

Traffic Information Systems: Efficient Message Dissemination via Adaptive Beaconing

C. Sommer^{*†}, O. K. Tonguz[†] and F. Dressler[‡]

^{*}Computer Networks and Communication Systems, Dept. of Computer Science, University of Erlangen, Germany

^{*}Electrical and Computer Engineering Dept., Carnegie Mellon University, Pittsburgh, PA, USA

[†]Computer and Communication Systems, Institute of Computer Science, University of Innsbruck, Austria

christoph.sommer@informatik.uni-erlangen.de, tonguz@ece.cmu.edu,

falko.dressler@uibk.ac.at

Abstract—Traffic Information Systems (TIS) are one of the key non-safety application areas of Vehicular Ad Hoc Networks (VANETs). As such, TIS are much less delay sensitive compared to safety applications, which have recently attracted a lot of attention in VANET research. We propose a new message dissemination protocol, Adaptive Traffic Beacon (ATB), which is *fully distributed* and uses adaptive beaconing based on two key metrics: 1) the message utility; and 2) the channel quality. It is shown that adaptive beaconing leads to a much broader dissemination of messages (in terms of penetration rate) than flooding-based approaches, albeit at a slower rate. Adaptive beaconing thus seems to be much more suitable for TIS than flooding-based protocols.

I. INTRODUCTION

Information dissemination schemes in Vehicular Ad Hoc Networks (VANETs) are commonly categorized into two different groups, according to their envisioned application: safety and non-safety messaging. Much attention in VANET research has been focused on safety applications, which, being highly demanding in terms of message delay, present a challenging field of study. In these systems, the frequency of messages is low, but each message, e.g., an emergency brake warning, has to reach a maximum number of nodes in a given area within a very short time interval. After this time interval, the message essentially becomes useless [1].

Message dissemination schemes based loosely on flooding are therefore perfectly suited for the provision of safety applications. Such approaches have been enhanced to keep the network load within reasonable bounds and, more recently, to also be resilient to network disconnections [2]. In flooding based schemes, information about an event spreads outward from the point of origin much like a shock wave, and all network activity takes place at the wave front. This frees up network resources in the area of nodes that have already received an event, and allows them to dedicate their resources to the dissemination of new events. While formerly disconnected parts might rejoin a network at a location where the shock wave has already passed, it can be assumed that they will likely do so at a time where the message has already become useless to them.

This assumption, however, does not hold for *non-safety applications* of VANETs. Here, messages are, by their very nature, valid for a significantly longer time interval: up to

several minutes for Traffic Information Systems (TIS) [3]. Events are more frequent, but delay constraints are not as stringent [1]. Disseminating these messages in sparse and frequently-disconnected networks using flooding will therefore lead to large parts of the network being simply skipped over by the shock wave and nodes in these areas thus being completely unaware of the message content. For the domain of non-safety applications, schemes outside the domain of flooding-based approaches thus appear very attractive.

The class of schemes that provide distributed non-safety message dissemination is commonly further categorized in terms of whether the design is focusing on deployment in highway or in urban scenarios.

Highway scenarios typically exhibit much lower densities and offer a much more reliable topology, as roughly half of the vehicles are traveling in the same direction and vehicles exhibit a natural tendency to cluster [2]. At the same time, interconnection times with vehicles traveling in the opposite direction are extremely short. VANETs on highways thus exhibit both the properties of well connected and sparsely connected networks at the same time, which has led to them being characterized as exhibiting bipolar behavior. Finally, the distances that messages have to be disseminated, i.e., required hop counts, are comparatively high. Messages commonly need to reach as far back as the next exit to make sure vehicles are informed in time to pick another route.

Urban scenarios, on the other hand, present a completely different set of requirements and opportunities. Topology dynamics in this setting are much less predictable and a network will constantly oscillate between high-density, fully connected states when vehicles are queuing in front of a traffic light and low density, disconnected states when vehicles are driving. Moreover, in urban scenarios such potentially disconnected clusters of driving vehicles will frequently pass high-density clusters of vehicles, namely while crossing an intersection where other vehicles queue.

Other than in highway scenarios, VANETs in urban scenarios thus exhibit the properties of both a disconnected and a well-connected network within a very short time interval, but not necessarily at the same time. On the other hand, compared to highway scenarios, the region of interest for a given message is noticeably smaller and an event needs to go through less hops.

It can be argued [4], however, that such distinctions cannot be made for real-life systems. This means that any system for the dissemination of non-safety messages in VANETs will inevitably have to adapt to highly dynamic, heterogeneous environments and will likely exhibit suboptimal performance when designed with rigid assumptions about the environment in mind.

In fact, adaptive beaconing schemes for non-safety applications have become a hot topic in the VANET research community [5], [6] and numerous proposals have been made to adapt the transmit power or beacon rate; this is done either according to the current channel load, or based on assumptions about other vehicles' knowledge base contents or road topology. Yet, the basic conclusion is that wrong predictions are often worse than no predictions at all and that isolated approaches are unlikely to accurately capture the overall environment conditions [6]. Moreover, a conclusive evaluation of the benefits and drawbacks of these schemes, in particular when compared with flooding-based approaches, is still lacking, as evidenced by the recent debate that is surrounding them.

In this article – to fill this gap and the shortcomings of current schemes pointed out in [6] – we present an extended version of our work on adaptive beaconing [7], [8] including an evaluation of, and comparisons with, flooding approaches. We focus on a solution that captures both message utility and channel conditions in the past, present, and future, in order to adapt to highly dynamic environments that change from fully-connected to disconnected states in the time it takes a vehicle to pass queued vehicles at an intersection.

This adaptivity is achieved by nodes continuously sensing their surroundings in order to quickly and dynamically react to changes. As is common, this sensing is supported by the exchange of beacon messages – however, beaconing in this scheme is not employed for topology maintenance, the applicability of which is actively debated for VANET scenarios [5], [9].

Rather, beaconing is employed for data transport. Thus, no additional overhead is needed for gathering the channel metrics required for such an adaptive beaconing scheme, it being simply a byproduct of information exchange. Not maintaining topology information means that no preconceived assumptions about the deployment scenario need to be built into this system, nor does the system need to rely on it being able to correctly derive an abstract model of the environment it is operating in, as failure to do either would result in a mismatch between system behavior and environment. Lastly, with no topology information to maintain, participating nodes need not be assigned unique identifiers that might severely compromise the privacy of users. However, the beaconing protocol needs to carefully use the available capacity of the channel not to cause overload and packet loss [10].

The contribution of this work is thus twofold:

- we introduce Adaptive Traffic Beacon (ATB), a novel scheme for rapidly adapting the frequency of knowledge base exchange according to a comprehensive set of metrics, capturing both message utility and network conditions in the past, present, and future;

- we evaluate and compare related approaches from the domain of flooding with the proposed scheme under both synthetic and highly realistic conditions.

II. RELATED WORK

Numerous fully distributed solutions for broadcasting traffic information to neighboring vehicles, either periodically or triggered by new events, have been investigated in the past. They are supported by aggregation and other data pre-processing techniques, developed to optimize the quality of traffic information and to reduce the necessary communication load [3], [11], [12]. One of the most sophisticated solutions in this class is the Self-Organizing Traffic Information System (SOTIS) [3]. It stores information in the form of annotated maps with varying resolution, depending on the distance from the current position and the age of the information, performing information exchange using a specialized MAC protocol. This approach has been further elaborated and merged with ideas from the peer-to-peer domain to form PeerTIS [13]. Likewise, the MobTorrent [14] scheme has been reported, which also provides mobile Internet access using Roadside Units (RSUs).

Completely different solutions are investigated in the context of multi-hop broadcast based approaches, which are well suited for the dissemination of messages with delay bounds [2]. In order to reduce unnecessary broadcast transmissions, directional broadcast (as opposed to pure flooding) can be used [15]. 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 via extensive simulations in [10], showing that with increasing distance, the success ratio decreased quickly. Most recently, 2-hop beaconing has been described to acquire topology knowledge for opportunistic forwarding using the selected best target forwarder in the TO-GO scheme [11] and reliable beaconing has been studied in [5]. Still, 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, leading to the creation of REACT [16], which can skip individual beacon transmissions to support emergency applications. Furthermore, fundamental scalability criteria need to be considered in order to make the protocol applicable in the target scenario [12].

As a promising solution to the inherent sensitivity of beaconing protocol parameters to rapid changes in network topology and utilization, a protocol that can adapt to these changes according to a comprehensive set of metrics appears to be the next logical step [5], [6]. ATB goes beyond available broadcast-based solutions by adaptively updating the beacon frequency not only based on the importance of messages but mainly based on the available capacity of the wireless channel. The key idea is to send beacons as frequently as possible but without causing the channel to become overloaded.

III. ADAPTIVE TRAFFIC BEACON (ATB)

The main objective of ATB is to exchange information in knowledge bases by sending beacons as frequently as

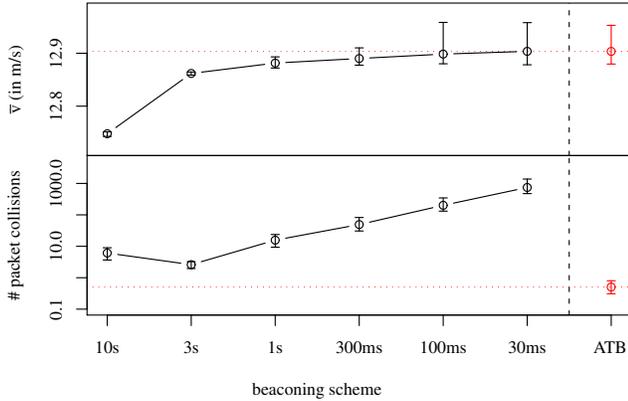


Fig. 1. TIS performance of static beaoning schemes, compared to adaptive beaoning (red), recorded in a 256 km² grid scenario. Effective average speed is as good, but packet collisions are several orders of magnitude lower.

possible, but to maintain a congestion-free wireless channel. ATB achieves this by employing two different metrics, the *channel quality* C and the *message utility* P , to calculate the beacon interval I with which to disseminate messages.

The benefits that adapting beacon intervals can bring to beaoning schemes are obvious: increasing the TIS performance of static beaoning schemes to a level comparable with that of adaptive beaoning schemes causes packet collisions to rise to intolerably high values. In fact, channel load can be up to several orders of magnitude higher [8], as illustrated in Figure 1.

In the following, we briefly introduce the different metrics ATB uses to assess channel quality and message utility. Each metric is derived by considering one particular measure of either channel quality or message utility and calculating its value relative to a fixed maximum value. We further illustrate how the metrics work together to adapt the beacon interval and present how nodes manage their local knowledge bases.

a) Channel quality C : The channel quality is estimated by means of three metrics, which are indicative of network conditions in the past, present, and future, respectively. First, a node observes the number of collisions on the channel, deriving a value K which is a measure of past channel conditions. We made our protocol very sensitive to this metric to prevent overload situations. Secondly, a node continuously measures the Signal to Noise Ratio (SNR) on the channel to derive S , which reflects current channel use. Obviously, this is only an indicator for the channel quality. If the car is far away, thus, the SNR being small, we anticipate a larger beacon interval, hoping to see other cars beaoning first. This is in line with findings published for example in [2]. Lastly, a node observes other nodes' beacons, deriving a measure for the number of neighbors N and thus enabling it to factor in, to a certain degree, the outcome of channel access in the near future.

Based on these metrics, which capture the quality of the channel in the past, present, and future, a node is able to derive a metric of the overall channel quality C , which is a linear combination of K , S , and N , ranging in the interval $[0, 1]$ (lower values describing a better channel quality). In order

to tune the reactivity of the protocol, flexible weight should be given to the estimation of future channel use N , while equal weight should be given to the more immediate metrics K and S .

b) Message utility P : The message utility P is derived from two metrics: First, a node accounts for the distance of a vehicle to an event as D_e , which is the most direct indication of message utility. Secondly, it accounts for message age A , thus allowing newer information to spread faster. Both D_e and A are of equal value for determining the compound utility metric P . As part of ongoing research on RSU supported operation of VANETs, we are aiming to further expand the calculation of P to take into account how useful a particular message might be in the presence of RSUs, as well as its utility to nearby RSUs. Similar to C , the value of P can also range from 0 to 1, lower values describing a higher priority messages.

c) Beacon interval calculation: Based on the two aforementioned metrics, C and P , ATB continuously adapts the beacon interval in a range from I_{min} to I_{max} . As the channel quality metric C in turn depends on the value of I that was chosen by nearby vehicles, ATB exhibits some properties of a *self-organizing system* [17]: on a macroscopic scale, vehicles participating in the VANET will independently arrive at beacon intervals that enable them to use the shared channel commensurate to their own and other nodes' needs. Hence, proper rules at the local level (car level) lead to emergent behavior at the global level.

ATB adjusts I such that it becomes minimal only for the highest message utility and the best channel quality. In all other cases, channel use is reduced drastically, allowing uninterrupted use of the channel by other applications. In addition, the relative impact of both parameters w_I is designed to be configurable, e.g., in order to calibrate ATB for different MAC protocols.

In our experiments, we used $w_I = 0.75$, i.e., weighting the channel quality higher than the message priority. That means that the beacon interval is very sensitive to the conditions of the radio channel. As described in a more detailed technical report [7], already a few collisions will cause ATB to backoff in order to efficiently use the remaining capacity of the wireless channel. This also means that ATB inherently can co-exist with other applications using the same channel.

In our implementation this is achieved by re-evaluating C and P and deriving a new beacon interval I after each beacon sent or received, according to the following equation – please note that I is thus in the range $[I_{min}, I_{max}]$.

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

That is, if both C and P are low, i.e. the channel is free and there are high priority messages, the resulting beacon interval I is low. If either C or P is high, then the result is somewhere in the middle; and if both are high, the resulting beacon interval is close to I_{max} , i.e., it is high. We also use a higher weight for C compared to P as the channel quality is clearly more important. The quadratic form is used to make ATB very sensitive in “bad” situations (e.g., channel overload), whereas to only slightly adapt I in the “best” case scenarios.

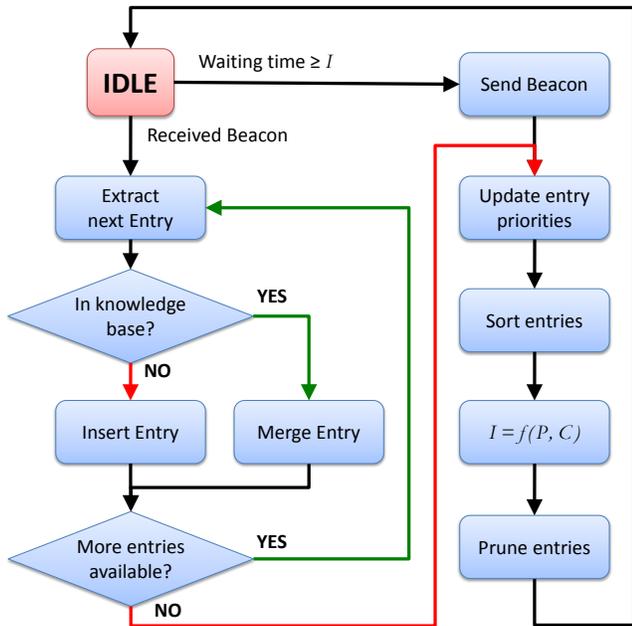


Fig. 2. Core functionality of the ATB protocol. Vehicles update a local knowledge base with entries from received beacons. The interval between beacons is continuously adjusted based on P and C .

A detailed sensitivity analysis of all these parameters is presented in [7].

d) Knowledge base management: In order to maintain the scalability of the TIS, the transmission of irrelevant information needs to be suppressed and each knowledge base needs to contain all received traffic information in aggregated form [12]. The operation of ATB, however, is independent from the scheme used for the selection of knowledge base entries and their aggregation, so any of the numerous approaches in the relevant literature can be used in an implementation.

As a baseline for our evaluation of the message dissemination characteristics, we implemented a simplistic scheme: The knowledge base stores only the most recent information for each route segment, i.e., each new event either updates an existing record or it is appended to the knowledge base. A garbage collection process continuously expunges entries that are older than a configurable timeout. Each node prioritizes available information according to the age of an entry, as well as the distance to the event. Using the calculated priorities, a node can then generate beacon messages by selecting as many entries as there is room in a single link layer frame from the top of the list, i.e., those with the highest priority. The most important message is used to calculate the message priority for the beacon interval. This way, the frame size is optimally used and problems with stateful handling of messages split into multiple frames are inherently avoided. Nodes that receive these beacons can then in turn update their knowledge bases and beacon intervals according to the algorithm illustrated in Figure 2.

Clearly, the presented handling of the knowledge base is not yet optimal – the key aspect of ATB is the efficient use to the available capacity of the wireless channel.

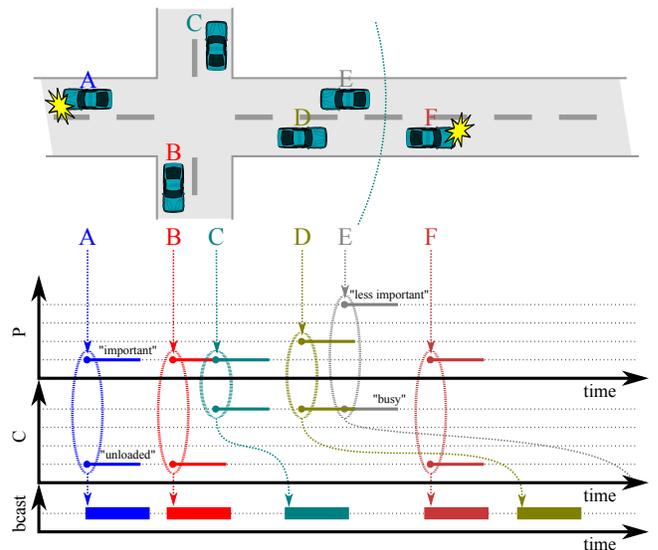


Fig. 3. Behavior of ATB on the microscopic level, plotted in the style of a timing diagram. One new event each is to be disseminated by car A, then car F. Metrics P and C work in harmony to keep beacon frequency high, but channel use low, allowing car F to instantly broadcast.

IV. BEHAVIOR ON A MICROSCOPIC LEVEL

As the behavior of ATB on the system level is dictated by the local behavior of individual nodes, we first illustrate how ATB behaves on the microscopic level by means of a simple example. Assume that ATB is deployed in the scenario shown in Figure 3. At the beginning of the observed time interval, car A senses a new event, which it will distribute to the nearby cars B, C, D, and E. Suppose some time later, car F, which is just out of range of car A, will sense a new event, too, which it will proceed to distribute.

Figure 3 shows how this scenario would play out in the style of a timing diagram. Please note that small values for C and P denote a free channel and a high message priority, respectively. From top to bottom, we plot the following four metrics of ATB: first, the time when an entry is inserted into a vehicle's knowledge base, i.e., the time when a new event was created (cars A and F) or when a broadcast was received (cars B, C, D, and E). Secondly, we plot the values of the message utility metric P and the channel quality metric C , which are used by ATB to calculate the beacon interval. Lastly, we indicate the time when an event is broadcast on the channel.

As can be observed, car A registers the new event at maximum values of both message utility P and channel quality C . Car A therefore chooses its minimum beacon interval for its next transmission, broadcasting the event almost immediately.

For the sake of example, we assume that car B is the first vehicle to process the reception of this broadcast (taking operating system and other issues into account), i.e., it becomes ready to rebroadcast first, followed shortly by car C. Thus, car B is free to transmit its next beacon instantly – unlike car C, which observes these two broadcasts within a short time interval and thus deduces a low value for the channel quality metric C . Following the presented algorithm, it reacts by increasing its beacon interval, postponing its broadcasting

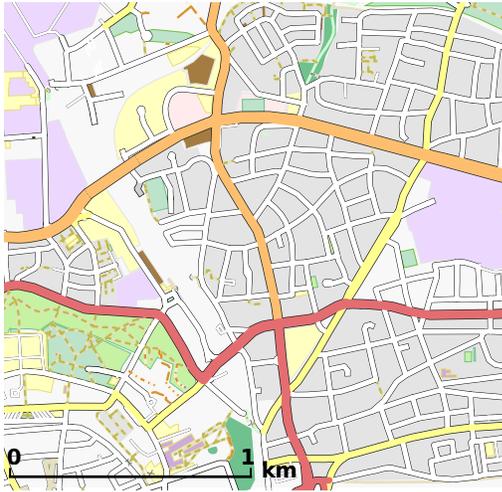


Fig. 4. Scenario used for the performance evaluation and comparison: fully-detailed 4 km² section of Ingolstadt, containing a heterogeneous mixture of high- and low-density areas, along with a crossing of two federal highways, B 13 and B 16a.

of the event.

This broadcast, however, interferes with the desire of cars D and E to broadcast the event themselves. Both cars now derive not only a low channel quality metric C , but (based on the increased event distance and age) also a lower message utility metric P . This leads to them postponing their broadcast even further, leaving the channel idle for now.

Thus, car F, which needs to disseminate *a new event*, is presented with ideal channel conditions and, thus, arrives at a maximum value for the channel quality metric C . Taken together with the message utility metric P , which assumes its maximum value because of the event being newly created, this leads to car F choosing its minimum value for the beacon interval. It can hence instantly broadcast a beacon containing the new event. Of course, each message may include more than one event.

This example thus illustrates how the chosen metrics work in harmony to keep nodes' beacon frequency high whenever necessary, but the number of collisions on the channel low, helping ATB adapt to highly dynamic network conditions – both proactively and reactively.

V. COMPARISON WITH FLOODING-BASED APPROACHES

We implemented ATB in our *Veins*¹ simulation environment, which is based on OMNeT++ for event-driven network simulation and SUMO for road traffic microsimulation [18]. Both simulators have been extended by modules that allow the road traffic simulation to communicate with its network simulation counterpart. 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. In our implementation, the knowledge base is checked after processing each received beacon to identify events on the current route of the vehicle. If an incident is found, an alternative road traffic 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.

In this set of evaluations, we evaluate the performance of ATB in terms of dissemination speed, the classic domain of flooding-based schemes. Other metrics such as the collisions on the wireless channel have been also recorded. We only report the most interesting results. We therefore compare its performance to that of the most closely related and most recent of these schemes, DV-CAST [2], which operates in a fully distributed manner, can keep the network load within reasonable bounds, and is resilient to network disconnections. Even though the concept of DV-CAST was developed for highway scenarios, its flooding algorithm makes no assumptions about the underlying road topology, and, hence, can be deployed in any traffic scenario. In particular, DV-CAST incorporates most recent findings on flooding approaches in dense scenarios to eliminate the broadcast storm problem, and adaptively switches to a DTN-like solution in sparse scenarios. We implemented DV-CAST in *Veins*, calibrating its behavior to that of its original NS-2 model.

The scenario we use for this evaluation simulates traffic of varying density, from 14 to 170 vehicles per km², in the city of Ingolstadt, collating data points from up to 100 independent runs to ensure the statistical significance of results. The road network was based on the comprehensive road topology and attribute database available from the OpenStreetMap project and adapted to reflect realistic intersection management and timing. Traffic was generated by randomly selecting source-destination pairs and iteratively applying the dynamic user assignment algorithm implemented in SUMO until it reported a stable, optimal distribution of flows.

In the evaluation, we focus on the 4 km² Region of Interest (ROI) shown in Figure 4, which contains a heterogeneous mixture of high- and low-capacity roads, traffic lights and unregulated intersections, as well as high- and low-density areas. While traffic is simulated in the whole city of Ingolstadt to avoid border effects, only vehicles within the ROI are considered to be participating in the network. In order to determine how realistic this generated road traffic was, we implemented induction loops in the ROI and compared the measured values with real values provided courtesy of the local authorities. We observed that the range of traffic densities spans from off-peak densities, e.g., the beginning of the morning rush hour at 6:00 a.m., all the way down to sparse traffic at low equipment rates.

All vehicles use ATB to exchange information about obstructions, the parameters of ATB being configured as follows: minimum and maximum beacon interval are set to $I_{min} = 30$ ms and $I_{max} = 60$ s, the weighting w_I of channel quality set to 0.75, the relative weight of K and S vs. N to 2. Vehicles generate a traffic incident event after being blocked for 10 s and set its validity to 120 s. An IEEE 802.11b NIC transmitting at 11 Mbit/s for exchanging beacons and a radio channel to provide a beacon range of up to 180 m are modeled by the OMNeT++ *INET Framework*. In future work, we plan to further compare the results of these evaluations with ones obtained using an IEEE 802.11p NIC, but expect no qualitative

¹<http://veins.car2x.org/>

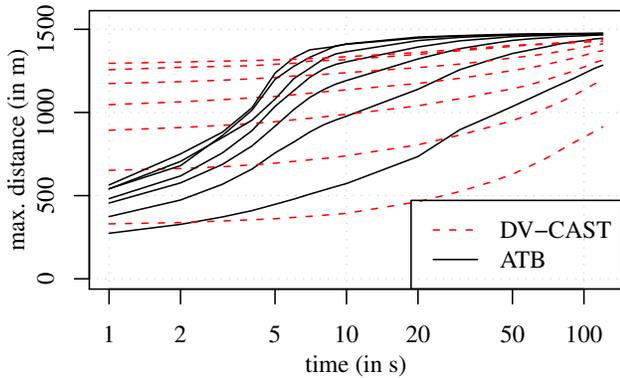


Fig. 5. Dissemination distance vs. time for 7 groups of traffic densities with means of 24, 51, 73, 98, 119, 151, and 170 vehicles per km^2 . Beacons achieves higher dissemination distances in the long run. Flooding outperforms beacons in terms of speed.

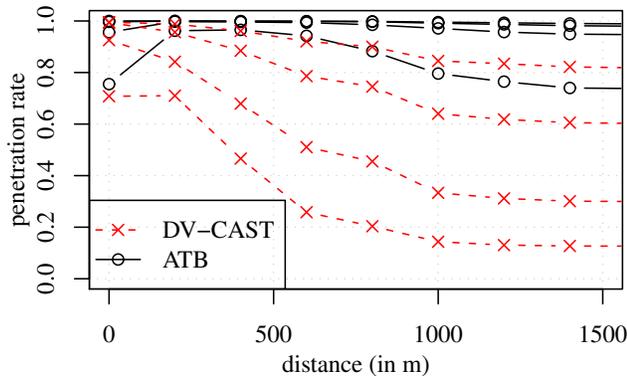


Fig. 6. Dissemination quality vs. distance for 4 groups of traffic densities with means of 19, 49, 97, and 159 vehicles per km^2 . Flooding leaves temporarily-disconnected clusters unaware of events.

difference; quantitatively, there will certainly be a difference, which applies to all the discussed protocol variants.

We evaluate the message dissemination speed by triggering the transmission of a warning message at the intersection of the two federal highways crossing this area, B 13 and B 16a (shown in dark red on Figure 4). We keep track of whether each of the simulated cars received the event and record its distance to the event's origin for the duration of the event's validity. This metric specifically includes cases where the event was not picked up by any car. Moreover, it includes the progress of message dissemination achieved simply by virtue of cars moving across the map.

As is common, we first examine the maximum distance covered by each message as time progresses. Figure 5 illustrates how ATB compares to DV-CAST in this respect, plotting the mean value of this distance for seven groups of the simulated traffic densities. The interquartile range of results remained below 4% at all times. We observe that in terms of dissemination speed flooding clearly outperforms our beacons approach independent of the traffic density, having already bridged several hundreds of meters before the first distance sample was taken. Yet, we also observe that, due to the extremely high network dynamics in this scenario, the beacons approach is more robust, leading to more

widespread message propagation in the long run. Nevertheless, it is not until several seconds after an event took place that ATB reaches the same dissemination distance as DV-CAST.

However, we argue that examining the maximum dissemination distance alone does not capture the benefit beacons brings to frequently disconnected networks: there is no clearly defined shock wave (the front of epidemically spreading information) where all message transmissions take place, but rather a broad continuum where nodes frequently exchange parts of their knowledge bases. This means that disconnected clusters re-joining the network, even after several seconds of disconnection, still have a very good chance of receiving and, in turn, being able to disseminate these missed events.

We illustrate this in Figure 6, plotting for four groups of the simulated traffic densities the ratio of nodes that received an event in a certain area, depending on its size. The interquartile range of results remained below 5% at all times. We observe that, in particular at low densities and considering a large area around the point of origin, adaptive beacons achieves a noticeably higher event penetration ratio, whereas flooding leaves large clusters of nodes unaware of an event. However, we also observe how the lower speed of event dissemination of ATB leads to its maximum of penetration ratios being shifted some hundreds of meters away from the event source.

VI. CONCLUSION

The dissemination of information in VANETs in a completely distributed fashion faces unique challenges in the context of highly dynamic, heterogeneous environments. For a beacons based protocol to function in such scenarios, it is indispensable that it adapts to changing conditions according to complementary metrics of channel conditions and message utility, rather than considering only one of these metrics. We demonstrate that, by considering message utility, observations of past channel conditions, measurements of current conditions, and tentative estimations of future conditions, such a protocol can successfully adapt to highly dynamic environments both proactively and reactively.

We further show that adaptive beacons leads to a much broader dissemination of messages, the predominant metric for non-safety applications. On the other hand, flooding clearly outperforms beacons in terms of dissemination speed, the predominant metric for safety applications. Adaptive beacons and flooding based schemes thus offer unique benefits and drawbacks, depending on the application scenario. This distinction becomes even more pronounced in scenarios with low traffic density or low equipment rates, where the shock wave like dissemination of messages in flooding based schemes simply skips temporarily-disconnected clusters.

REFERENCES

- [1] K. Dar, M. Bakhouya, J. Gaber, M. Wack, and P. Lorenz, "Wireless Communication Technologies for ITS Applications," *IEEE Communications Magazine*, vol. 48, no. 5, pp. 156–162, May 2010.
- [2] O. K. Tonguz, N. Wisitpongphan, and F. Bai, "DV-CAST: A distributed vehicular broadcast protocol for vehicular ad hoc networks," *IEEE Wireless Communications*, vol. 17, no. 2, pp. 47–57, April 2010.
- [3] 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.

- [4] F. Bai and B. Krishnamachari, "Exploiting the Wisdom of the Crowd: Localized, Distributed Information-Centric VANETs," *IEEE Communications Magazine*, vol. 48, no. 5, pp. 138–146, May 2010.
- [5] F. J. Ros, P. M. Ruiz, and I. Stojmenovic, "Reliable and Efficient Broadcasting in Vehicular Ad Hoc Networks," in *IEEE VTC2009-Spring*, Barcelona, Spain: IEEE, April 2009, pp. 1–5.
- [6] R. K. Schmidt, T. Leinmüller, E. Schoch, F. Kargl, and G. Schäfer, "Exploration of adaptive beaconing for efficient intervehicle safety communication," *IEEE Network Magazine*, vol. 24, no. 1, pp. 14–19, January 2010.
- [7] C. Sommer, R. German, and F. Dressler, "Adaptive Beaconing for Delay-Sensitive and Congestion-Aware Traffic Information Systems," University of Erlangen, Dept. of Computer Science, Technical Report CS-2010-01, January 2010.
- [8] C. Sommer, O. K. Tonguz, and F. Dressler, "Adaptive Beaconing for Delay-Sensitive and Congestion-Aware Traffic Information Systems," in *2nd IEEE Vehicular Networking Conference (VNC 2010)*. Jersey City, NJ: IEEE, December 2010, pp. 1–8.
- [9] D. Rossi, R. Fracchia, and M. Meo, "VANETs: Why Use Beaconing at All?" in *IEEE International Conference on Communications (IEEE ICC 2008)*, May 2008, pp. 2745–2751.
- [10] 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.
- [11] K. Lee, U. Lee, and M. Gerla, "Geo-Oppportunistic Routing for Vehicular Networks," *IEEE Communications Magazine*, vol. 48, no. 5, pp. 164–170, May 2010.
- [12] B. Scheuermann, C. Lochert, J. Rybicki, and M. Mauve, "A Fundamental Scalability Criterion for Data Aggregation in VANETs," in *ACM MobiCom 2009*. Beijing, China: ACM, September 2009.
- [13] 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.
- [14] 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.
- [15] 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.
- [16] 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.
- [17] F. Dressler, *Self-Organization in Sensor and Actor Networks*. John Wiley & Sons, December 2007.
- [18] C. Sommer, R. German, and F. Dressler, "Bidirectionally Coupled Network and Road Traffic Simulation for Improved IVC Analysis," *IEEE Transactions on Mobile Computing*, vol. 10, no. 1, pp. 3–15, January 2011.



Ozan K. Tonguz (tonguz@ece.cmu.edu) is a tenured full professor in the Electrical and Computer Engineering Department of Carnegie Mellon University (CMU), Pittsburgh, PA. He currently leads substantial research efforts at CMU in the broad areas of telecommunications and networking. He has published about 300 papers in IEEE journals and conference proceedings in the areas of wireless networking, optical communications, and computer networks. He is the author (with G. Ferrari) of the book *Ad Hoc Wireless Networks: A Communication-Theoretic Perspective* (Wiley, 2006). He co-founded Virtual Traffic Lights (VTL), LLC, a CMU spin-off, in December 2010, which specializes in providing solutions to several transportation problems, such as safety and traffic information systems, using V2V and V2I communications paradigms. His current research interests include vehicular ad hoc networks, wireless ad hoc and sensor networks, self-organizing networks, bioinformatics, and security. He currently serves or has served as a consultant or expert for several companies, major law firms, and government agencies in the United States, Europe, and Asia.



Falko Dressler (dressler@informatik.uni-erlangen.de) is a full professor for computer science at the University of Innsbruck. Dr. Dressler is an Editor for journals such as Elsevier *Ad Hoc Networks*, *ACM/Springer Wireless Networks (WINET)*, and Elsevier *Nano Communication Networks*. Among other, Dr. Dressler wrote the textbooks *Self-Organization in Sensor and Actor Networks*, published by Wiley in 2007. Dr. Dressler is an IEEE Distinguished Lecturer in the fields of inter-vehicular communication, self-organization, and bio-inspired networking. Dr. Dressler is a Senior Member of the IEEE (COMSOC, CS, VTS) as well as a Senior Member of ACM (SIGMOBILE), and member of GI (KuVS). His research activities are focused on adaptive wireless networking and self-organization methods addressing issues in wireless ad hoc and sensor networks, inter-vehicular communication systems, bio-inspired networking, and adaptive network security techniques.



Christoph Sommer (christoph.sommer@informatik.uni-erlangen.de) received his M.Sc. in computer science in 2006 and is now about to finish his Ph.D. at the Computer Networks and Communication Systems group at the Department of Computer Science, University of Erlangen. At the time of writing, Christoph was a visiting scholar with the Electrical and Computer Engineering Department of Carnegie Mellon University. His research is focused on questions regarding efficiency and security aspects of Car-to-X communication in heterogeneous environments.