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

Christoph Sommer*, Ozan K. Tonguz† and Falko Dressler*
*Computer Networks and Communication Systems, Dept. of Computer Science, University of Erlangen, Germany
†Electrical and Computer Engineering Dept., Carnegie Mellon University, Pittsburgh, PA
tonguz@ece.cmu.edu, {csommer,dressler}@informatik.uni-erlangen.de

***Abstract*—We present Adaptive Traffic Beacon (ATB), a fully-decentralized car-to-X protocol built around the central ideas of delay sensitivity and congestion awareness. From previous research findings, we see that centralized solutions, VANETs, and broadcasting based approaches each show benefits and drawbacks depending on traffic density, penetration, network utilization, and other parameters; intelligent transportation systems are therefore tuned for specific settings. In order to overcome this limitation, we developed a broadcast-based solution that has been designed to carefully use only the remaining capacity of the wireless channel. Thus, it will not influence other applications using the same radio network. ATB is adaptive in two dimensions: First, the beacon interval is adapted dynamically according to the channel quality and the importance of messages in the local knowledge base, and, secondly, the protocol can dynamically make use of available infrastructure elements. Simulation experiments demonstrate that ATB performs well in a broad range of settings. It maintains a non-congested wireless channel to prevent collisions during the data exchange.**

I. Introduction

Intelligent transportation systems rely on accurate and timely information about road traffic, in particular about possible congestions and accidents. Traffic Information Systems (TIS) are currently one of the most interesting application domains – both from a scientific point of view and from a business perspective. In this paper, we focus on both the collection and the distribution of traffic information in the context of fully decentralized 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 [1].

In the last couple of years, many efforts have been reported that study quite diverse strategies for Inter-Vehicle Communication (IVC), resulting in a variety of specialized IVC protocols [2]. However, some of the most challenging problems are still not fully solved: First, protocols have to cope with rapid changes in network topology and utilization. Secondly, available resources have to be coordinated in a self-organizing, distributed way, dynamically incorporating infrastructure elements and even centralized information repositories. Thirdly, delay-sensitive transmission of emergency messages has to be balanced against channel load, so as not to overload the channel [3]. Basically, two main research lines are being investigated for efficient data dissemination in the scope of TIS: dynamic broadcast-based solutions and infrastructure-based peer-to-peer approaches [4].

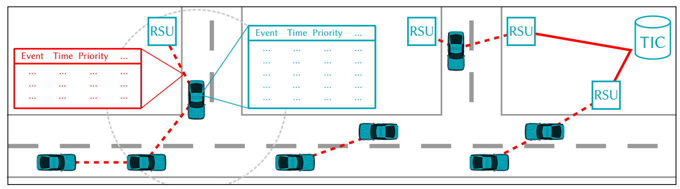

Fig. 1: ATB system architecture

We present and evaluate the Adaptive Traffic Beacon (ATB) protocol, which is designed to ensure an uncongested channel, i.e. prevent packet loss due to collisions, and reduce the end-to-end delay of the information transfer. ATB uses periodic beacons, i.e. single broadcast packets, to exchange information among neighboring cars. The key aspect of ATB is to continuously adapt the interval between two beacons to carefully use only the remaining capacity of the wireless channel, without influencing other protocols. The quality of the wireless channel is estimated using three aspects: Observing collisions, recent overload situations can be detected. The signal-to-noise ratio provides a rough estimate of the current channel conditions. Finally, the vehicle density is an indicator of transmissions to be expected in the next time interval. We further take a message priority into account, i.e. a value depending on the distance to an event or the availability of infrastructure. In addition to supporting fully decentralized information exchange among participating vehicles, ATB can also make use of available infrastructure. As illustrated in Figure 1, support may be completely absent or range from disconnected units participating only in the wireless network, up to a network of Roadside Units (RSUs) and servers.

The main contributions of this paper can be summarized as follows: ATB uses a variable interval for the dissemination of elements in a local knowledge base, dynamically adapting the rate to a wide range of parameters such as wireless channel conditions, vehicle density, communication reliability, and delay. It does not assume, nor create, network topology or roadmap information. Using this adaptive beaconing concept, collisions on the channel become negligible. Thus, ATB is able to efficiently operate the wireless channel even at high vehicle densities. To the best of our knowledge, ATB is the first solution that adapts the beacon interval in a fully self-organizing manner and that is tolerant to other IVC protocols using the same channel.

978-1-4244-9524-5/10/$26.00 ©2010 IEEE

II. Related Work

Inter-Vehicle Communication (IVC) can be categorized into several classes according to application scenarios, each leading to unique traffic patterns and protocol requirements [5]. In this paper, we concentrate on TIS applications in general and cooperative driving and safety applications in particular.

Early solutions relied on a Traffic Information Center (TIC) and pre-deployed RSUs, using Mobile Ad Hoc Network (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 DSR or DYMO [6]. 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, yet the 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 avoid extreme network congestion, albeit at high operational costs. Moreover, it has been discovered that Vehicular Ad Hoc Networks (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 [7].

In the following, we briefly point out IVC approaches described in the literature that take different approaches to dealing with these problems.

Delay Tolerant Network (DTN) related approaches [8] support connectivity of nodes to a TIC by 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 [3]. Another approach, and currently the only commercially successful alternative, is the use of cellular networks to connect vehicles, as implemented in the TomTom HD traffic service. Recent studies show that capabilities of 3G/4G networks, especially the availability of multicast communication [9], [10], can be beneficial for the development of TIS applications. However, all these approaches are heavily dependent on available network infrastructure elements and support only an efficient downlink to the vehicles. Moreover, they rely on a central 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 [5], [6].

The use of cellular networks without the need for a TIC was investigated by incorporating ideas from the peer-to-peer domain in the *peers on wheels* vision and further refined in PeerTIS [4]. Conceptually, it is possible to build extremely robust traffic information systems supporting publish/subscribe interfaces managed by a Distributed Hash Table (DHT) maintained by the vehicles. Further, the MobTorrent approach has been published [11], which also provides mobile (BitTorrent-like) Internet access from vehicles using RSUs (building on the ideas of drive-thru Internet [8], but exploiting state-of-the-art data management functions). LOUVRE [12], on the other hand, provides overlay routing in vehicular environments. It has also been shown that lightweight RSUs, called stationary support units or repeaters, may be used to replace expensive RSUs having a permanent backbone connection [13].

As a further alternative, decentralized infrastructure-less solutions have been investigated. One of the most sophisticated solutions in this class is the Self-Organizing Traffic Information System (SOTIS) [14]. 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 information.

Common to these solutions is the broadcast of traffic information to neighboring vehicles, either periodically or triggered by new events [13]. Such traffic information can be surveyed in a decentralized manner, e.g., based on spatio-temporal data obtained from vehicle position traces [15]. The dissemination process can also be supported using directed (i.e. geographic) flooding, which makes lightweight information encoding about both target areas and preferred routes a necessity [16]. 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 [14], [17], [18]. Multi-hop broadcast is thus a promising technique, especially for emergency message propagation with delay bounds [19].

Beaconing, or 1-hop broadcast, is an inherent feature of most of the discussed systems. For example, neighborhood information is collected using beaconing. The exploitation of periodic information exchange using such beacons, with special focus on safety applications, has been analyzed in extensive simulations in [20], showing that with increasing distance, the success ratio decreased quickly. Combined with a position based forwarding strategy, however, the approach could be improved. It was also shown that network load can be significantly reduced by selectively suppressing broadcasts based on 1-hop neighbor information and that reliability can be increased via the use of explicit acknowledgments [21]. Most recently, 2-hop beaconing has been described to acquire topology knowledge for opportunistic forwarding using the selected best target forwarder [18].

The main challenge for all such 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 [22]. Based only on neighbor detection, it can skip intervals for beacon transmission to support emergency applications. Furthermore, fundamental scalability criteria need to be considered in order to make the protocol applicable in the target scenario [23].

We believe that beaconing systems are in general well suited for TIS data exchange. The key issue is that their optimal configuration is highly dependent on environmental conditions, e.g., number of vehicles or channel load. A fully adaptive protocol was thus identified as the next logical step [24].

III. Adaptive Traffic Beacon

In the following, we describe the key concepts of ATB and motivate the chosen estimation criteria for the beacon interval. We also outline the capabilities to use available infrastructure elements to further improve TIS data dissemination.

A. System architecture

From previous work, we learned that centralized solutions and broadcast based approaches each show benefits and drawbacks depending on a wide range of system parameters. The feasibility, but also the quality, of transmissions depends mainly on the vehicle density and penetration rate. This has also been confirmed in [7]. For our new 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 only those information elements most relevant to the node. 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 quality and the importance of the message to send. As a simple rule, ATB tries to send beacons, i.e. TIS data fragments, as frequently as possible to ensure fast and reliably delivery, but always checks the channel quality to prevent collisions and interference with other protocols using the same wireless channel.

Figure 1 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 utility, which is based on the importance of the message and the estimated benefit to other vehicles. Each beacon contains a subset of these entries. Furthermore, infrastructure support can be exploited for improved information exchange.

B. Adaptive beacon intervals

ATB uses two different metrics to calculate the interval parameter I: the *channel quality* C and the *message priority* P. Like all metrics of ATB, smaller values of C and P represent a better channel and a higher priority, respectively. 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 $[0, 1]$) is calculated according as:

$$I = (1 - w_I) \times P^2 + (w_I \times C^2) \tag{1}$$

We experimented using linear combinations of the parameters and finally deduced that the interval parameter I matches the environmental conditions best if C and P are included in squared form. Similarly, the weighting factor w_I needs to either emphasize on C or P: In our experiments, we always used $w_I > 0.5$ to make ATB more sensitive to the channel quality. Figure 2 shows the behavior of I for $w_I = 0.75$. As can be seen, the interval parameter becomes 1 only for the lowest message priority and the worst channel quality. In all other cases, I quickly falls to values below 0.5.

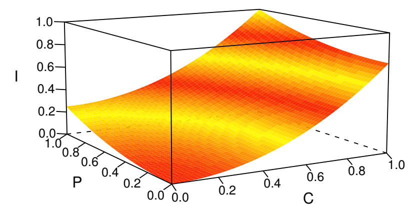

Fig. 2: Interval parameter I for interval weighting $w_I = 0.75$

From the interval parameter, the beacon interval ΔI is then derived (I_{min} and I_{max} represent the minimum and the maximum beacon interval, respectively):

$$\Delta I = I_{min} + (I_{max} - I_{min}) \times I \tag{2}$$

a) Channel quality C: The channel quality is a metric indicating the availability of channel resources for ATB transmissions. We observe three almost independent parameters to estimate the congestion probability in three time scales.

Based on the number of collisions or bit errors observed in the last time interval, the recent load on the channel can be estimated. The key objective of ATB is to ensure congestion-aware communication, i.e. not to interfere with other protocols using the same wireless channel or with other cars using ATB. We model $K = 1 - \frac{1}{1+\#\text{ collisions}}$.

An estimate for the current transmission quality is the Signal to Noise Ratio (SNR) as perceived for the last transmission. We model $S = \max\left\{0; (\frac{\text{SNR}}{\text{max. SNR}})^2\right\}$. In measurements, it has been shown that the error rate of WiFi communication quickly increases if the SNR drops below 25 dB [25]. Therefore, we configured a maximum SNR (50 dB) so that S already decreases to 0.25 for a SNR equal to 25 dB.

Finally, we need to predict the probability of other transmissions in the next time interval. Here, we use the density of vehicles, i.e. the number of neighbors, to estimate the congestion probability (the more neighbors, the higher the probability for simultaneous transmissions). We model $N = \min\left\{(\frac{\#\text{ neighbors}}{\text{max. }\#\text{ neighbors}})^2; 1\right\}$. The parameter quadratically approaches 1, scaled by a pre-defined maximum.

Finally, the channel quality C can be calculated as follows (the factor $w_C \geq 1$ is used to weight the measured parameters K and S higher than the estimated congestion probability N):

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

b) Message priority P: The message priority is an indicator for the demand to broadcast messages early and frequently. Basically, the message priority allows nodes to schedule the next transmission in a way that nodes having high priority messages will be able to transmit first. In our current model, 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 utility in the local knowledge base (see Section III-D).

First, information age is accounted for 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 represent the node's proximity to the event $D_e = \min\left\{(\frac{\text{distance to event}/v}{I_{max}})^2; 1\right\}$ as well as 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 proximity in the form of an estimated travel time. This distance estimation can be further enhanced using map and location information as described in the TO-GO approach [18].

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

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 Traffic Information Centers (TICs) are inherently supported by ATB. In principle, ATB-enabled vehicles and ATB-enabled RSUs operate in a similar fashion. RSUs participate in the beaconing process and adapt the beacon interval according to the same rules 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 [17]. The RSUs can further be connected to a backbone network. We assume that these RSUs know their own geographic position and those of their neighboring RSUs. Therefore, data muling concepts can be realized [3].

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 [26], or aggregation based on the distance to the event, e.g., the SOTIS approach relying on annotated maps [14].

Our current implementation of ATB simply stores only the most recent information for each route segment, i.e. new information elements either update records for an existing route segment (either in part or as a whole) or they are appended to the knowledge base. In order to deliver better performance and scalability it can, however, be readily extended to employ advanced data management and aggregation techniques found in the literature [14], [17], [18].

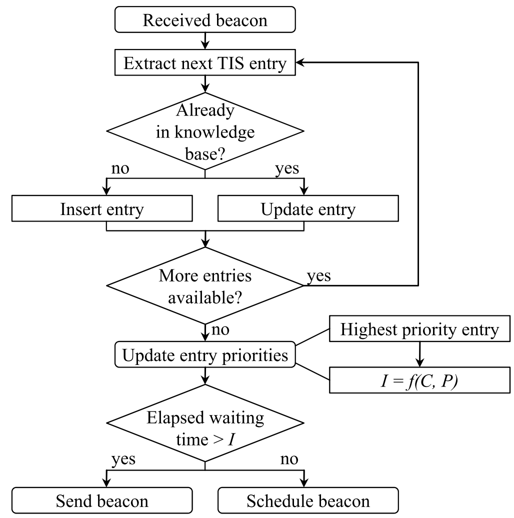

Fig. 3: Handling of received beacon messages

In our implementation, the knowledge base is updated with every received beacon, each of which may contain multiple information elements. We prioritize entries to be transmitted in a beacon according to their age $\delta t_{entry} = t - t_{entry}$, the proximity to the event $t_c = \frac{\text{distance to event}}{v}$ and the proximity to the next RSU $t_r = \frac{\text{distance to RSU}}{v}$. Based on these measures, the priority of each entry can be calculated as follows:

$$p_{entry} = \delta t_{entry} - t_c + t_r \tag{5}$$

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 [27] from the top of the list, i.e. those with highest priority. 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 3. 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.

E. Security and privacy issues

ATB does not include specific security measures. However, as discussed in [28], beaconing can be adequately secured using signatures and certificates added to "selected" messages, e.g., with the help of WAVE security services [29]. In general, the computational and the protocol overhead for this is not negligible. However, this data can be omitted, e.g., if transmitting multiple beacons among the same stations [28].

Also highly relevant are questions surrounding privacy issues [30]. The transmitted TIS data, however, does not contain the ID of any vehicle. Therefore, the operation of ATB does not further interfere with privacy enhancing schemes implemented on lower protocol layers. The only identifier used in the traffic information is that of used RSUs, which we assume does not raise specific privacy concerns.

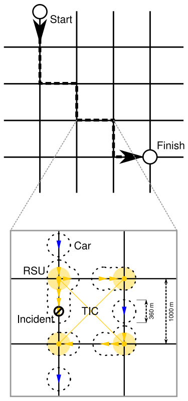

(a)

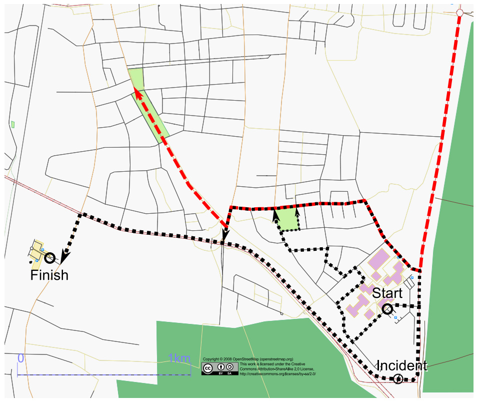

(b)

Fig. 4: Scenarios used for the performance evaluation: (a) a grid-shaped road network of $1\,\mathrm{km}$ grid width, (b) a small section of the road network in Erlangen

TABLE I: Common simulation parameters

Parameter	Value
minimum beacon interval I_{min}	$30\,\mathrm{ms}$
maximum beacon interval I_{max}	$60\,\mathrm{s}$
channel quality weighting w_C	2
interval weighting w_I	0.75
number of neighbors for $N = 1$	50
SNR for $S = 1$	$50\,\mathrm{dB}$
neighborship data expiry	$60\,\mathrm{s}$
TIS data expiry t_{store}	$120\,\mathrm{s}$
report traffic incident after queuing	$10\,\mathrm{s}$
TIC radius of interest	$5\,\mathrm{km}$
processing delay	$1\,\mathrm{ms} \ldots 10\,\mathrm{ms}$
channel bitrate	$11\,\mathrm{Mbit/s}$
approx. transmission radius (Friis)	$180\,\mathrm{m}$
vehicle mobility model	Krauss
max. speed	$14\,\mathrm{m/s}$
max. acceleration	$2.6\,\mathrm{m/s^2}$
driver imperfection σ	0.5
max. deceleration	$4.5\,\mathrm{m/s^2}$
vehicle length	$5\,\mathrm{m}$

IV. Performance Evaluation

We evaluated ATB in several simulation experiments to investigate the influence of different protocol parameters, to compare it with traditional approaches, and to show the feasibility in a real-world scenario. In the following, 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 [31], [32]. OMNeT++ is a simulation framework free for academic use, runs discrete, event-based simulations of communicating nodes and is becoming increasingly popular in the field of network simulation. SUMO is a road traffic simulation environment that can import city maps from a variety of file formats including freely available *OpenStreetMap* data.

With the help of this simulation environment, we configured two classes of settings for the evaluation of ATB, illustrated in Figure 4. For each part of the analysis, we used an appropriate simulation scenario in order to clearly show the specific protocol characteristics. The knowledge base of ATB-equipped vehicles is checked after processing each 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. The full set of simulation parameters common to all evaluated scenarios can be found in Table I. The configured timeout values for TIS data expiry have been selected according to the accident lengths used in the simulations.

[1] http://www7.informatik.uni-erlangen.de/veins/
[2] http://www.omnetpp.org/
[3] http://sumo.sourceforge.net/

For each scenario, we performed multiple simulation runs for statistical validity and to identify outliers, but no less than 10 runs, and assessed the impact of TIS operation using two primary performance metrics: First, we tracked 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. Its impact on individual vehicles, smoothing the 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 [33], [34].

B. Comparative evaluation

In order to evaluate the performance of ATB in terms of its impact on road traffic, we rely on similar scenarios as used in the literature, e.g., in [31]. As illustrated in Figure 4a, two grid-shaped road networks of $5\,\mathrm{km}$ and $16\,\mathrm{km}$ width were prepared with horizontal and vertical roads spaced $1\,\mathrm{km}$ 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 configured 30 vehicles to drive on the $5\,\mathrm{km}$ grid, one departing every $4\,\mathrm{s}$. An artificial traffic incident is created by stopping the lead vehicle for $60\,\mathrm{s}$. We used 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 based on fixed beacon intervals, we also simulated configurations of ATB using $I_{min} = I_{max}$. The resulting fixed beaconing scheme, however, still includes our optimizations to start beaconing only if data is available. Furthermore, we simulated two baseline scenarios without any radio communications, one with and one without the artificial traffic incident.

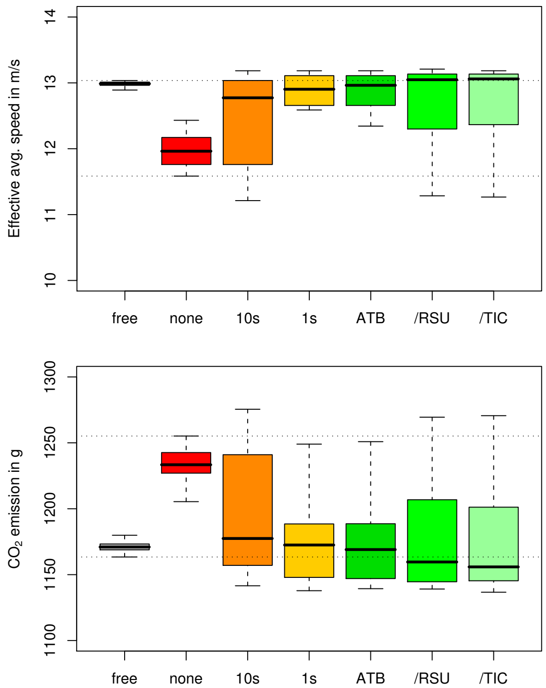

Fig. 5: Effective average speed of vehicles and their CO_2 emissions in the 5 km grid setup. Plotted are two baseline settings, as well as communication with a fixed beacon interval, and with three ATB configurations: no, partial and full infrastructure support

Box plots from this set of simulation runs are shown in Figure 5. For each data set, a box is drawn 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, two interesting effects can be observed from the evaluation: enabling car-to-X capabilities of simulated vehicles leads to some of them 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 at corners or intersections [32]. Yet, warning some 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 faster than is possible using fixed beacon intervals, even as short as 1 s.

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

We first examined the impact of TIS operation on vehicles' speeds and CO_2 emissions for fixed beacon intervals and for ATB. The results are plotted in Figure 6a, which shows the metrics' mean values as well as the 10 % and 90 % quantiles. As can be seen it is not until beacon intervals of 1 s are used that results become comparable with those of ATB.

We further examined this scenario by comparing for the same setups the end to end delays of generated traffic information. Figure 6b shows this metric in the form of CDF plots. As can be seen, ATB performance can just match that of fixed beaconing with intervals. The depicted delays represent the typical store-carry-forward behavior in VANETs, which is greatly influenced by the mobility of vehicles. It also needs to be mentioned that the absolute measures need to be carefully evaluated because lost messages do not contribute to the shown CDF.

We therefore also examined the load and congestion of the wireless channel by measuring the number of collisions per packet received. Figure 6c depicts the results in a log-scale graph: periodic beaconing always leads to a significant number of collisions, especially for periods smaller than 1 s. In contrast, ATB carefully manages the channel to operate below congestion threshold.

From these results, we see that static beaconing with a period clearly smaller than 1 s allows a similar range and quality of the TIS information exchange. However, as can be seen from Figure 6c, the number of collisions caused by the static beaconing exponentially increases for smaller periods. ATB is clearly able to perform well in all the investigated scenarios. Thus, we can conclude that ATB succeeds at managing access to the radio channel – which, according to the used quality metrics, also holds if other devices or applications start sharing the same wireless channel.

C. 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 are connected by two trunk roads, but are reachable also via several residential roads, as illustrated in Figure 4b. On this network, we configure 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 as it passes 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 the first plot of Figure 7 (default scenario) 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 at. 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 for ATB.

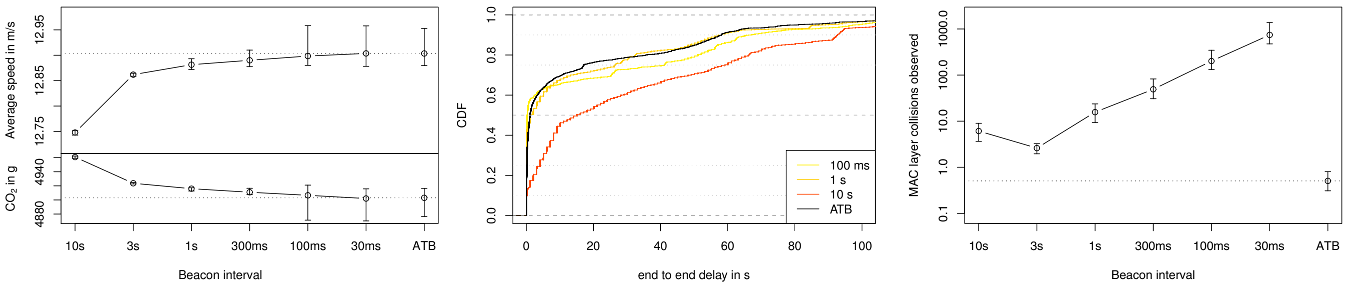

(a) traffic performance (b) end to end delay (c) MAC layer collisions

Fig. 6: TIS performance metrics recorded in the 16 km grid scenario for various fixed beacon intervals as well as for ATB

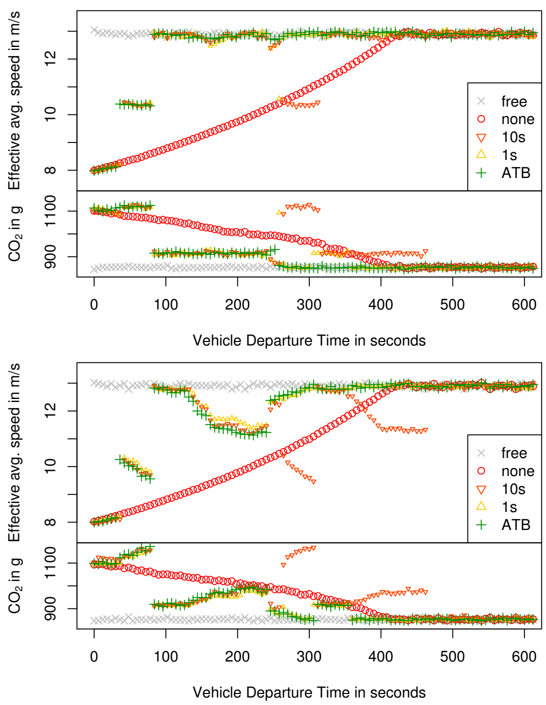

Fig. 7: Traffic performance in Erlangen scenarios; top: default scenario; bottom: additional vehicle flow saturating detours

TABLE II: Beacon intervals in 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

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 because they are already driving on the single-lane road segment, as well as a group of vehicles that can avoid the incident, but have 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 some vehicles from immediately entering the area of the jam.

No secondary jams could be observed in this scenario. The traffic density in this scenario was low enough that all vehicles can be accommodated by the various detours and continue to their destination unobstructed. In a second set of simulation runs, we therefore set up an additional flow of vehicles which saturates the region crossed by popular detours, indicated by a red, dashed arrow in Figure 4b.

Results gathered from this series of simulations, restricted to those gathered from vehicles of the original traffic flow, are shown in the second plot of Figure 7. 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. A fixed beacon interval of 1 s improves overall traffic performance beyond the capabilities of ATB. The reason for this behavior, however, is illustrated in Table II: In order to avoid collisions on the radio channel, ATB has to operate with a much higher mean beacon interval 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 vehicle density, vehicles' speed, radio communication reliability and delay to optimize the beacon period. ATB is also adaptive in a second dimension. It can automatically use optionally available infrastructure elements such as RSUs or TIC servers. We evaluated the protocol performance in extensive simulation experiments.

From the results, we conclude that ATB fulfills its task to support efficient TIS data exchange with support for delay-sensitive and congestion-aware communication. ATB performs well in both synthetic scenarios and realistic urban networks, without assuming the presence of (or creating and maintaining) network topology or roadmap information.

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. Secondly, the implementation will need to be extended to include an established aggregation and data management scheme to meet scalability demands. Furthermore, we are working on an extended simulation framework that will be able to more accurately model physical layer effects by taking obstacles to radio communication, such as buildings and other cars, into account.

References

[1] L. Iftode, S. Smaldone, M. Gerla, and J. Misener, "Active Highways," in *19th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC 2008)*, Cannes, France, September 2008, pp. 1–5.

[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] A. Skordylis and N. Trigoni, "Delay-bounded Routing in Vehicular Ad-hoc Networks," in *9th ACM International Symposium on Mobile Ad Hoc Networking and Computing (ACM Mobihoc 2008)*. Hong Kong, China: ACM, May 2008, pp. 341–350.

[4] J. Rybicki, B. Scheuermann, M. Koegel, and M. Mauve, "PeerTIS - A Peer-to-Peer Traffic Information System," in *6th ACM International Workshop on Vehicular Inter-Networking (VANET 2009)*. Beijing, China: ACM, September 2009, pp. 23–32.

[5] 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.

[6] C. Sommer and F. Dressler, "The DYMO Routing Protocol in VANET Scenarios," in *66th IEEE Vehicular Technology Conference (VTC2007-Fall)*. Baltimore, MD: IEEE, September/October 2007, pp. 16–20.

[7] 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.

[8] J. Ott and D. Kutscher, "A Disconnection-Tolerant Transport for Drive-thru Internet Environments," in *24th IEEE Conference on Computer Communications (IEEE INFOCOM 2005)*, Miami, FL, March 2005.

[9] D. Valerio, F. Ricciato, P. Belanovic, and T. Zemen, "UMTS on the Road: Broadcasting Intelligent Road Safety Information via MBMS," in *67th IEEE Vehicular Technology Conference (VTC2008-Spring)*, May 2008, pp. 3026–3030.

[10] C. Sommer, A. Schmidt, Y. Chen, R. German, W. Koch, and F. Dressler, "On the Feasibility of UMTS-based Traffic Information Systems," *Elsevier Ad Hoc Networks, Special Issue on Vehicular Networks*, vol. 8, no. 5, pp. 506–517, July 2010.

[11] B. B. Chen and M. C. Chan, "MobTorrent: A Framework for Mobile Internet Access from Vehicles," in *28th IEEE Conference on Computer Communications (IEEE INFOCOM 2009)*. Rio de Janeiro, Brazil: IEEE, April 2009.

[12] K. C. Lee, M. Le, J. Härri, and M. Gerla, "LOUVRE: Landmark Overlays for Urban Vehicular Routing Environments," in *68th IEEE Vehicular Technology Conference (VTC2008-Fall)*, September 2008, pp. 1–5.

[13] G. Korkmaz, E. Ekici, and F. Özgüner, "An Efficient Fully Ad-Hoc Multi-Hop Broadcast Protocol for Inter-Vehicular Communication Systems," in *IEEE International Conference on Communications (IEEE ICC 2006)*, Istanbul, Turkey, June 2006, pp. 423–428.

[14] 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.

[15] J. Yoon, B. Noble, and M. Liu, "Surface Street Traffic Estimation," in *5th ACM International Conference on Mobile Systems, Applications, and Services (ACM MobiSys 2007)*, San Juan, Puerto Rico, June 13 2007.

[16] P. Costa, D. Frey, M. Migliavacca, and L. Mottola, "Towards lightweight information dissemination in inter-vehicular networks," in *3rd ACM International Workshop on Vehicular Ad Hoc Networks (VANET 2006)*. Los Angeles, CA: ACM, 2006, pp. 20–29.

[17] C. Lochert, B. Scheuermann, C. Wewetzer, A. Luebke, and M. Mauve, "Data Aggregation and Roadside Unit Placement for a VANET Traffic Information System," in *5th ACM International Workshop on Vehicular Inter-Networking (VANET 2008)*. San Francisco, CA: ACM, 2008, pp. 58–65.

[18] K. Lee, U. Lee, and M. Gerla, "Geo-Opportunistic Routing for Vehicular Networks," *IEEE Communications Magazine*, vol. 48, no. 5, pp. 164–170, May 2010.

[19] G. Resta, P. Santi, and J. Simon, "Analysis of MultiHop Emergency Message Propagation in Vehicular Ad Hoc Networks," in *8th ACM International Symposium on Mobile Ad Hoc Networking and Computing (ACM Mobihoc 2007)*, Montreal, Quebec, Canada, September 2007, pp. 140–149.

[20] F. Schmidt-Eisenlohr, M. Torrent-Moreno, J. Mittag, and H. Hartenstein, "Simulation Platform for Inter-Vehicle Communications and Analysis of Periodic Information Exchange," in *4th IEEE/IFIP Conference on Wireless On demand Network Systems and Services (WONS 2007)*. Obergurgl, Austria: IEEE, January 2007, pp. 50–58.

[21] F. J. Ros, P. M. Ruiz, and I. Stojmenovic, "Reliable and Efficient Broadcasting in Vehicular Ad Hoc Networks," in *69th IEEE Vehicular Technology Conference (VTC2009-Spring)*. Barcelona, Spain: IEEE, April 2009, pp. 1–5.

[22] 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.

[23] B. Scheuermann, C. Lochert, J. Rybicki, and M. Mauve, "A Fundamental Scalability Criterion for Data Aggregation in VANETs," in *15th ACM International Conference on Mobile Computing and Networking (MobiCom 2009)*. Beijing, China: ACM, September 2009.

[24] 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.

[25] 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.

[26] C. Lochert, B. Scheuermann, and M. Mauve, "Probabilistic Aggregation for Data Dissemination in VANETs," in *4th ACM International Workshop on Vehicular Ad Hoc Networks (VANET 2007)*. Montreal, Quebec, Canada: ACM, September 2007, pp. 1–8.

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

[28] F. Kargl, E. Schoch, B. Wiedersheim, and T. Leinmüller, "Secure and Efficient Beaconing for Vehicular Networks," in *5th ACM International Workshop on Vehicular Inter-Networking (VANET 2008), Poster Session*. San Francisco, CA: ACM, 2008, pp. 82–83.

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

[30] 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.

[31] C. Sommer, Z. Yao, R. German, and F. Dressler, "Simulating the Influence of IVC on Road Traffic using Bidirectionally Coupled Simulators," in *27th IEEE Conference on Computer Communications (INFOCOM 2008): IEEE Workshop on Mobile Networking for Vehicular Environments (MOVE 2008)*. Phoenix, AZ: IEEE, April 2008, pp. 1–6.

[32] C. Sommer, R. German, and F. Dressler, "Bidirectionally Coupled Network and Road Traffic Simulation for Improved IVC Analysis," *IEEE Transactions on Mobile Computing*, 2010, available online: 10.1109/TMC.2010.133.

[33] A. Cappiello, I. Chabini, E. Nam, A. Lue, and M. Abou Zeid, "A statistical model of vehicle emissions and fuel consumption," in *5th IEEE International Conference on Intelligent Transportation Systems (IEEE ITSC)*, September 2002, pp. 801–809.

[34] C. Sommer, R. Krul, R. German, and F. Dressler, "Emissions vs. Travel Time: Simulative Evaluation of the Environmental Impact of ITS," in *71st IEEE Vehicular Technology Conference (VTC2010-Spring)*. Taipei, Taiwan: IEEE, May 2010.