Position-Based Data Traffic Prioritization in Safety-Critical, Real-Time Vehicle-to-Infrastructure Communication

Annette Böhm and Magnus Jonsson

Böhm A, Jonsson M. Position-based data traffic prioritization in safety-critical, real-time vehicle-to-infrastructure communication. In: IEEE International Conference on Communications Workshops, 2009. ICC Workshops 2009. IEEE; 2009. p. 1-6.

DOI: http://dx.doi.org/10.1109/ICCW.2009.5208065

Böhm, A. and M. Jonsson, "Position-based data traffic prioritization in safety-critical, real-time vehicle-to-infrastructure communication," Proc. IEEE Vehicular Networking and Applications Workshop (VehiMobil 2009) in conjunction with the IEEE International Conference on Communications (ICC), Dresden, Germany, June 14, 2009.

Position-Based Data Traffic Prioritization in Safety-Critical, Real-Time Vehicle-to-Infrastructure Communication

Annette Böhm and Magnus Jonsson
CERES (Centre for Research on Embedded Systems)
Halmstad University, Halmstad, Sweden
{Annette.Bohm, Magnus.Jonsson}@hh.se
http://ceres.hh.se

***Abstract* – Future active-safety applications in vehicular networks rely heavily on the support for real-time inter-vehicle communication. The Medium Access Control (MAC) mechanism proposed for the upcoming IEEE 802.11p standard intended for Intelligent Transport Systems (ITS) applications does not offer deterministic real-time support, i.e., the channel access delay is not upper bounded. We therefore propose a vehicle-to-infrastructure (V2I) communication solution extending IEEE 802.11p, by introducing a collision-free MAC phase with an enhanced prioritization mechanism based on vehicle positions and the overall road traffic density. A road side unit using a polling mechanism is then able to provide real-time support such that it can guarantee collision-free channel access within its transmission range. Part of the bandwidth remains unchanged such that best-effort services like ongoing vehicle-to-vehicle (V2V) applications may continue. Our solution guarantees that all communication deadlines of the V2I applications are met, while minimizing the required length of the collision-free phase. This in turn maximizes the amount of bandwidth available for best-effort services and ongoing V2V applications. The position-based prioritization mechanism further improves the throughput of both real-time and best-effort data traffic by focusing the communication resources to the most hazardous areas. The concept is evaluated analytically based on a realistic task set from a V2I merge assistance scenario.**

I. INTRODUCTION

In proactive Intelligent Transport Systems (ITS) applications, information is gathered from surrounding vehicles or from static roadside units (RSU) to provide both driver and vehicle with an improved understanding of the current traffic situation and its potential hazards. These types of applications rely heavily on the timely delivery of safety-critical real-time data. We define real-time data as deadline-dependent data traffic that must be delivered to its destination before a given deadline in order to be of use for its, often safety-critical, target application.

Wireless Access in the Vehicular Environment (WAVE) is considered a key enabling technology for ITS safety applications. The details of the physical and link layers are currently being standardized as IEEE 802.11p [1], a variation of the IEEE 802.11 Wireless LAN standard. However, the contention-based Medium Access Control (MAC) method of IEEE 802.11p does not provide support for deadline-dependent real-time data traffic, required by proactive traffic safety applications. This is a major drawback pointed out in several studies, e.g., [2]-[3]. Previous solutions to this problem suggest simply replacing the MAC method, i.e., changing the standard, or introducing data traffic smoothing in the upper layers to reduce the occurrence of the problem with data traffic collisions.

We base our non-intrusive solution on a RSU and an additional real-time layer, placed on top of the unaltered IEEE 802.11p protocol, such that a deterministic, polling-based MAC method is introduced. This real-time layer divides the available bandwidth into two different phases: in the contention-based phase (CBP) nodes compete for access to the medium according to the original IEEE 802.11p MAC method, but in the collision-free phase (CFP) the RSU assumes the responsibility for scheduling the data traffic and polls the mobile nodes for data. As previously presented by the authors [4], this real-time layer guarantees the timely delivery of safety-critical data, as long as the CFP is long enough. We apply real-time schedulability analysis to determine the minimum length of the CFP such that all deadlines are guaranteed and the remaining bandwidth is maximized to cater for best-effort services. In this paper, we further enhance our solution by introducing geographical priority zones around an area of hazard (e.g. an intersection), defining different levels of priority and, thereby, determining a vehicle's communication parameters based on its geographical position. We adapt the schedulability analysis to our prioritization scheme in order to focus communication resources to vehicles and data traffic classes that need them the most.

Mak *et al.* [5] propose a similar V2I MAC solution with a polling-based phase for safety data exchange and a phase for non-safety services. However, since no schedulability analysis is made, no guarantees for timely delivery of real-time data can be provided. Further, [5] assumes the availability of several separate data channels for ITS applications, which is unrealistic since the dedicated ITS frequency band often is very limited (30

MHz is currently reserved in Europe). In [6] different priority classes based on the urgency of the message are suggested, such that the most critical data packets are sent with the highest update frequency, but no connection between the vehicle's location and the criticality of the data is made. The position of the vehicle has been used in frequency division MAC methods, e.g. [7], where geographical zones are coupled to certain frequency bands, again assuming availability of several frequency bands.

The rest of the paper is organized as follows, section II introduces the MAC method coupled with the position-based prioritization mechanism and gives an overview of the example scenario. In section III, details of the real-time analysis are given. An analytical evaluation based on a merge assistance example scenario is presented in section IV before the paper is concluded in section V.

II. PROTOCOL DETAILS

Traffic accidents are typically concentrated to certain road sections such as highway entrance ramps, intersections or areas with road work. In such areas it is justified to install a (static or temporary) RSU offering V2I communication. Further, if vehicular networks also provide other, non-safety-critical data exchange it may partly cover the costs of infrastructure deployment and maintenance. An integration of both types of data traffic is therefore desirable. Our proposed V2I system hence involves a RSU in charge of scheduling real-time data traffic from proactive safety applications, while still reserving bandwidth for other, non-safety-critical data exchange. We assume one 10 MHz frequency channel per driving direction, shared by safety-critical and best-effort services. For financial reasons, a seamless coverage of RSU cannot be expected in the near future and we therefore see our solution as an addition to V2V applications at particularly dangerous road sections. A merge assistance application at a highway entrance ramp was chosen as the type scenario for protocol description and evaluation in this paper and its details are given below.

A. Merge Assistance Application Scenario

The merge assistance scenario is based on V2I communication involving a RSU at a highway entrance ramp that supports both entering and passing vehicles with heterogeneous communication services covering a variety of Quality of Service (QoS) classes. Through short, periodic heartbeat messages from the vehicles, the RSU is informed about individual positions, states and intentions. The RSU broadcasts recommendations concerning the merging process back to the vehicles. Moreover, vehicles passing a RSU are provided with updates on e.g. current road, traffic or weather conditions. Vehicle heartbeats, RSU recommendations and road information updates are all considered safety-critical. At the same time, bandwidth should be reserved to support ongoing V2V applications or various, V2I-based, best-effort data traffic classes as e.g. digital map updates or commercial services. All V2V applications (even safety-critical ones as e.g. collision-avoidance) are best-effort as they employ the contention-based MAC method of 802.11p. We can only make sure that these applications can be maintained while the vehicle moves through a region with a RSU. In order to make themselves known to the RSU when entering within transmission range, vehicles send out hello messages. Once recognized by the RSU, vehicles can be integrated into a polling scheme. The data traffic classes are summarized in Table 1.

TABLE 1: DATA TRAFFIC CLASSES

	Safety-Critical	**Best-Effort**
Vehicle → RSU	▪ Vehicle heartbeat	▪ Hello messages
RSU → Vehicle	▪ RSU recommendation ▪ Road information updates	▪ Non-safety-critical RSU-based services
Vehicle → Vehicle		▪ RSU-independent V2V communication

B. Protocol Details

In IEEE 802.11p, nodes compete for access to the medium according to the Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) MAC method. Further, 802.11p uses Enhanced Distributed Channel Access (EDCA) with QoS support first introduced in IEEE 802.11e [8]. Despite the different QoS classes enforcing different waiting times (inter-frame spacing) before transmission attempts, collisions within a class are still possible and therefore, no timing guarantees can be given. Unlike most IEEE 802.11 standards, 802.11p does not provide an optional CFP, centrally controlled by an access point through polling. We propose to reintroduce the CFP by placing a real-time layer on top of 802.11p. This layer schedules the real-time data traffic before handing down best-effort packets to the unaltered 802.11p protocol. In our layer, time is divided into superframes (SF) of fixed size, each consisting of a CBP and a CFP, Fig. 1. The CFP needs support from a RSU that takes responsibility for scheduling the data traffic and polling the mobile nodes for data. A mobile node is thereby assigned sole right to use the channel for a specified amount of time. A RSU can also assign these rights to itself without polling. Based on the RSU's knowledge of the QoS demands in the network, real-time data for the next superframe is scheduled according to Earliest Deadline First (EDF) [9]. As no collisions can occur, the CFP is deterministic and suitable for the delay-sensitive real-time data traffic needed in many ITS safety applications. A certain percentage of the SF is set aside for the CBP where nodes compete for access using the regular CSMA/CA. The ratio between the CBP and the CFP within an SF can be determined by schedulability analysis.

In order for the data traffic of a vehicle to be scheduled, the RSU needs to know about its presence. Vehicles therefore announce their presence through hello messages sent during the

CBP. There is a small probability that a vehicle fails to become part of the RSU polling scheme before leaving the transmission range, due to competing for access in the CBP. This is another reason for attempting to keep the CBP as long as possible, although the probability of packet collision can be reduced by giving hello messages the highest priority defined for the CBP in 802.11p. However, it should also be noted that in the initial stages of an ITS introduction, it cannot be assumed that all vehicles are equipped with the proper communication technology. Detection and integration of these vehicles into the merge assistance application therefore need done using various sensors on ITS-enabled vehicles. A vehicle that fails to communicate its hello messages to the RSU would thus be treated as an unequipped vehicle by the application, with the difference that it still would be able to receive broadcasts from the RSU and take part in best-effort data transmissions or ongoing V2V applications.

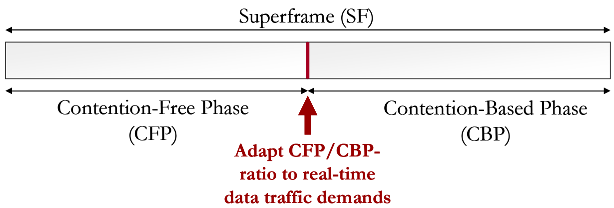

Figure 1: Adaptable ratio between CFP and CBP.

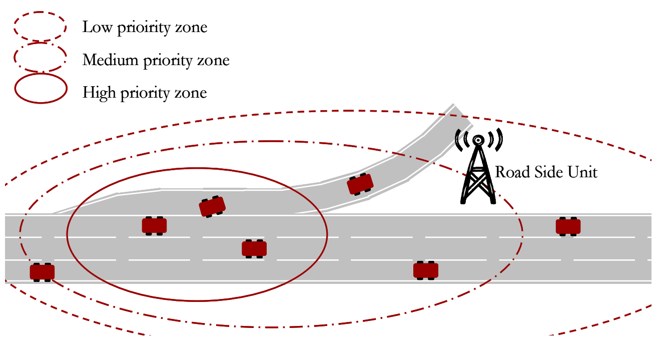

Figure 2: Position-based prioritization zones.

Through vehicle heartbeats, the RSU knows about the individual position for each vehicle in relation to the zone of hazard. We have defined a number of zones of descending priority around the hazard zone, Fig. 2. Each zone defines the period and deadline for real-time data sent from a vehicle residing in this zone. The closer to the hazard zone a vehicle is positioned, the shorter are its period and deadline (which through EDF scheduling directly influences its priority). Decreasing the amount of bandwidth used in less important areas of the RSU's transmission range, frees resources that can either be used to increase the data transfer in prioritized areas or to accommodate an increased overall number of vehicles. The unreliable nature of the wireless communication channel may justify extra resources for redundancy, e.g. sending two identical packets of real-time data in order to increase transmission reliability or providing more time for hello messages.

III. Real-Time Analysis

Real-time schedulability analysis, first introduced by [10] in the field of processor scheduling, contains two tests to determine if the bandwidth reserved for the CFP guarantees that no real-time deadlines are missed. Firstly, the utilization of the wireless channel must not exceed one, and secondly, the workload function, i.e., the sum of the transmission times for all packets with an absolute deadline less than or equal t, must be less than or equal to t, [11]. If the schedulability tests are passed, the CFP may be reduced by a predefined number and the analysis repeated until the minimum CFP is found. Durations and deadlines of all real-time messages are thus needed and a schedulability analysis is therefore performed for a certain task set. In our case the task set is derived from the merge assistance application. Each data traffic class is defined by a period, a deadline, a maximum packet length and a direction of transmission, vehicle to RSU ($V \rightarrow RSU$), initiated by polling, or RSU to vehicle ($RSU \rightarrow V$), without polling.

A. Definitions and Timing Analysis

Let our scenario contain a total of Q different, logical real-time (RT) channels, each representing a packet flow with real-time requirements from a specific source and for a specific traffic class. A RT channel is defined by its source, destination, period P_i, packet length L_i and deadline D_i, where $i = 1, 2, \ldots, Q$. Any packet transmission is preceded by a SIFS, of duration T_{SIFS}, set as a system parameter. Assuming a bit rate R, the total transmission time, T_i, of a packet is:

$$T_i = \begin{cases} \dfrac{L_i}{R} + T_{SIFS} & \text{if} \quad RSU \rightarrow V \\ \dfrac{L_i + L_{poll}}{R} + 2T_{SIFS} + 2T_{prop_i} & \text{if} \quad V \rightarrow RSU \end{cases} \qquad (1)$$

with a polling packet of length L_{poll} and propagation delay T_{prop}. As stated, a superframe of length T_{SF} consists of a contention-based phase of length T_{CBP} and a collision-free phase of length T_{CFP}. Only T_{CFP} can be used for transmitting real-time data. Towards the end of T_{CFP}, there might not be enough time for a full packet to be scheduled. This is accounted for by reducing T_{CFP} by $T_{blocking}$, the transmission time of the longest packet specified by any RT channel:

$$T_{blocking} = \max_i (T_i)\,. \qquad (2)$$

The remaining fraction F_{CFP} of the total bandwidth available for RT traffic is:

$$F_{CFP} = \frac{T_{CFP} - T_{blocking}}{T_{SF}}. \quad (3)$$

The original bit rate, R, is thereby reduced to an experienced bit rate of R_{CFP}, while the experienced transmission time, E_i, of a real-time packet is extended accordingly:

$$R_{CFP} = R \cdot F_{CFP}, \quad (4)$$

$$E_i = \frac{T_i}{F_{CFP}}. \quad (5)$$

The immediate transmission of a packet may be delayed by T_{CBP}. If the time before the start of a CBP is too short to accommodate another packet, the waiting time is increased by T_i. Further, the worst-case blocking delay due to an ongoing lower-priority packet transmission is $T_{blocking}$. The original deadline, D_i, of a packet is therefore reduced to an adapted deadline, D'_i. Since the propagation delay is not included in T_i for $RSU \rightarrow V$ traffic, this must be accounted for, as:

$$D'_i = \begin{cases} D_i - T_{CBP} - T_{blocking} - T_i - T_{prop} & \text{if} \quad RSU \rightarrow V \\ D_i - T_{CBP} - T_{blocking} - T_i & \text{if} \quad V \rightarrow RSU \end{cases} \quad (6)$$

B. Real-Time Schedulability Analysis

The analysis is done in two steps. A necessary, but not sufficient, condition is that the utilization U of the wireless link must never exceed one. According to EDF scheduling theory [9], the utilization of periodic real-time traffic is:

$$U = \sum_{i=1}^{Q} \frac{E_i}{P_i}. \quad (7)$$

The second step entails the workload function, $h(t)$, which must be less than or equal to t. The following definitions are needed:

- The hyperperiod (HP) is the least common multiple of all periods of all RT channels.
- The busyperiod (BP) is any interval within the HP during which the link is not idle.

Then $h(t)$ is the sum of the transmission times for all packets of all RT channels with an absolute deadline less than or equal to t, where t signifies the number of time units elapsed since the beginning of the HP [10] [11]:

$$h(t) = \sum_{\substack{i=1 \\ D'_i \le t}}^{Q} \left(1 + \left\lfloor \frac{t - D'_i}{P_i} \right\rfloor\right) \cdot E_i \le t \quad (8)$$

where the number of instances of evaluation can be reduced to the instances of t where a deadline occurs and that falls into the first BP in the first HP of the schedule where all periods start at time zero.

C. Schedulability Analysis with Position-Based Priorities

We now extend the real-time schedulability analysis described above with the concept of position-based priorities. Vehicle heartbeats hence adopt period and deadline based on the priority zone the vehicle currently resides in. RSU merge recommendations are always sent with the period and deadline of the highest priority zone. Road information broadcasts use the period of the lowest priority zone but employ a shorter deadline. The parameters for the different real-time traffic classes are then deadline for merge heartbeats, D_{merge}, merge recommendations, D_{rec}, and road information updates, D_{info}. Periods are denoted accordingly as, P_{merge}, P_{rec}, and P_{info}. The number of priority zones lies within the interval {1, 2,…, Z}, where Z is the number of zones chosen for the specific road site. According to equation (6), we can write the periods and adapted deadlines as:

$$D'_{merge,z} = D_{merge,z} - T_{CBP} - T_{blocking} - T_{merge}, \quad (9)$$

$$P_{merge,z} = D_{merge,z} \quad \text{where } (1 \le z \le Z), \quad (10)$$

$$D'_{rec} = D_{rec} - T_{CBP} - T_{blocking} - T_{rec} - T_{prop}, \quad (11)$$

$$P_{rec} = D_{merge,z} \quad \text{where } (D_{rec} = D_{merge,1}), \quad (12)$$

$$D'_{\text{info}} = D_{\text{info}} - T_{CBP} - T_{blocking} - T_{\text{info}} - T_{prop}, \quad (13)$$

$$P_{\text{info}} = D_{merge,z} \quad \text{where } (D_{info} = D_{merge,Z}), \quad (14)$$

The two steps of the real-time schedulability analysis introduced generally in equations 7 and 8 are adapted to the parameters of the specific scenario as follows. The utilization is calculated as:

$$U = \sum_{z=1}^{Z} \sum_{i=1}^{C_z} \frac{E_i}{P_{merge,z,i}} + \frac{E_{rec,i}}{P_{rec,i}} + \frac{E_{\text{info},i}}{P_{\text{info},i}}, \quad (15)$$

where C is the number of channels per traffic class and zone and E_{merge}, E_{rec}, and E_{info} denote the experienced transmission times of each traffic class, respectively. Similarly, $h(t)$ is expressed as:

$$\begin{aligned} h(t) = \sum_{z=1}^{Z} \sum_{\substack{i=1 \\ D'_i \le t}}^{C_z} & \left(1 + \left\lfloor \frac{t - D'_{merge,z,i}}{P_{merge,z,i}} \right\rfloor\right) \cdot E_{merge,i} + \\ & + I(D'_i \le t) \cdot \left(1 + \left\lfloor \frac{t - D'_{rec,i}}{P_{rec,i}} \right\rfloor\right) \cdot E_{rec,i} + \\ & + I(D'_i \le t) \cdot \left(1 + \left\lfloor \frac{t - D'_{\text{info},i}}{P_{\text{info},i}} \right\rfloor\right) \cdot E_{\text{info},i}, \end{aligned} \quad (16)$$

where I(A) = 1 if A = true, otherwise zero.

IV. Performance Evaluation

The proposed protocol was evaluated analytically based on a set of RT channels derived from a Matlab simulation of the merge assistance scenario defined in Section II. Table 2 gives the full list of parameter values the evaluation is based upon.

Fig. 3 shows the analytical results for the case without priority zones and thus constitutes a baseline for comparison with our novel, position-based priority scheme. In Fig 3., the percentage of the bandwidth used for the CBP can be seen as a function of the number of vehicles within transmission range of the RSU. Note that a CBP of this length implies that all safety-critical real-time data packets are accommodated without any deadline misses. Results for bit rates of 6, 12 and 24 Mbit/s are shown, approximately spanning over the achievable bit rates for 802.11p. We assume that a minimum reserved bandwidth of 20% is needed for the CBP, which is indicated by the bold line in the figure. With a CFP/CBP-ratio of 80/20, 82 vehicles can be accommodated at a bit rate of 6 Mbit/s, 160 vehicles at a bit rate of 12 Mbit/s and 292 vehicles at a bit rate of 24 Mbit/s..

TABLE 2: List of Simulation Parameters

RSU transmission radius	400 m
Superframe length	100 ms
Vehicle speed interval	100 - 150 km/h
Propagation delay	0.01 ms
Bit rate	6 - 24Mbit/s
SIFS delay	0.016 ms
Polling packet length	20 byte
Vehicle heartbeat: Packet length	500 byte
Road information update: Packet length	1.5 Kbyte
RSU recommendation: Packet length	1.5 Kbyte

In Fig. 4 and 5, we introduce position-based prioritization zones and vehicle mobility. In order to determine the amount of bandwidth required for real-time data traffic (i.e., the CFP), a schedulability test for a certain set of vehicles needs to be conducted. Fig. 4 and 5 show the results of an average of 1000 schedulability tests for a specific number of vehicles. The task set is found through Matlab simulations according to the following. A constant speed (between 100 and 150 km/h) and an initial position (within the transmission range of the RSU) are randomly determined for each vehicle. Each time instance a schedulability test is required, each vehicle's position is updated according to its speed. Periods and deadlines are set depending on the priority zone this updated position belongs to. The radius of a priority zone is defined by $r_{zone} = r_{RSU} / z$, where r_{RSU} denotes the transmission radius of the RSU and z denotes the priority (where z = 1, 2,…, Z, with Z as the total number of zones and 1 as the highest priority zone). In other words, we assume that the most hazardous zone lies in the center of the RSU's transmission range. Period and deadline of vehicle heartbeat packets ranges from 50 to 1000 ms depending on the number of priority zones. Period and deadline of RSU recommendations, as well as the deadline of road information updates correspond to the parameters of merge heartbeats in the zone with highest priority. The period of road information data, on the other hand, corresponds to the period of merge heartbeats in the lowest priority zone. The propagation delay of 0.01 ms includes delays at the sender and receiver. A new vehicle is introduced whenever another one leaves the transmission range, in order to keep the total number of vehicles at a constant value.

In Fig. 4, three priority classes, zones, are introduced. The periods are set to 50, 100 and 1000 ms, respectively. With this particular choice of periods, the message update frequency in the zone with highest priority is doubled compared to the single zone case (Fig. 3). With 80 vehicles within the RSU transmission range (which can be considered very dense traffic even on a highway) approximately 20% of the bandwidth would be left for best-effort traffic at 6 Mbit/s, 42% at 12 Mbit/s and 57% at 24 Mbit/s. Due to the fact that the deadline is set equal to the period, real-time communication within the area of hazard (the zone with highest priority) is not only taking place more frequently, but is even guaranteed to always take priority over the remaining zones and data traffic classes. Note also the slight increase in accepted vehicles as compared to Fig. 3.

The effects of introducing several priority zones can be seen in Fig. 5, with results for one, three and five priority zones. For one zone, the period is 100 ms, for three zones, the periods are 100, 300 and 1000 ms and for five zones, they are 100, 200, 300, 600 and 1000 ms. A higher number of zones offers an increased flexibility to adjust the message update frequency to the actual road and traffic conditions at the specific road site. In Fig. 5, the use of five priority zones frees 70% of the total bandwidth for best-effort services, while still guaranteeing real-time support for 80 vehicles inside the RSU's transmission range, i.e. for a highway scenario with dense traffic. The savings in resources due to lower update frequencies in low priority zones can be used in various ways. Increasing communication in hazardous areas is one option, supporting a larger amount of vehicles in the network, another. Alternatively, the threshold for the minimum amount of best-effort data traffic in the network can be increased.

V. Conclusion

We have presented a communication protocol for safety-critical V2I communication, complementing the upcoming IEEE 802.11p standard with support for real-time data traffic with position-based priorities. In our solution, guarantees for the timely delivery of safety-critical data are given. Additionally, the position-based priorities insure a decrease of the period and deadline of data packets to and from vehicles close to a zone of hazard. Communication resources are thereby focused on areas where they are most needed and the response time of proactive safety applications (and thereby ultimately even the reaction time of the driver) to critical situations is improved. Real-time schedulability analysis is used to minimize the bandwidth reserved for safety-critical data traffic while still guaranteeing timely delivery. Further, the remaining bandwidth can be used for various best effort services (infrastructure-based services or ongoing V2V applications), thereby reducing the interference on the safety-critical V2I applications to a minimum. The evaluation show that the proposed solution is suitable for the bit rates supported by the 802.11p standard and the road traffic capacity expected from a 2 – 4 lane highway. The number of priority zones, their respective period and the threshold for reserved bandwidth for best-effort services can be fine-tuned to fit the conditions of the individual road. Thereby, our solution provides the flexibility to e.g. adjust the number of supported vehicles, the reactiveness of certain ITS safety applications or the amount of possible best-effort data traffic in the network.

References

[1] Task Group p. *IEEE 802.11p: Wireless Access in Vehicular Environments (WAVE)*, draft standard, IEEE Computer Soc. 2007.

[2] K. Bilstrup, E. Uhlemann, E. G. Ström and U. Bilstrup, "Evaluation of the IEEE 802.11p MAC method for vehicle-to-vehicle communication," *Proc. IEEE Int. Symposium on Wireless Vehicular Communications*, Calgary, Canada, Sept. 2008.

[3] S. Eichler, "Performance evaluation of the IEEE 802.11p WAVE communication standard," *Proc. IEEE Vehicular Techn. Conf.*, Baltimore, MD, Sept. 2007, pp. 2199-2203.

[4] A. Böhm and M. Jonsson, "Supporting real-time data traffic in safety-critical, vehicle-to-infrastructure communication," *Proc. IEEE LCN Workshop on User Mobility and Vehicular Networks*, Montreal, Canada, Oct. 2008.

[5] T. K. Mak *et al.*, "A multi-channel VANET providing concurrent safety and commercial services," *Proc. ACM Int. Workshop on Vehicular Ad Hoc Networks*, New York, NY, Sept. 2005, pp. 1-9.

[6] C. Suthaputchakun and A. Ganz, "Priority-based inter-vehicle communication in vehicular ad-hoc networks using IEEE 802.11e," *Proc. IEEE Vehicular Techn. Conf*, Dublin, Ireland, Apr. 2007, pp. 2595-2599.

[7] S. Katragadda, *et al.*, "A decentralized location-based channel access protocol for inter-vehicle communication," *Proc. IEEE Vehicular Techn. Conf*, Seoul, Korea, Apr. 2003, pp. 1831-1835.

[8] Standards Committee, *Wireless LAN Medium Access Control (MAC) and Physical layer (PHY) specifications: Amendment 8: Medium Access Control (MAC) Quality of Service Enhancements*, IEEE Computer Society, 2005.

[9] C. L. Liu and J. W. Layland, "Scheduling algorithms for multi-programming in a hard-real-time environment," *Journal of the ACM*, vol. 20, no. 1, pp. 46-61, Jan. 1973.

[10] M. Spuri, "Analysis of deadline scheduled real-time systems," *Technical Report RR No. 2772*, Institute National de Recherche en Informatique et en Automatique (INRIA), France, 1996.

[11] J. A. Stankovic, M. Spuri, K. Ramamritham, and G. C. Buttazzo, *Deadline Scheduling for Real-Time Systems - DF and Related Algorithms*, Kluwer Academic Publishers, Boston, MA, 1998.

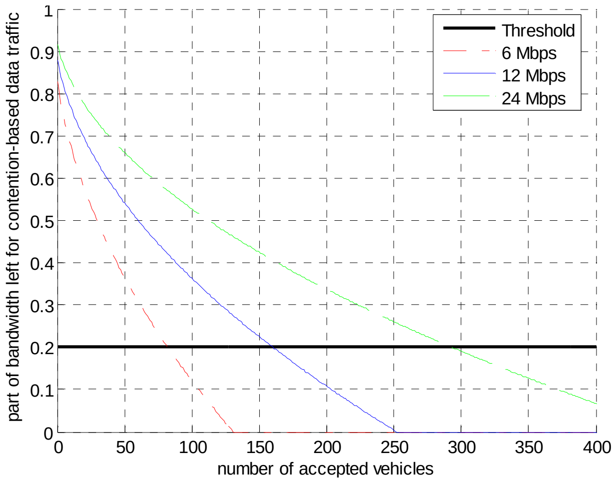

Figure 3: Fraction of bandwidth left for CBP for various numbers of accepted vehicles. No position-based priorities are used.

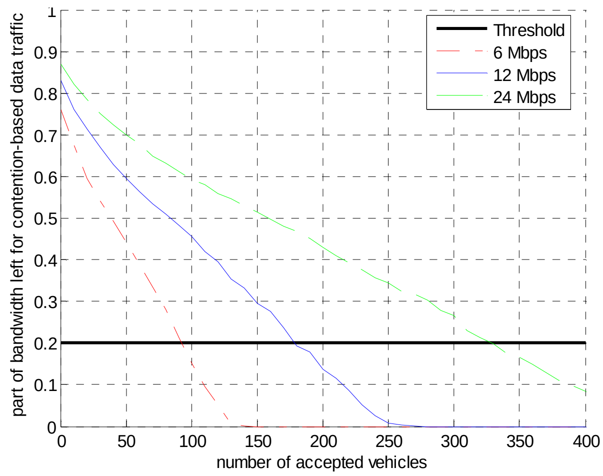

Figure 4: Fraction of bandwidth left for CBP using 3 priority zones with period 50, 100 and 1000 ms.

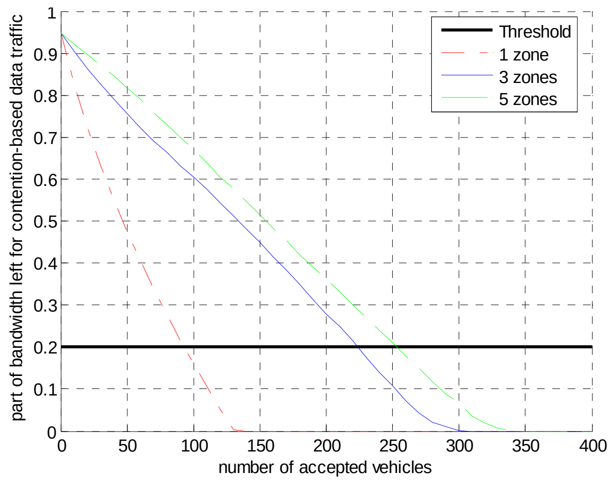

Figure 5: Fraction of bandwidth left for CBP for various numbers of priority zones and a bandwidth of 6 Mbps.