 *Christian Heinlein*
Safely Extending Procedure Types to Allow Nested Procedures as Values
- 2003-04 *Stefanie Rinderle, Manfred Reichert, Peter Dadam*
On Dealing With Semantically Conflicting Business Process Changes.
- 2003-05 *Christian Heinlein*
Dynamic Class Methods in Java
- 2003-06 *Christian Heinlein*
Vertical, Horizontal, and Behavioural Extensibility of Software Systems
- 2003-07 *Christian Heinlein*
Safely Extending Procedure Types to Allow Nested Procedures as Values
(Corrected Version)
- 2003-08 *Changling Liu, Jörg Kaiser*
Survey of Mobile Ad Hoc Network Routing Protocols)
- 2004-01 *Thom Frühwirth, Marc Meister (eds.)*
First Workshop on Constraint Handling Rules
- 2004-02 *Christian Heinlein*
Concept and Implementation of C+++, an Extension of C++ to Support User-Defined Operator Symbols and Control Structures
- 2004-03 *Susanne Biundo, Thom Frühwirth, Günther Palm(eds.)*
Poster Proceedings of the 27th Annual German Conference on Artificial Intelligence
- 2005-01 *Armin Wolf, Thom Frühwirth, Marc Meister (eds.)*
19th Workshop on (Constraint) Logic Programming
- 2005-02 *Wolfgang Lindner (Hg.), Universität Ulm , Christopher Wolf (Hg.) KU Leuven*
2. Krypto-Tag – Workshop über Kryptographie, Universität Ulm
- 2005-03 *Walter Guttmann, Markus Maucher*
Constrained Ordering
- 2006-01 *Stefan Sarstedt*
Model-Driven Development with ACTIVECHARTS, Tutorial
- 2006-02 *Alexander Raschke, Ramin Tavakoli Kolagari*
Ein experimenteller Vergleich zwischen einer plan-getriebenen und einer leichtgewichtigen Entwicklungsmethode zur Spezifikation von eingebetteten Systemen
- 2006-03 *Jens Kohlmeyer, Alexander Raschke, Ramin Tavakoli Kolagari*
Eine qualitative Untersuchung zur Produktlinien-Integration über Organisationsgrenzen hinweg
- 2006-04 *Thorsten Liebig*
Reasoning with OWL - System Support and Insights –
- 2008-01 *H.A. Kestler, J. Messner, A. Müller, R. Schuler*
On the complexity of intersecting multiple circles for graphical display

- 2008-02 *Manfred Reichert, Peter Dadam, Martin Jurisch, Ulrich Kreher, Kevin Göser, Markus Lauer*
Architectural Design of Flexible Process Management Technology
- 2008-03 *Frank Raiser*
Semi-Automatic Generation of CHR Solvers from Global Constraint Automata
- 2008-04 *Ramin Tavakoli Kolagari, Alexander Raschke, Matthias Schneiderhan, Ian Alexander*
Entscheidungsdokumentation bei der Entwicklung innovativer Systeme für produktlinien-basierte Entwicklungsprozesse
- 2008-05 *Markus Kalb, Claudia Dittrich, Peter Dadam*
Support of Relationships Among Moving Objects on Networks
- 2008-06 *Matthias Frank, Frank Kargl, Burkhard Stiller (Hg.)*
WMAN 2008 – KuVS Fachgespräch über Mobile Ad-hoc Netzwerke
- 2008-07 *M. Maucher, U. Schöning, H.A. Kestler*
An empirical assessment of local and population based search methods with different degrees of pseudorandomness
- 2008-08 *Henning Wunderlich*
Covers have structure
- 2008-09 *Karl-Heinz Niggl, Henning Wunderlich*
Implicit characterization of FPTIME and NC revisited
- 2008-10 *Henning Wunderlich*
On span- P^{cc} and related classes in structural communication complexity
- 2008-11 *M. Maucher, U. Schöning, H.A. Kestler*
On the different notions of pseudorandomness
- 2008-12 *Henning Wunderlich*
On Toda's Theorem in structural communication complexity
- 2008-13 *Manfred Reichert, Peter Dadam*
Realizing Adaptive Process-aware Information Systems with ADEPT2
- 2009-01 *Peter Dadam, Manfred Reichert*
The ADEPT Project: A Decade of Research and Development for Robust and Flexible Process Support
Challenges and Achievements
- 2009-02 *Peter Dadam, Manfred Reichert, Stefanie Rinderle-Ma, Kevin Göser, Ulrich Kreher, Martin Jurisch*
Von ADEPT zur AristaFlow[®] BPM Suite – Eine Vision wird Realität “Correctness by Construction” und flexible, robuste Ausführung von Unternehmensprozessen

- 2009-03 *Alena Hallerbach, Thomas Bauer, Manfred Reichert*
Correct Configuration of Process Variants in Provop
- 2009-04 *Martin Bader*
On Reversal and Transposition Medians
- 2009-05 *Barbara Weber, Andreas Lanz, Manfred Reichert*
Time Patterns for Process-aware Information Systems: A Pattern-based Analysis
- 2009-06 *Stefanie Rinderle-Ma, Manfred Reichert*
Adjustment Strategies for Non-Compliant Process Instances
- 2009-07 *H.A. Kestler, B. Lausen, H. Binder H.-P. Klenk, F. Leisch, M. Schmid*
Statistical Computing 2009 – Abstracts der 41. Arbeitstagung
- 2009-08 *Ulrich Kreher, Manfred Reichert, Stefanie Rinderle-Ma, Peter Dadam*
Effiziente Repräsentation von Vorlagen- und Instanzdaten in Prozess-Management-Systemen
- 2009-09 *Dammertz, Holger, Alexander Keller, Hendrik P.A. Lensch*
Progressive Point-Light-Based Global Illumination
- 2009-10 *Dao Zhou, Christoph Müssel, Ludwig Lausser, Martin Hopfensitz, Michael Kühl, Hans A. Kestler*
Boolean networks for modeling and analysis of gene regulation
- 2009-11 *J. Hanika, H.P.A. Lensch, A. Keller*
Two-Level Ray Tracing with Recordering for Highly Complex Scenes
- 2009-12 *Stephan Buchwald, Thomas Bauer, Manfred Reichert*
Durchgängige Modellierung von Geschäftsprozessen durch Einführung eines Abbildungsmodells: Ansätze, Konzepte, Notationen
- 2010-01 *Hariolf Betz, Frank Raiser, Thom Frühwirth*
A Complete and Terminating Execution Model for Constraint Handling Rules
- 2010-02 *Ulrich Kreher, Manfred Reichert*
Speichereffiziente Repräsentation instanzspezifischer Änderungen in Prozess-Management-Systemen
- 2010-03 *Patrick Frey*
Case Study: Engine Control Application
- 2010-04 *Matthias Lohrmann und Manfred Reichert*
Basic Considerations on Business Process Quality
- 2010-05 *HA Kestler, H Binder, B Lausen, H-P Klenk, M Schmid, F Leisch (eds):*
Statistical Computing 2010 - Abstracts der 42. Arbeitstagung
- 2010-06 *Vera Künzle, Barbara Weber, Manfred Reichert*
Object-aware Business Processes: Properties, Requirements, Existing Approaches

- 2011-01 *Stephan Buchwald, Thomas Bauer, Manfred Reichert*
Flexibilisierung Service-orientierter Architekturen
- 2011-02 *Johannes Hanika, Holger Dammertz, Hendrik Lensch*
Edge-Optimized \hat{A} -Trous Wavelets for Local Contrast Enhancement with Robust Denoising
- 2011-03 *Stefanie Kaiser, Manfred Reichert*
Datenflussvarianten in Prozessmodellen: Szenarien, Herausforderungen, Ansätze
- 2011-04 *Hans A. Kestler, Harald Binder, Matthias Schmid, Friedrich Leisch, Johann M. Kraus (eds):*
Statistical Computing 2011 - Abstracts der 43. Arbeitstagung
- 2011-05 *Vera Künzle, Manfred Reichert*
PHILharmonicFlows: Research and Design Methodology
- 2011-06 *David Knuplesch, Manfred Reichert*
Ensuring Business Process Compliance Along the Process Life Cycle
- 2011-07 *Marcel Dausend*
Towards a UML Profile on Formal Semantics for Modeling Multimodal Interactive Systems
- 2011-08 *Dominik Gessenharter*
Model-Driven Software Development with ACTIVECHARTS - A Case Study
- 2012-01 *Andreas Steigmiller, Thorsten Liebig, Birte Glimm*
Extended Caching, Backjumping and Merging for Expressive Description Logics
- 2012-02 *Hans A. Kestler, Harald Binder, Matthias Schmid, Johann M. Kraus (eds):*
Statistical Computing 2012 - Abstracts der 44. Arbeitstagung
- 2012-03 *Felix Schüssel, Frank Honold, Michael Weber*
Influencing Factors on Multimodal Interaction at Selection Tasks
- 2012-04 *Jens Kolb, Paul Hübner, Manfred Reichert*
Model-Driven User Interface Generation and Adaption in Process-Aware Information Systems
- 2012-05 *Matthias Lohrmann, Manfred Reichert*
Formalizing Concepts for Efficacy-aware Business Process Modeling
- 2012-06 *David Knuplesch, Rüdiger Pryss, Manfred Reichert*
A Formal Framework for Data-Aware Process Interaction Models
- 2012-07 *Clara Ayora, Victoria Torres, Barbara Weber, Manfred Reichert, Vicente Pelechano*
Dealing with Variability in Process-Aware Information Systems: Language Requirements, Features, and Existing Proposals
- 2013-01 *Frank Kargl*
Abstract Proceedings of the 7th Workshop on Wireless and Mobile Ad-Hoc Networks (WMAN 2013)

- 2013-02 *Andreas Lanz, Manfred Reichert, Barbara Weber*
A Formal Semantics of Time Patterns for Process-aware Information Systems
- 2013-03 *Matthias Lohrmann, Manfred Reichert*
Demonstrating the Effectiveness of Process Improvement Patterns with Mining Results
- 2013-04 *Semra Catalkaya, David Knuplesch, Manfred Reichert*
Bringing More Semantics to XOR-Split Gateways in Business Process Models Based on Decision Rules
- 2013-05 *David Knuplesch, Manfred Reichert, Linh Thao Ly, Akhil Kumar, Stefanie Rinderle-Ma*
On the Formal Semantics of the Extended Compliance Rule Graph
- 2013-06 *Andreas Steigmiller, Birte Glimm, Thorsten Liebig*
Nominal Schema Absorption
- 2013-07 *Hans A. Kestler, Matthias Schmid, Florian Schmid, Dr. Markus Maucher, Johann M. Kraus (eds)*
Statistical Computing 2013 - Abstracts der 45. Arbeitstagung
- 2013-08 *Daniel Ott, Dr. Alexander Raschke*
Evaluating Benefits of Requirement Categorization in Natural Language Specifications for Review Improvements
- 2013-09 *Philip Geiger, Rüdiger Pryss, Marc Schickler, Manfred Reichert*
Engineering an Advanced Location-Based Augmented Reality Engine for Smart Mobile Devices
- 2014-01 *Andreas Lanz, Manfred Reichert*
Analyzing the Impact of Process Change Operations on Time-Aware Processes
- 2014-02 *Andreas Steigmiller, Birte Glimm, and Thorsten Liebig*
Coupling Tableau Algorithms for the DL SROIQ with Completion-based Saturation Procedures
- 2014-03 *Thomas Geier, Felix Richter, Susanne Biundo*
Conditioned Belief Propagation Revisited: Extended Version
- 2014-04 *Hans A. Kestler, Matthias Schmid, Ludwig Lausser, Johann M. Kraus (eds)*
Statistical Computing 2014 - Abstracts der 46. Arbeitstagung
- 2014-05 *Andreas Lanz, Roberto Posenato, Carlo Combi, Manfred Reichert*
Simple Temporal Networks with Partially Shrinkable Uncertainty (Extended Version)
- 2014-06 *David Knuplesch, Manfred Reichert*
An Operational Semantics for the Extended Compliance Rule Graph Language
- 2015-01 *Andreas Lanz, Roberto Posenato, Carlo Combi, Manfred Reichert*
Controlling Time-Awareness in Modularized Processes (Extended Version)

2015-03

Raphael Frank, Christoph Sommer, Frank Kargl, Stefan Dietzel, Rens W. van der Heijden

Proceedings of the 3rd GI/ITG KuVS Fachgespräch Inter-Vehicle Communication (FG-IVC 2015)

Ulmer Informatik-Berichte

ISSN 0939-5091

Herausgeber:

Universität Ulm

Fakultät für Ingenieurwissenschaften und Informatik

89069 Ulm

