Department Informatik

Technical Reports CS-2010-01

Christoph Sommer, Reinhard German and Falko Dressler

Adaptive Beaconing for Delay-Sensitive and Congestion-Aware Traffic Information Systems

January 2010

Adaptive Beaconing for Delay-Sensitive and Congestion-Aware Traffic Information Systems

Christoph Sommer, Reinhard German and Falko Dressler
Computer Networks and Communication Systems
Dept. of Computer Science, University of Erlangen, Germany
`{christoph.sommer,german,dressler}@informatik.uni-erlangen.de`

Abstract—**We present a new car-to-X communication protocol, Adaptive Traffic Beacon (ATB), which supports the exchange of delay-sensitive traffic information in a wide range of scenarios by flexibly adapting to the availability of infrastructure elements as well as to the network load. From previous work, we see that centralized solutions and flooding based approaches each show benefits and drawbacks depending on traffic density, penetration, network utilization, and other parameters. This observation is in line with findings about intelligent transportation systems that have been developed for specific settings. In order to overcome this limitation, we designed ATB to be adaptive in two dimensions: First, the beacon interval is adapted dynamically and, secondly, the protocol can dynamically make use of available infrastructure elements. We concentrate on a Traffic Information System (TIS) with a focus on congestion-aware communication. Simulation experiments clearly demonstrate that ATB performs well in a broad range of settings. It maintains a non-congested wireless channel to prevent collisions during TIS data exchange.**

I. INTRODUCTION

Intelligent transportation systems rely on accurate and timely information about the current road traffic and about possible congestions and accidents. In the last couple of years, many efforts have been undertaken to study quite diverse strategies for Inter-Vehicle Communication (IVC). Traffic Information Systems (TISs) are currently one of the most interesting application domains [1], [2] both from a scientific point of view and as a business case. In this paper, we concentrate on both the collection and the distribution of traffic information in the context of a TIS and the support for delay-sensitive and congestion-aware wireless communication. Our objective is to provide support for intelligent roads or active highways that optimize routing of vehicles [3].

Early inter-vehicle communication protocols concentrated on the establishment of Vehicular Ad Hoc Networks (VANETs), i.e. the application of routing protocols and coordination techniques known from Mobile Ad Hoc Networks (MANETs) [2], in conjunction with pre-deployed infrastructure elements such as Roadside Units (RSUs). The objective was to set up a path between the vehicle and a central server. Another approach to TISs, which is currently the only commercially successful alternative, is the use of cellular networks for the information exchange. Currently available solutions such as TomTom HD rely on 2.5G or 3G communication networks. Recent studies show that special capabilities of 3G/4G networks, especially the availability of multicast communication [4], [5], can be beneficial for the development of TIS applications. However, all these approaches are fully dependent on available network infrastructure elements and only support an efficient downlink to the vehicles. The described TIS approaches rely on a centralized application server that serves as a sink for new traffic information and that also transmits the currently available information (or at least the samples currently relevant to a particular region) back to the participating vehicles. Such a centralized service can become a bottleneck or may not be available in some situations [6], [7].

Therefore, fully decentralized solutions have been investigated. As information flooding is no feasible alternative, geographic flooding and periodic beaconing have been studied. The common idea is to broadcast traffic information to neighboring vehicles, either periodically or triggered by new events [8]. This process can be supported using directed, i.e. geographic, flooding [9]. Furthermore, aggregation and other data pre-processing techniques have been developed to optimize the quality of traffic information and to reduce the necessary communication load [10]–[12].

Obviously, there are many IVC protocols out there, but some aspects are still not solved: First, protocols

have to cope with very dynamic networks and, secondly, available resources have to be coordinated in a self-organizing distributed way, incorporating both infrastructure elements and even centralized information repositories.

In this paper, we show that this complexity can be handled using an adaptive beaconing system. We present the Adaptive Traffic Beacon (ATB) protocol, which supports distributed exchange of traffic information. ATB is designed to use IEEE 802.11p [13] at the MAC layer. Progress beyond state-of-the-art solutions can be broken down into two aspects:

- ATB uses a variable beacon period which dynamically adapts the frequency of information exchange to a wide range of parameters such as vehicle density, vehicles' speed, radio communication reliability and delay.
- In addition to supporting fully decentralized information exchange among participating vehicles, ATB also automatically makes use of available infrastructure, starting with intelligent RSUs that participate in the ATB network up to a centralized Traffic Information Center (TIC) connected to a network of RSUs.

The rest of the paper is organized as follows. Section II summarizes related approaches. Section III introduces ATB and outlines its capabilities w.r.t. adaptive beaconing and on demand incorporation of infrastructure. The protocol is evaluated in Section IV using different scenarios and a variety of metrics. Finally, conclusions are drawn in Section V.

II. RELATED WORK

Inter-Vehicle Communication (IVC) can be categorized into several classes according to the nature of the envisioned application scenarios [7]. For TIS applications, classes of particular interest are safety applications, cooperative driving systems, and driver information systems. Additional classes that can be distinguished are Internet access and point-to-point communication among cars within the network. In general the traffic patterns exhibited by the selected application classes greatly influence the protocol design [6]. In this paper, we concentrate on the first two IVC classes to support the transport of delay-sensitive and congestion-aware traffic information. In the following, we briefly point out approaches described in the literature that also support the exchange of traffic information using IVC protocols. All these solutions take different approaches to dealing with dynamic network

characteristics, e.g. velocity, movement patterns, node density, and employ different communication patterns, e.g. unicast, geocast, and beaconing. Some integrate methods for information aggregation. We also distinguish between infrastructure-based, ad hoc, and hybrid solutions.

Early VANET based solutions relied on a centralized TIC and pre-deployed RSUs, using MANET routing techniques to set up a path between the vehicle and a central server. Experiments have been performed for several routing protocols and configurations, e.g. using DSRor DYMO [14]. The main problem with these solutions is a lack of scalability along multiple dimensions [2]. First, the quality of transmissions decreases with increasing path lengths. Secondly, this approach only works at sufficient node densities, whereby this node density also has an upper bound, as wireless ad hoc communication suffers from high collision probabilities in congested areas. Infrastructure elements such as RSUs help to avoid extreme network congestion, albeit at high operational costs.

Internet connectivity can be supported by Delay Tolerant Network (DTN) related approaches [15] following a *store-carry-forward* communication principle. Furthermore, data muling concepts can be applied to intelligent transportation systems performing carry-and-forward of TIS data between vehicles and dedicated infrastructure nodes [16]. Recently, it has been shown that lightweight RSUs [17], called stationary support units, or repeaters [8] may be used to replace expensive RSUs with permanent backbone connection.

It has been discovered [18] that VANETs may exhibit bipolar behavior, i.e. the network can either be fully connected or sparsely connected depending on the time of day or the penetration rate. Therefore, lightweight information encoding about both target areas and preferred routes in combination with flooding-like dissemination, such as probabilistic diffusion, seems to be an appropriate solution [19]. In this context, probabilistic aggregation techniques have been investigated for maintaining the distributed message stores [20]. This can be further combined to provide optimal aggregation with infrastructure-support (RSU placement) [10].

One of the most sophisticated solutions is the Self-Organizing Traffic Information System (SOTIS) [11]. Its main aspects are information exchange using a specialized MAC protocol as well as storage of information in the form of annotated maps with variable resolution, depending on distance from the current

position and age of the information. This approach has been further elaborated and merged with ideas from the peer-to-peer domain in the *peers on wheels* vision [21]. Conceptually, it is possible to build extremely robust traffic information sharing systems supporting publish/subscribe interfaces managed by a Distributed Hash Table (DHT) which is maintained by the vehicles. Recently, the MobTorrent approach has been published [22], which also provides mobile (torrent-like) Internet access from vehicles using RSUs (building on the ideas of drive-thru Internet [15], but exploiting state-of-the-art data management functions). LOUVRE [23], on the other hand, provides overlay routing in vehicular environments.

Completely different solutions are investigated in the context of broadcast based approaches. Multi-hop broadcast is a promising technique, especially for emergency message propagation with delay bounds [24]. In order to reduce unnecessary broadcast transmissions, directional broadcast (as opposed to pure flooding) can be used [8]. This can be further optimized using geocasting approaches [9].

Beaconing, or 1-hop broadcast, is an inherent feature of most of the discussed systems. For example, neighborhood information is collected by exchanging beacons. The exploitation of periodic information exchange, with special focus on safety applications, has been first analyzed in extensive simulations in [25], showing that with increasing distance, the success ratio decreased quickly. Combined with a position based forwarding strategy, however, the approach could be improved [26]. Most recently, 2-hop beaconing has been described to acquire topology knowledge for opportunistic forwarding using the selected best target forwarder [12]. The main challenge for all of the introduced beacon systems is that they are very sensitive to environmental conditions such as vehicle density and network load. A first adaptive beaconing system was REACT [27]. Based only on neighbor detection, it can skip intervals for beacon transmission to support emergency applications.

Based on current IVC approaches, two trends can therefore be identified for future TIS solutions: peer-to-peer like information management and beaconing techniques. Both are highly distributed and delay-tolerant, but are targeting two different application domains. Peer-to-peer solutions are highly applicable for extremely delay-tolerant applications. On the other hand, beaconing systems are better suited for TIS data exchange with support for delay-sensitive safety applications.
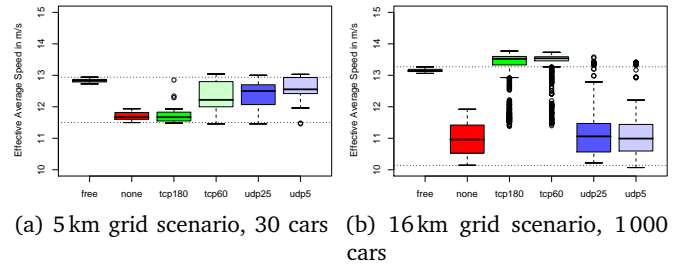

(a) 5 km grid scenario, 30 cars (b) 16 km grid scenario, 1 000 cars

Fig. 1: Comparison of TIS performance between a centralized TIC with MANET routing and decentralized flooding [28]

III. ADAPTIVE TRAFFIC BEACON

In the following section, we outline the system architecture of ATB and introduce the mechanisms for adaptive beacon interval selection and for flexible use of available infrastructure elements.

A. Motivation and Architecture

In previous work, we evaluated the behavior of MANET protocols vs. completely distributed broadcast based operation for TIS applications [28]. It turned out that centralized solutions and flooding based approaches each show benefits and drawbacks depending on a wide range of system parameters. The feasibility, but also the quality, of the transmission depends mainly on the vehicle density and penetration rate. This has also been shown in [18]. So, depending on the scenario, either infrastructure-based or distributed ad hoc solutions can show their strengths.

Figure 1 shows selected results from our earlier study [28]. As can be seen, the average speed of vehicles is optimized by re-routing around traffic congestions based on information from a TIS. Comparing the results for a centralized solution ("tcp180") and a decentralized one ("udp5") shows that depending on the scenario either one is more beneficial.

For our new Adaptive Traffic Beacon (ATB) protocol, which we designed to be adaptive according to the current scenario and traffic conditions, we chose to rely on a beacon system. ATB distributes information about traffic related events, e.g. accident or congestion information, by means of 1-hop broadcasts. These beacon messages are prepared to contain those information elements most relevant to the node (see Section III-D). In order to avoid congestion of the wireless channel while ensuring good information distribution, the interval between two messages is adapted based on two metrics: the perceived channel congestion and the importance
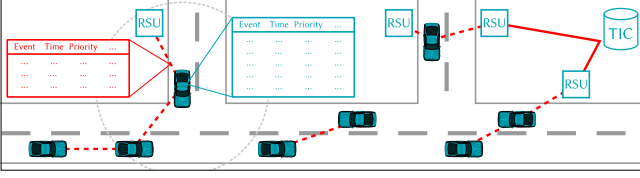
Fig. 2: ATB system architecture

of the message to send (see Section III-B). Figure 2 shows the envisioned system architecture. Vehicles continuously exchange beacon messages containing TIS data. The locally maintained knowledge bases are sorted w.r.t. the message priority, which is a measure of the importance of the message and the estimated utility to other vehicles. Each beacon contains a subset of these entries. Furthermore, infrastructure support can be exploited for improved information exchange.

ATB thus supports two dimensions of adaptivity: First, the traffic parameters related to road traffic and the wireless network are considered (see Section III-B). Secondly, available infrastructure elements are integrated automatically (see Section III-C).

B. Adaptive Beaconing

As described before, ATB provides TIS data exchange by active beaconing, but instead of using fixed beacon periods, the *beacon interval* is dynamically adapted to the currently estimated *channel quality* and the *message priorities*. The main objective is to transmit TIS data as frequently as possible, while making sure the wireless channel does not become congested. The beacon structure, the data format, and the management of the local knowledge base are described in Section III-D.

ATB uses two different metrics to calculate the interval parameter I: the *channel quality C* and the *message priority P*. The relative impact of both parameters is configured using an interval weighting factor w_I that can also be used to calibrate ATB for different MAC protocol variants. The interval parameter I (in the range $0 \ldots 1$) is calculated according to Equation 1.
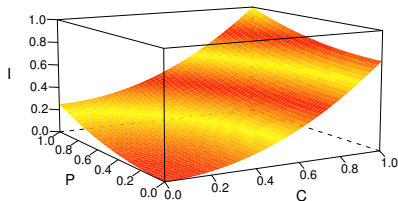

Fig. 3: Interval parameter I for an interval weighting $w_I = 0.75$

Figure 3 shows the behavior of I for $w_I = 0.75$. As can be seen, the interval parameter becomes 1 only for the highest message priority and the best channel quality. In all other cases, I quickly falls to values below 0.5. From the interval parameter, the beacon interval ΔI is then derived (Equation 2; I_{min} and I_{max} represent the minimum and the maximum beacon interval, respectively):

$$I = (1 - w_I) \times P^2 + (w_I \times C^2) \qquad (1)$$
$$\Delta I = I_{min} + (I_{max} - I_{min}) \times I \qquad (2)$$

In the following, we briefly introduce the different parameters affecting the channel quality and the message priority.

a) Channel quality C: The channel quality is estimated by means of three measures. First, the number of neighbors is used to estimate the congestion probability (the more neighbors, the higher the probability for simultaneous transmissions). We model $N = \min\left\{(\frac{\text{\# neighbors}}{\text{max. \# neighbors}})^2; 1\right\}$. The parameter quadratically approaches 1, scaled by a pre-defined maximum.

Besides this estimation, we also take the measured number of collisions into account as $K = 1 - \frac{1}{1 + \text{\# collisions}}$. The objective is to keep the number of observed collisions close to zero, i.e. to ensure a congestion-aware communication.

The third criterion, S, is based on the measured Signal to Noise Ratio (SNR). We modeled $S = \max\left\{0; (\frac{\text{SNR}}{\text{max. SNR}})^2\right\}$. In measurements, is has been shown that the error rate of WiFi communication quickly increases if the SNR drops below 25 dB [29]. Therefore, we configured a maximum SNR (50 dB) so that S already decreases to 0.25 for a SNR equal to 25 dB.

Finally, the channel quality C can be calculated as follows (the factor w_C is used to weight the measured parameters K and S higher than the estimated congestion probability N):

$$C = \frac{N + w_C \times \frac{S+K}{2}}{1 + w_C} \qquad (3)$$

b) Message priority P: We calculate the message priority P as a function of the age of the TIS data, the distance to the event source, the distance to the next RSU, and how well the information is already disseminated. The message priority is calculated for the TIS data with the highest priority in the local knowledge base (see Section III-D).

The message age is accounted by weighting it with the maximum beacon interval I_{max}: $A = \min\left\{(\frac{\text{message age}}{I_{max}})^2; 1\right\}$. The older the information is, the less frequently it should be distributed (bounded by the maximum beacon interval I_{max}).

The next metrics are related to the estimated distance of the vehicle to the event $D_e = \min\left\{(\frac{\text{distance to event}/v}{I_{max}})^2; 1\right\}$ and to the distance to the next RSU $D_r = \max\left\{0; 1 - \sqrt{\frac{\text{distance to RSU}/v}{I_{max}}}\right\}$. Both metrics take the current speed v of the vehicle into account to measure the distance in the form of an estimated travel time.

Finally, the message priority is weighted based on how well its contents are already disseminated. Taking into account how much of the information to be sent was not received via an RSU, this factor is calculated as $B = \frac{1}{1+\# \text{ unknown entries}}$. This measure, which is only used if the current beacon comes from an RSU, ensures that messages are quickly forwarded to the local RSU if it lacks information carried by the vehicle.

The message priority P can now be calculated as follows:

$$P = B \times \frac{A + D_e + D_r}{3} \tag{4}$$

C. Flexible use of infrastructure elements

ATB has been designed keeping in mind the possible exploitation of available infrastructure elements. Thus, deployed RSUs and even central TICs are inherently supported by ATB.

In principle, ATB-enabled RSUs operate similar to ATB-enabled vehicles. They participate in the beaconing process and adapt the beacon interval according to the same rules as described in Section III-B. Thus, an RSU can simply be deployed as a standalone system, e.g. with an attached solar-cell for autonomous operation. This is similar to the concept of stationary support units [10], [17].

As shown in Figure 2, the RSUs can also be connected to a backbone network. This connection is used by ATB to inform other RSUs about received traffic information. In turn, the other RSUs update their local knowledge base accordingly, using the same procedure as when receiving a regular beacon (see Section III-D). We further assume that these RSUs also know their geographic position and the positions of the neighboring RSUs. Therefore, data muling concepts as described in [16] can be realized.

The main difference between RSUs and vehicles is the calculation of the beacon interval. RSUs are not able to estimate their travel time to a traffic congestion. Thus, these metrics are ignored resulting in slightly shorter beacon intervals:

$$P_{RSU} = B \times A \tag{5}$$
$$I_{RSU} = (1 - w_I) \times P_{RSU} + (w_I \times C) \tag{6}$$

Last but not least, the RSUs can be connected to one or more central TICs. A TIC disseminates received TIS data differently than the vehicles and RSUs. Using the available topology information of the connected RSUs, only relevant, i.e. geographically related, information is transmitted to each RSU. We modeled this mechanism by defining a circular area for the RSUs within which information is considered relevant.

D. TIS data management

The concept of ATB is to maintain local knowledge bases that contain all received traffic information in aggregated form. Most recent approaches select either a probabilistic aggregation scheme for message store maintenance [20], or aggregation based on the distance to the event, e.g. the SOTIS approach relying on annotated maps [11].

In comparison, ATB prioritizes the available information according to the age of an entry and the distance to the event. In principle, only the most recent information is stored for each route segment, i.e. new information elements either update already existing records or are appended to the knowledge base. Furthermore, a garbage collection process continuously expunges entries that are older than a configurable timeout t_{store}. In urban environments, this timeout should be kept small in order to efficiently bypass short-living congestions, whereas, congestions on highways usually last for a longer time and require extended timeouts.

The knowledge base is updated with every received beacon, each of which contains multiple TIS messages. We prioritize entries in the knowledge base to be transmitted in a beacon according to their age, the distance to the event, and the availability of RSUs. In particular, we measure the age of each entry as $\delta t_{entry} = t - t_{entry}$. Both distance measures are estimated in the form of an estimated travel time according to the current speed of the vehicle: $t_c = \frac{\text{distance to event}}{v}$ and $t_r = \frac{\text{distance to RSU}}{v}$ (the distances are estimated using the geographical positions of the vehicle and that of the reported event or the RSU).

Based on these measures, the priority of each entry can be calculated as follows:
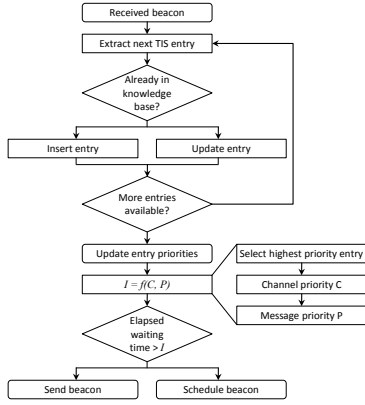
$$p_{entry} = \delta t_{entry} - t_c + t_r \tag{7}$$
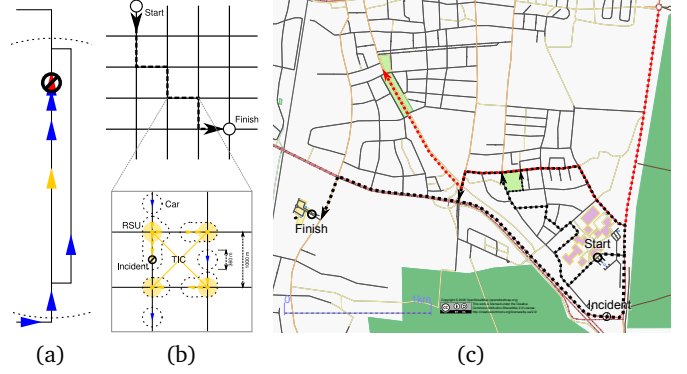
Fig. 4: Handling of received beacon messages

Fig. 5: Scenarios used for the performance evaluation: (a) a simple detour, (b) a grid-shaped road network of 5 km and 16 km width, (c) a small section of the road network in Erlangen

Using the calculated priorities, beacon messages can be generated by selecting as many entries as there is room in a single IEEE 802.11p frame [13] from the top of the list, i.e. those with highest priority. This way, the frame size is optimally used and problems with stateful handling of messages split into multiple frames are inherently avoided. A single entry comprises at least the following elements: Event type, time, position, priority, and RSU identifier.

The handling of received beacons is depicted in Figure 4. Basically, after receiving a beacon each entry is compared with the local knowledge base. If the event is not yet known, the entry is simply appended. Otherwise it is updated appropriately. Each update results in the re-calculation of the priorities of all entries and the calculation of the next beacon interval (see Section III-B). Furthermore, the knowledge base is checked after processing the received beacon to identify events on the current route of the vehicle. If an incident is found, an alternative route is calculated using the Dijkstra shortest path algorithm. Similarly, resolved traffic congestions trigger a re-calculation of the route to check whether there is now a shorter route to the destination.

E. Security and privacy issues

ATB does not include specific security measures. However, as discussed in [30], beaconing can be adequately secured using signatures and certificates added to "selected" messages, e.g. with the help of WAVE security services [31]. In general, the computational and the protocol overhead for this is not negligible. However, certificates or signatures can be omitted, e.g. if transmitting multiple beacons among the same stations [30].

More relevant is the question about privacy issues [32], [33]. Locally, the identity of the vehicle can be derived in form of the used MAC address. The transmitted TIS data, however, does not contain any IDs. Therefore, no privacy concerns apply to forwarded beacons. The only identifier used in the traffic information is that of used RSUs, which we assume does not raise specific privacy concerns.

IV. PERFORMANCE EVALUATION

We evaluated ATB in several simulation experiments to investigate the influence of different ATB protocol parameters, to compare it with traditional approaches, and to show the feasibility in a real-world scenario. In the following section, we outline the simulation environment, describe the experiments, and discuss the results of the evaluations.

A. Simulation environment and parameters

We investigated the performance of ATB with the help of our *Veins*[1] simulation environment, which is based on OMNeT++[2] for event-driven network simulation and SUMO[3] for road traffic microsimulation [28]. OMNeT++ is a simulation environment free for academic use and its *INET Framework* extension offers a set of GPL-licensed simulation modules for simulating computer networks. OMNeT++ runs discrete, event-based simulations of communicating nodes and is becoming increasingly popular in the field of network simulation. SUMO is a GPL-licensed microscopic road

---

[1]http://www7.informatik.uni-erlangen.de/veins/

[2]http://www.omnetpp.org/

[3]http://sumo.sourceforge.net/

traffic simulation environment. It performs simulations both running with and without a GUI and can import city maps from a variety of file formats including freely available *OpenStreetMap* data. SUMO allows high-performance simulations of huge networks with roads consisting of multiple lanes, as well as of intra-junction traffic on these roads, either using simple right-of-way rules or traffic lights. Vehicle types are freely configurable with each vehicle following statically assigned routes, dynamically generated routes, or driving according to a configured timetable.

Both simulators have been extended by modules that allow the road traffic simulation to communicate with its network simulation counterpart via a TCP/IP connection. In particular, this also allows the network simulation to directly control the road traffic simulation and thus to simulate the influence of the TIS on road traffic [28], [34].

With the help of this simulation environment, we configured three classes of settings for the evaluation of ATB, illustrated in Figure 5.

We modeled vehicles capable of exchanging information using ATB by implementing the protocol in Veins according to the principles described in Section III. ATB data was encapsulated in UDP/IP packets and sent via broadcast messages. The radio channel and an IEEE 802.11b NIC transmitting at 11 Mbit/s was modeled by the *INET Framework*. Vehicles moved with a maximum speed of 14 m/s and according to the Krauss mobility model. The full set of simulation parameters common to all evaluated scenarios can be found in Table I. The configured timeout values used for TIS data expiry have been selected according to the accident lengths used in the simulations.

For each scenario, we performed multiple simulation runs for statistical validity and to identify outliers, but no less than 10 runs, and assess the impact of TIS operation using two primary performance metrics: First, we track the effective average speed of vehicles, i.e. the time it takes a vehicle to reach its destination in relation to the traveling time on the shortest route. This metric reflects the benefit of the TIS on traffic as a whole. The impact individual vehicles by smoothing traffic flow is reflected in a second metric, the amount of emitted CO_2. For calculating the CO_2 emissions we employ our implementation of the EMIT emission model presented in [35].

B. ATB parameters

In the first set of simulation experiments, we investigate how ATB adapts to varying conditions occurring in the course of a very simple scenario. We set up a simple single-lane road network, shown in Figure 5a, which consisted essentially of a 300 m main road and a 350 m detour. Traveling along the main road are 101 cars, one entering the simulation every 5 s. The first vehicle is configured to stop from $t = 45\,\text{s} \ldots 105\,\text{s}$, near the end of the main road, creating an artificial traffic incident. All vehicles use ATB to exchange information about this obstruction, so that oncoming cars are able to use the detour and avoid the incident.

During the course of this scenario, we record the internal variables ATB uses to adapt the beacon interval as presented in Section III, i.e. to calculate the message priority P of congestion notifications and to assess the channel quality C. For this evaluation, we present the values of these variables in the form of a scatter plot derived from one exemplary simulation run. In order to represent the large number of observations corresponding to the very same time and value, each individual observation is plotted with a low opacity, so a larger number of similar observations results in a more pronounced dot.

Figure 6 illustrates in this fashion how the two main traffic events, the detection of presence and absence of the obstruction, are reflected in the ATB criterion P, message priority. As no RSUs are present in this scenario, the criteria age A and event distance D_e are

TABLE I: Common simulation parameters

Parameter	Value
minimum beacon interval I_{min}	30 ms
maximum beacon interval I_{max}	60 s
channel quality weighting w_C	2
interval weighting w_I	0.75
number of neighbors for $N = 1$	50
SNR for $S = 1$	50 dB
neighborship data expiry	60 s
TIS data expiry t_{store}	120 s
report traffic incident after queuing	10 s
TIC radius of interest	5 km
processing delay	1 ms ... 10 ms
channel bitrate	11 Mbit/s
approx. transmission radius	180 m
vehicle mobility model	Krauss
max. speed	14 m/s
max. acceleration	2.6 m/s^2
driver imperfection σ	0.5
max. deceleration	4.5 m/s^2
vehicle length	5 m
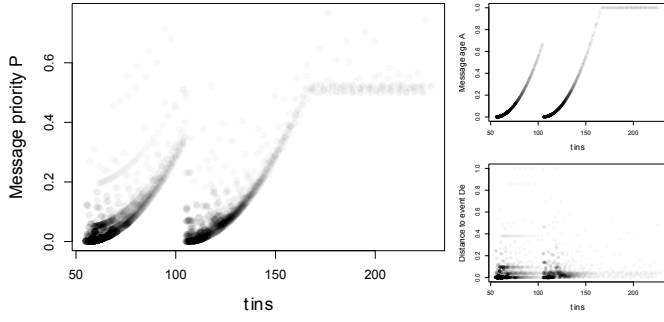
Fig. 6: Message priority criterion P; base criteria age A and distance D_e

Fig. 7: Channel quality criterion C; base criteria neighbors N and SNR S

the only parameters influencing P. At $t = 55\,\text{s}$, the incident is detected by the first vehicle and registered with a value of $A = 0$ and $D_e = 0$. Consistent with the principles of the protocol, A then starts to increase quadratically. At the same time, the spread between observed values of D_e rises quickly as the message is disseminated among the present vehicles. The distinctive bands visible in the plot of D_e correspond to corners in the road network, where vehicles have to decelerate. The number of observed events then starts to decrease until, $60\,\text{s}$ after its creation, the traffic incident is resolved and, hence, A resets to 0. No new events are detected after $t = 105\,\text{s}$ and A again starts to increase quadratically until it reaches is maximum, $A = 1$, where it remains until the end of the simulation.

The second criterion employed by ATB for the calculation of the beacon interval is C, the channel quality, shown in Figure 7. As detailed in Section III, the criterion C is in turn calculated based on three criteria related to neighbors, collisions and SNR. In the scenario, a negligible number of collisions was observed, so the figure only shows the component criteria N and S. As can be seen in the plot N only assumes discrete values, as it is based on the number of known neighbors, a discrete quantity. The value of N increases almost quadratically for the vehicles caught in the jam because, as time advances, a linearly increasing number of vehicles has passed them on the detour. At $t = 105\,\text{s}$ information about the absence of the incident is primarily disseminated by the vehicle just approaching it, which is visible as a narrow band, as well as by the vehicles already in the jam. Throughout the simulation, the value of S varies widely, but the majority of observations indicates a very good SNR. Even as the number of vehicles in the simulation increases, no degradation of the SNR can be seen. Still, throughout the simulation a light band of observations
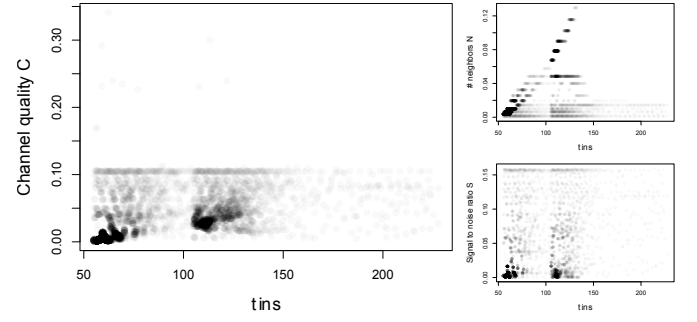
is visible at $S = 0.16$, when high message priority forces the beacon interval to be smaller than the channel quality would advise. Taken together, N and S lead to highly dynamic values of the channel quality criterion C, but only few extreme observations. The outliers of C that are faintly visible in the figure are due to recorded collisions on the channel, each of which instantly increases C to far outside its regular range of values.

C. Comparative evaluation

In order to evaluate the performance of ATB in terms of its impact on road traffic, we use the same scenarios we already employed for the performance study in [28]. As illustrated in Figure 5b, two grid-shaped road networks of $5\,\text{km}$ and $16\,\text{km}$ width were prepared with horizontal and vertical roads spaced $1\,\text{km}$ apart. Starting in one corner, vehicles can then use dynamic routing to avoid obstructions on their way to the opposite corner.

In a first set of simulation runs, we configure 30 vehicles to drive on the $5\,\text{km}$ grid, one starting every $4\,\text{s}$. An artificial traffic incident is created by stopping the lead vehicle for $60\,\text{s}$. We use this scenario to compare the performance of ATB in three network configurations. One offers no infrastructural support, one supports TIS operation by a network of RSUs spread over the intersections, and one contains an additional TIC connected to the network. In order to compare the performance of ATB with that of a protocol using fixed beacon intervals we also simulate two configurations of ATB using $I_{min} = I_{max}$, either set to $1\,\text{s}$ or $10\,\text{s}$. Furthermore, we simulated two baseline scenarios without any radio communications, one with and one without the artificial traffic incident.

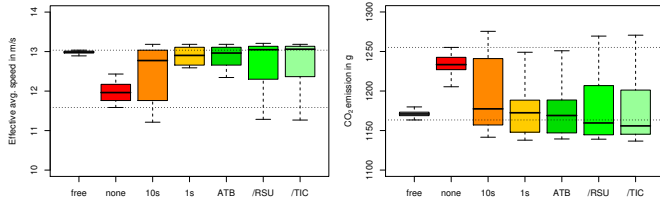Box plots of the results of this set of simulation runs are shown in Figure 8. For each data set, a box is drawn

Fig. 8: Effective average speed of vehicles and their CO_2 emissions in the 5 km grid setup. Plotted are two baseline settings, as well as communication with a fixed beacon interval, and with three ATB configurations: no, partial and full infrastructure support

Fig. 9: TIS performance metrics recorded in the 16 km grid scenario for various fixed beacon intervals as well as for ATB

from the first quartile to the third quartile, and the median is marked with a thick line. Whiskers extend from the edges of the box towards the minimum and maximum of the data set, but no further than 1.5 times the interquartile range. Dotted lines mark the best and worst cases observed in the baseline scenarios.

Aside from the obvious improvements that ATB demonstrates in particular when supported by infrastructure, two interesting effects can be observed from the evaluation. Enabling the car-to-X communication capabilities of simulated vehicles leads to some reaching their destination even faster than the fastest vehicles in an obstruction-free scenario. This is because re-routing around the obstruction leads to traffic flows being more evenly distributed over the network, avoiding micro-jams which commonly occur at corners. At the same time, warning vehicles too late, but also too early, causes them to take unnecessary detours, which results in some vehicles arriving at their destination later than they would have when sticking to their original route. Still, the use of ATB, in particular when supported by infrastructure, typically leads both to lower emissions and to vehicles reaching their destination significantly faster than is possible using fixed beacon intervals, even as short as 1 s.

In a second set of simulation runs, we therefore examine how ATB performs when compared with beacon protocols using even shorter fixed beacon intervals. We also increase the size of the road grid to 16 km and the number of cars to 1 000 to obtain meaningful results for message delays.

Again, we first examine the impact of TIS operation on vehicles' speeds and their CO_2 emission for decreasing values of fixed beacon intervals and for ATB. The results are plotted in Figure 9, which shows the metrics' mean values as well as the 10 % and 90 % quantile. As can be seen it is not until beacon intervals of 1 s
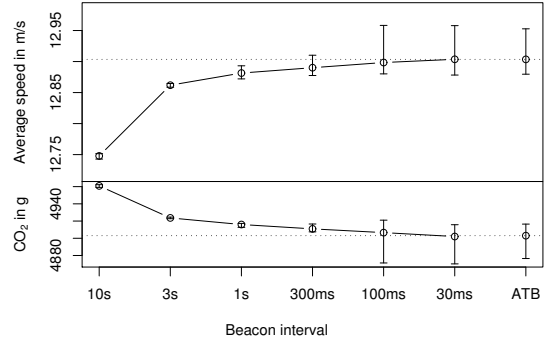
are used that results become comparable with those of ATB.

D. Realistic city scenario

In order to evaluate the performance of ATB in a less synthetic scenario, we chose a road network based on OpenStreetMap data of the city of Erlangen. The modeled section of the city comprises the university campus and a business park about 5 km away. Both were connected by two trunk roads, but were reachable also via several residential roads, as illustrated in Figure 5c. On this network we configured a flow of 200 vehicles, one starting every 6 s, leaving the university campus and heading to the business park. We introduce a traffic obstruction by stopping the lead vehicle for 240 s when it just passed a short one-lane section in the road network. All vehicles following the lead vehicle are therefore either caught in the jam, or, if informed early enough, are able turn back and pick an alternate route.

Shown in Figure 10a are the results from this series of simulations, plotting in the style of a scatter plot for one exemplary run the effective average speed and the CO_2 emission of each vehicle versus the time it entered the simulation. Again we plot results for unobstructed traffic, no car-to-X communication capabilities, for fixed beacon intervals of 10 s and 1 s length, and results for ATB. As can be seen from the plot, the traffic obstruction significantly delays all vehicles caught in the jam.

Only vehicles departing later than just over 400 s enter the simulation late enough to be uninfluenced by the 240 s incident. Both the protocols using a fixed beacon interval and ATB again manage to inform most vehicles of the obstruction in time. Also visible is a group of vehicles that can simply not avoid the incident
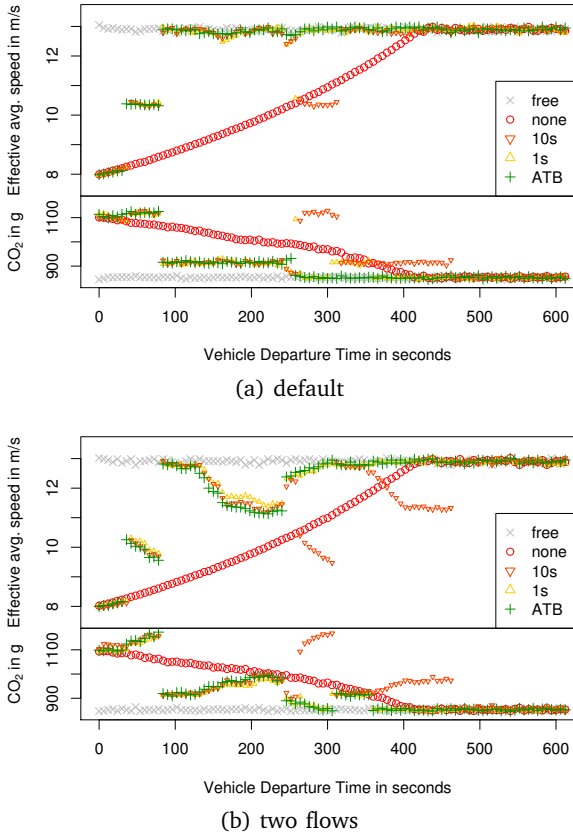
(a) default

(b) two flows

Fig. 10: Traffic performance in the Erlangen scenario

because they are already driving on the single-lane road segment, as well as a group of vehicles that could avoid the incident, but had to turn around to do so. Only at one point in the simulation, when the artificial jam begins to dissolve, a fixed beacon interval of 10 s proves too coarse to keep a number of vehicles from immediately entering the area of the jam.

No secondary jams can be observed in this scenario. The traffic density in this scenario was low enough that all vehicles could be accommodated by the various detours and continue to their destination unobstructed. For a second set of simulation runs we, therefore, set up an additional flow of vehicles which saturated the region crossed by popular detours, indicated by a red arrow in Figure 5c.

Results gathered from this series of simulations, restricted to those gathered from vehicles of the original traffic flow, are shown in Figure 10b. Here, the additional flow of vehicles can be seen to lead to the detours quickly becoming congested. This results in numerous secondary jams and, thus, continuous re-routing of vehicles. As can be seen, a fixed beacon interval of 1 s could improve overall traffic performance

TABLE II: Calculated beacon intervals in the Erlangen scenarios (in s)

	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
default	0.03	0.05	0.06	1.20	0.14	15.55
two flows	0.03	0.06	0.10	3.72	5.39	16.02

beyond what could be managed with ATB. The reason for this behavior is illustrated in Table II: In order to avoid collisions on the radio channel, ATB increased the beacon interval to a mean value of 3.72 s.

V. Conclusion

We presented a new car-to-X communication protocol, Adaptive Traffic Beacon (ATB), which progresses beyond state-of-the-art solutions by providing a self-organizing system architecture that automatically adapts to various settings and conditions. ATB is based on a beaconing approach taking into account the vehicle density, vehicles' speed, radio communication reliability and delay to optimize the beacon period. ATB is adaptive in a second dimension. It automatically makes use of available infrastructure elements such as simple RSUs or even centralized TIC services. We evaluated the protocol performance in extensive simulation experiments covering the analysis of internal protocol parameters, an easy to understand grid setup for comparative evaluation, and a real-world scenario using a map of the city of Erlangen. From the simulation results, we conclude that ATB fulfills its task to support efficient TIS data exchange with support for congestion aware communication. Still an open issue is how to reliably estimate the prospective duration of a detected traffic obstruction, which will influence the knowledge base garbage collection timeout t_{store}. Ongoing work includes a modification of the beacon interval to penalize vehicles close to the sender of a beacon to repeat information less frequently. This will lead to a forwarding scheme similar to greedy forwarding.

References

[1] T. L. Willke, P. Tientrakool, and N. F. Maxemchuk, "A Survey of Inter-Vehicle Communication Protocols and Their Applications," *IEEE Communications Surveys and Tutorials*, vol. 11, no. 2, pp. 3–20, 2009.

[2] M. L. Sichitiu and M. Kihl, "Inter-Vehicle Communication Systems: A Survey," *IEEE Communications Surveys and Tutorials*, vol. 10, no. 2, pp. 88–105, 2008.

[3] L. Iftode, S. Smaldone, M. Gerla, and J. Misener, "Active Highways," in *IEEE PIMRC 2008*, Cannes, France, September 2008, pp. 1–5.

[4] D. Valerio, F. Ricciato, P. Belanovic, and T. Zemen, "UMTS on the Road: Broadcasting Intelligent Road Safety Information via MBMS," in *IEEE VTC 2008-Spring*, May 2008, pp. 3026–3030.

[5] C. Sommer, A. Schmidt, R. German, W. Koch, and F. Dressler, "Simulative Evaluation of a UMTS-based Car-to-Infrastructure Traffic Information System," in *IEEE GLOBECOM 2008, AutoNet Workshop*, New Orleans, LA, December 2008.

[6] E. Schoch, F. Kargl, M. Weber, and T. Leinmüller, "Communication Patterns in VANETs," *IEEE Communications Magazine*, vol. 46, no. 11, pp. 119–125, November 2008.

[7] W. Kiess, J. Rybicki, and M. Mauve, "On the nature of Inter-Vehicle Communication," in *WMAN 2007*, Bern, Switzerland, March 2007, pp. 493–502.

[8] G. Korkmaz, E. Ekici, and F. Özgüner, "An Efficient Fully Ad-Hoc Multi-Hop Broadcast Protocol for Inter-Vehicular Communication Systems," in *IEEE ICC 2006*, vol. 1, Istanbul, Turkey, June 2006, pp. 423–428.

[9] A. Boukerche, H. Oliveira, E. Nakamura, and A. Loureiro, "Vehicular Ad Hoc Networks: A New Challenge for Localization-Based Systems," *Elsevier Computer Communications*, vol. 31, no. 12, pp. 2838–2849, July 2008.

[10] C. Lochert, B. Scheuermann, C. Wewetzer, A. Luebke, and M. Mauve, "Data Aggregation and Roadside Unit Placement for a VANET Traffic Information System," in *ACM VANET 2008*, San Francisco, CA, 2008, pp. 58–65.

[11] L. Wischhof, A. Ebner, and H. Rohling, "Information dissemination in self-organizing intervehicle networks," *IEEE Transactions on Intelligent Transportation Systems*, vol. 6, no. 1, pp. 90–101, March 2005.

[12] K. C. Lee, U. Lee, and M. Gerla, "TO-GO: TOpology-assist Geo-Opportunistic Routing in Urban Vehicular Grids," in *IEEE/IFIP WONS 2009*. Snowbird, UT: IEEE, February 2009, pp. 11–18.

[13] "Wireless Access for Vehicular Environments," IEEE, Draft Standard 802.11p-D4.0, 2008.

[14] C. Sommer and F. Dressler, "The DYMO Routing Protocol in VANET Scenarios," in *IEEE VTC 2007-Fall*, Baltimore, MD, September/October 2007, pp. 16–20.

[15] J. Ott and D. Kutscher, "A Disconnection-Tolerant Transport for Drive-thru Internet Environments," in *IEEE INFOCOM 2005*, Miami, FL, March 2005.

[16] A. Skordylis and N. Trigoni, "Delay-bounded Routing in Vehicular Ad-hoc Networks," in *ACM Mobihoc 2008*, Hong Kong, China, May 2008, pp. 341–350.

[17] C. Lochert, B. Scheuermann, M. Caliskan, and M. Mauve, "The Feasibility of Information Dissemination in Vehicular Ad-Hoc Networks," in *IEEE/IFIP WONS 2007*, Obergurgl, Austria, January 2007, pp. 92–99.

[18] N. Wisitpongphan, F. Bai, P. Mudalige, V. Sadekar, and O. Tonguz, "Routing in Sparse Vehicular Ad Hoc Wireless Networks," *IEEE Journal on Selected Areas in Communications (JSAC)*, vol. 25, no. 8, pp. 1538–1556, October 2007.

[19] P. Costa, D. Frey, M. Migliavacca, and L. Mottola, "Towards lightweight information dissemination in inter-vehicular networks," in *ACM VANET 2006*, Los Angeles, CA, 2006, pp. 20–29.

[20] C. Lochert, B. Scheuermann, and M. Mauve, "Probabilistic Aggregation for Data Dissemination in VANETs," in *ACM VANET 2007*, Montreal, Quebec, Canada, September 2007, pp. 1–8.

[21] J. Rybicki, B. Scheuermann, W. Kiess, C. Lochert, P. Fallahi, and M. Mauve, "Challenge: Peers on Wheels - A Road to New Traffic Information Systems," in *ACM MobiCom 2007*, Montreal, Canada, September 2007, pp. 215–221.

[22] B. B. Chen and M. C. Chan, "MobTorrent: A Framework for Mobile Internet Access from Vehicles," in *IEEE INFOCOM 2009*, Rio de Janeiro, Brazil, April 2009.

[23] K. C. Lee, M. Le, J. Härri, and M. Gerla, "LOUVRE: Landmark Overlays for Urban Vehicular Routing Environments," in *IEEE VTC 2008-Fall*, September 2008, pp. 1–5.

[24] G. Resta, P. Santi, and J. Simon, "Analysis of MultiHop Emergency Message Propagation in Vehicular Ad Hoc Networks," in *ACM Mobihoc 2007*, Montreal, Quebec, Canada, September 2007, pp. 140–149.

[25] F. Schmidt-Eisenlohr, M. Torrent-Moreno, J. Mittag, and H. Hartenstein, "Simulation Platform for Inter-Vehicle Communications and Analysis of Periodic Information Exchange," in *IEEE/IFIP WONS 2007*, Obergurgl, Austria, January 2007, pp. 50–58.

[26] M. Torrent-Moreno, "Inter-Vehicle Communications: Assessing Information Dissemination under Safety Constraints," in *IEEE/IFIP WONS 2007*, Obergurgl, Austria, January 2007, pp. 59–64.

[27] E. Van de Velde and C. Blondia, "Adaptive REACT protocol for Emergency Applications in Vehicular Networks," in *IEEE LCN 2007*, Dublin, Ireland, October 2007, pp. 613–619.

[28] C. Sommer, Z. Yao, R. German, and F. Dressler, "Simulating the Influence of IVC on Road Traffic using Bidirectionally Coupled Simulators," in *IEEE INFOCOM 2008, MOVE Workshop*, Phoenix, AZ, April 2008.

[29] E. J. Rivera-Lara, R. Herrerias-Hernandez, J. A. Perez-Diaz, and C. F. Garcia-Hernandez, "Analysis of the Relationship between QoS and SNR for an 802.11g WLAN," in *International Conference on Communication Theory, Reliability, and Quality of Service (CTRQ 2008)*. Bucharest, Romania: IEEE, June/July 2008, pp. 103–107.

[30] F. Kargl, E. Schoch, B. Wiedersheim, and T. Leinmüller, "Secure and Efficient Beaconing for Vehicular Networks," in *ACM VANET 2008, Poster Session*, San Francisco, CA, 2008, pp. 82–83.

[31] R. A. Uzcátegui and G. Acosta-Marum, "WAVE: A Tutorial," *IEEE Communications Magazine*, vol. 47, no. 5, pp. 126–133, May 2009.

[32] P. Papadimitratos, L. Buttyan, T. Holczer, E. Schoch, J. Freudiger, M. Raya, Z. Ma, F. Kargl, A. Kung, and J.-P. Hubaux, "Secure Vehicular Communication Systems: Design and Architecture," *IEEE Communications Magazine*, vol. 46, no. 11, pp. 100–109, November 2008.

[33] F. Dressler, T. Gansen, C. Sommer, and L. Wischhof, "Requirements and Objectives for Secure Traffic Information Systems," in *IEEE MASS 2008, WSNS Workshop*, Atlanta, GA, September 2008, pp. 808–814.

[34] A. Wegener, M. Piorkowski, M. Raya, H. Hellbrück, S. Fischer, and J.-P. Hubaux, "TraCI: An Interface for Coupling Road Traffic and Network Simulators," in *11th Communications and Networking Simulation Symposium (CNS'08)*, Ottawa, Canada, April 2008.

[35] A. Cappiello, I. Chabini, E. Nam, A. Lue, and M. Abou Zeid, "A statistical model of vehicle emissions and fuel consumption," in *IEEE ITSC 2002*, September 2002, pp. 801–809.