

# Design Methodology and Evaluation of Rate Adaptation Based Congestion Control for Vehicle Safety Communications

Tessa Tielert\*, Daniel Jiang<sup>†</sup>, Qi Chen<sup>†</sup>, Luca Delgrossi<sup>†</sup> and Hannes Hartenstein\*

\*Institute of Telematics and Steinbuch Centre for Computing, Karlsruhe Institute of Technology

Email: {tessa.tielert, hannes.hartenstein}@kit.edu

<sup>†</sup>Mercedes-Benz Research & Development North America, Inc.

Email: {daniel.jiang, qi.chen, luca.delgrossi}@daimler.com

**Abstract**—Vehicle Safety Communications (VSC) is advancing rapidly towards product development and field testing. While a number of possible solutions have been proposed, the question remains open as how such a system will address the issue of scalability in its actual deployment. This paper presents a design methodology for congestion control in VSC as well as the description and evaluation of a resulting rate adaption oriented protocol named PULSAR. We start with a list of design principles reflecting the state of the art that define why and how vehicles should behave while responding to channel congestion in order to ensure fairness and support the needs of safety applications. From these principles, we derive protocol building blocks required to fulfill the defined objectives. Then, the actual protocol is described and assessed in detail, including a discussion on the intricate features of channel load assessment, rate adaptation and information sharing. A comparison with other state-of-the-art protocols shows that “details matter” with respect to the temporal and spatial dimensions of the protocol outcome.

## I. INTRODUCTION

Vehicle Safety Communications (VSC) research and development is rapidly advancing to a stage where the focus shifts more and more to the actual deployment of the system. A key concern in this context is the question of system scalability. VSC based on Dedicated Short Range Communications (DSRC) will be dominated by pervasive and periodic safety message broadcasts, typically denoted as *Basic Safety Messages (BSMs)* in the U.S. and *Cooperative Awareness Messages (CAMs)* in Europe or generally as *beacons*, from all vehicles to their neighbors for cooperative awareness and tracking. Experiments have demonstrated that channel congestion can occur even in relatively simple traffic scenarios [1]. Therefore, it is important to introduce congestion control mechanisms to regulate vehicles’ BSMs in order to prevent them from actually drowning out each other.

Looking back at about 10 years of research in vehicular communications, a large number of dedicated congestion control solutions have been proposed, tailored not only towards VSC but also for traffic efficiency applications. These approaches often have differing design objectives and control dimensions, which makes it difficult to compare them and to eventually decide for one solution in the standardization process. In this work, we take one step back and provide:

1) A consolidated view on the design principles and require-

ments for congestion control for VSC. To integrate the views of previously proposed ideas, we review existing solutions with respect to these design principles and show that especially fairness aspects and safety applications’ requirements are often not covered sufficiently (Sections II and III).

2) An overall design methodology that takes into account the required transmission (Tx) range as an input from safety applications and optimizes the Tx rate with respect to a target Channel Busy Ratio (CBR) (Section IV).

3) Building blocks required to implement the design principles following the suggested methodology as well as an in-depth study of their respective configurations. We introduce a Tx rate adaptation based approach for VSC congestion control named PULSAR (Periodically Updated Load Sensitive Adaptive Rate control) as a direct implementation of the design principles and building blocks. We analyze thoroughly channel load assessment, rate adaptation and information sharing with its many intricate features and pitfalls particularly with respect to timing issues (Section V).

4) An Evaluation and comparison of PULSAR in dynamic scenarios with heterogeneous vehicular traffic densities. We show that even seemingly similar protocols show significant differences in behavior (Section VI).

We summarize our findings and conclusions in Section VII.

## II. CONGESTION CONTROL PRINCIPLES FOR VSC

In this section, we summarize the state of the art in general design principles for congestion control in VSC. Note that these principles are tailored specifically towards controlling the channel load generated by BSMs. Therefore, the objective of congestion control is to ensure that safety applications’ requirements are fulfilled rather than to maximize end-to-end throughput like for example in MANETs or sensor networks. Also, the focus is not on mitigating the broadcast storm problem resulting from retransmissions of multi-hop (emergency) messages or on regulating the channel load introduced by non-safety applications. Finally, we focus on the design of concrete algorithms rather than a general framework for the integration and regulation of different message types and priorities. We consider such efforts, e.g. [2], to be complementary to this work. In Section III, we are going to review the related work

with respect to the proposed principles.

1) *Decentralization*: DSRC-based VSC is meant to function in ad-hoc mode. Thus, it is straightforward that congestion control should be distributed. This principle is followed by all related approaches and is only stated for completeness.

2) *Participation*: In order to ensure fairness, it is intuitively clear that all nodes who *contribute* to congestion at a certain location should *participate* in congestion control. Commonly, the so-called Carrier Sense (CS) range is used as a boundary for this purpose. It is defined as the distance from a transmitting radio up to which the received signal strength can be distinguished from the noise floor by a receiving radio. Typically, it is calculated as a fixed distance assuming deterministic propagation [3][4][5][6]. An important implication of this principle is that all relevant nodes have to be *aware* of their contribution to congestion within their CS range.

3) *Local fairness*: Nodes near each other share the same channel. Therefore, it follows intuitively that nodes located physically close to each other should have similar congestion control *levels*. This principle is typically assumed implicitly in the related work, but not stated explicitly. Note that, as we will discuss later, this principle does not necessarily require nearby nodes to have the same Tx parameters.

4) *Global fairness*: [7] discusses the suitability of the proportional and max-min fairness notions for VSC. The authors conclude that, from a safety point of view, it does not make sense to increase the overall BSM throughput while potentially throttling individual nodes. These nodes who are not able to make themselves heard sufficiently may become a danger to others who are not aware of their presence. However, due to the unbounded and probabilistic nature of wireless communications, it is difficult to apply the max-min principle literally. Therefore, we use the term *global fairness* for a best-effort approach to fulfill max-min fairness.

5) *Deference to safety applications*: The objective of transmitting BSMs is to create a mutual awareness among vehicles. Thus, recent works increasingly demand to base congestion control measures on safety applications' requirements in different driving contexts. [8] suggests to adapt the Tx rate based on a vehicle's own movement, vehicle density and dangerous traffic situations. [9] describes a protocol which determines the min. required Tx rate based on the estimated tracking error of other vehicles. [10] goes one step further and derives the required Tx parameters by determining a min. warning distance based on a sample application's requirements. The overall channel load, however, is not controlled. [11] concludes that both, congestion and "awareness control", should be integrated into one solution. [2] suggests a general framework in which the application layer provides constraints within which the congestion control module can adapt the Tx parameters with respect to the channel load. We summarize these findings in the last principle: Congestion control should *defer to and work with guidance from safety applications on the space available for control adjustments*. Note that this principle implies that the adjustment space may be different for individual vehicles.

### III. RELATED WORK

Some recent approaches suggest to regulate beacons as well as event-driven messages based on their utility for the network [12][13]. They prioritize messages in the Tx queue based on factors like the distance to an event, message age, vehicle speed and the new area covered by a (re-)transmission. However, the term "beacon" in this context does not refer to BSMs but to periodic messages generated by non-safety applications, e.g. in order to disseminate road-state information. BSMs on the other hand contain information like a vehicle's current position and heading which is useful only for a very brief period of time after being generated. Therefore, it is preferable to avoid the queuing of outgoing BSMs by adapting their generation to channel conditions and safety applications' requirements.

The channel load resulting from BSMs is typically controlled by adapting Tx rate, power or both. Data rate and packet size are usually fixed following [14] and standardization guidelines. The dynamic adaptation of 802.11 parameters such as Clear Channel Assessment (CCA) threshold [2] and contention window size [15] is primarily suitable to prioritize packets and is therefore not considered here.

As detailed in [11], congestion control in VSC can be implemented as feedforward (proactive) or feedback (reactive) control. Proactive protocols estimate the future channel load and try to avoid congestion in the first place. One of the most cited solutions for VSC congestion control is D-FPAV [3] which focuses on proactive Tx power adaptation. In order to avoid congestion, vehicles running D-FPAV cooperatively calculate the max. Tx power for each individual while ensuring max-min fairness and not violating a system wide channel load threshold. The original D-FPAV design introduces a large overhead that is mostly reduced through a segment based information aggregation mechanism in [4].

For an actual deployment, it may be more practicable to react to congestion that has actually occurred, since this solution depends less on highly accurate information and suitable models for prediction. Therefore, the remainder of this paper focuses on reactive congestion control.

The protocol introduced in [9] determines the min. required Tx rate based on the estimated tracking error of other vehicles. In a second loop, Tx power is scaled to extend range until a defined CBR threshold is reached. Thereby, all information used for control is obtained locally, i.e., without information sharing. Consequently, some nodes may contribute to congestion at a location without being aware of it, since their own location is not congested.

The authors of [5] introduce a joint rate and power adaptation protocol designed to achieve fairness by assigning a fraction of resources to each node. That is, if a node increases its Tx power, it has to decrease its Tx rate and vice versa in order to occupy the same amount of resources. However, the same share of bandwidth does not necessarily imply the same amount of safety benefit for different vehicles [10]. Therefore, the underlying fairness assumption is likely too restrictive. In [16], the authors regulate the BSM Tx rate based on MAC

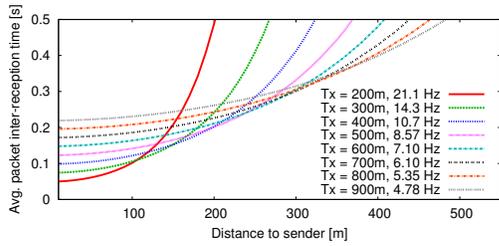


Fig. 1. Tx range and rate combinations resulting in 0.6 CBR (Rayleigh)

blocking, i.e., stopping all transmissions if the CBR exceeds a defined threshold. Based on congestion events, nodes adapt their Tx rate using Additive Increase Multiplicative Decrease (AIMD). The protocol is designed to maximize traffic generation fairness but does not take into account the awareness requirements of safety applications.

[6] and [17] are probably the most closely related approaches to the protocol presented in this work. Thus, a detailed description, analysis and evaluation as well as a comparison with our approach will be provided in Section VI.

#### IV. PROTOCOL DESIGN METHODOLOGY

Typically, the objective of VSC congestion control is to keep the channel load generated by BSMs below a certain threshold in order to reserve bandwidth for event-driven emergency messages [3][5][9]. In this context, the question is how transmission parameters, i.e., Tx rate and range, should be adapted to meet safety applications' requirements.

The underlying general optimization problem, i.e., to find the "optimal" Tx rates and ranges for a certain metric given the safety requirements of each individual vehicle as well as vehicle density and context parameters, is highly complex. Thus, a computationally tractable and distributed way of tackling the optimization problem is desirable. In the following, we outline the rationale of the proposed methodology in comparison to related work. Particularly, we focus on which metrics to select and which parameters to fix first.

In [18], the authors of [9] use the information dissemination rate (IDR), i.e., the number of packets received successfully by a node's neighbors, as a metric for application performance. They analyze different combinations of Tx rate and range and show that the max. achievable IDR is always the same, concluding that therefore IDR can be maximized through a separate control of both parameters. Furthermore, the authors analyze the relationship of IDR and CBR, i.e., the fraction of time the medium was sensed busy by the radio. They observe that any combination of Tx rate and range results in the same IDR vs. CBR curve. Thus, they conclude that CBR is a suitable feedback measure for maximizing IDR. Based on these conclusions, the authors present a design methodology for congestion control which first fixes Tx rate based on vehicle tracking performance and then adapts Tx range based on CBR. However, the authors face the problem that the optimal choice of the Tx range depends not only on the current Tx rate but also on vehicle density. Since vehicle density is hard to estimate correctly in reality, the authors "resort to a robust but suboptimal design" by adjusting the

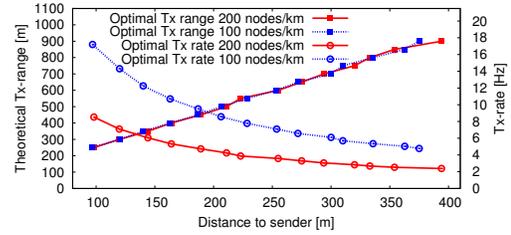


Fig. 2. Optimal Tx rate and range allocations for 0.6 CBR (Rayleigh)

channel load between a min. and max. CBR value of 0.4 and 0.8, even though they identified 0.65 as the optimal CBR value.

In this work, we take a different approach by first fixing Tx range and then adapting Tx rate with respect to CBR measurements. VSC applications based on BSMs are typically associated with a certain range up to which information is required to be received at high probability [10][19][20]. This *target range* is typically in the order of 100 to 300 meters [21]. A common metric for application performance in this context is the average *packet Inter-Reception Time (IRT)*, i.e., the average amount of time between two subsequently received messages for a sender-receiver pair [22]. Figure 1 illustrates the IRT for different pairs of Tx range<sup>1</sup> and rate which result in a CBR value of 0.6. The result was obtained from ns-2 simulations using a 20 km long circular linear road with a uniform node distribution of 100 nodes/km and Rayleigh fading (Nakagami model with  $m=1$ ). We observe that the optimal Tx rate/ range combinations in terms of minimizing the IRT for each distance from the sender form a linear line. Figure 2 plots these points as the optimal Tx rate and range with respect to the distance from the sender, i.e., the target range, for node densities of 100 and 200 nodes/km. The result indicates that, while the optimal choice of Tx rate depends on node density, the optimal Tx range does not. Therefore, our approach is to fix Tx range first depending on the currently required target range and to adapt Tx rate in a second step depending on channel conditions. Note that, since the target range changes with respect to a vehicle's driving context, the Tx range is also adapted dynamically based on guidance from the application layer. However, we expect these changes to occur at a lower pace than the adaptation of the Tx rate.

Given a fixed Tx range, the Tx rate can be used to adapt the IRT within the target range. Safety applications typically have a max. IRT value beyond which no further safety benefit can be achieved. Thereby, *safety benefit* could be expressed in different metrics, e.g., "awareness probability" [20], "application reliability" [10] or "vehicle tracking error" [18]. At the same time, there is likely to be a min. requirement below which the application cannot work. Depending on the driving context of a vehicle, we expect the safety benefit curve with respect to the Tx rate to resemble the ones shown in Figure 3. From this consideration, a min. and max. Tx rate requirement can be derived for each vehicle depending on its driving context as illustrated in the figure. If channel conditions don't allow

<sup>1</sup>The term (*theoretical*) Tx range denotes the distance at which the transition from reception to non-reception occurs using deterministic propagation.

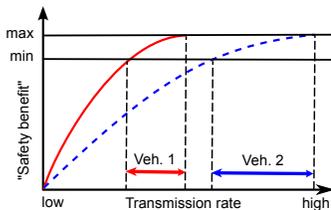


Fig. 3. Schematic illustration of safety benefit vs. Tx rate

for all vehicles to transmit at their max. rate requirement, it would be fair to provide everyone with the same relative “safety benefit”, i.e., to restrict each vehicle to a rate which corresponds to  $x\%$  of their min/max Tx rate interval.

## V. PROTOCOL DESCRIPTION

In this section, we translate the principles from Section II into actual protocol mechanisms following the methodology described in the Section IV and introduce PULSAR as a resulting reactive congestion control protocol for VSC. Note that, even though PULSAR is designed to support different Tx ranges and min/max Tx rate interval settings per node, in this work we focus on uniform configurations of all vehicles in order to better illustrate the characteristics of the different mechanisms. An extended description of the protocol is going to be presented in future work.

### A. Channel Load Assessment

While some protocols use metrics like the number of collisions and SNR as a metric for channel load [13], PULSAR, like most other approaches [5][6][9][16], makes use of CBR measurements. In [18], the authors conclude that CBR is a suitable feedback metric for maximizing the number of received packets (IDR) and recommend a *CBR target value* of 0.65. Our own analysis has shown that, independent of packet size and the fading model used, a CBR value between 0.6 and 0.7 achieves the best results in terms of IDR. In this work, we use a CBR target value of 0.6. Note that any CBR value is only meaningful considering the underlying CCA threshold. While some approaches use values near -85 dBm [5][23], our analysis has shown that a lower CCA-threshold results in better per-message reception performance. Today’s radios can set the CCA threshold very close to the noise floor. Assuming a noise floor of -99 dBm (-107 dBm thermal noise at 5.9 GHz and 5 dB hardware noise) and a tolerance of 4 dB, we use a CCA threshold of -95 dBm which is consistent with [18].

In PULSAR, the arrival of a new CBR measurement triggers a Tx rate adaptation. A key parameter in this context is the length of the interval during which the busy indications of the physical layer are evaluated. [17] uses a node’s Tx interval for this purpose. However, this approach can lead to (global) unfairness, for example since nodes using a lower Tx rate get less chances to increase their share. PULSAR therefore uses a fixed length *Channel Monitoring and Decision Interval (CMDI)* for all nodes. Since the CMDI influences the convergence time of the protocol, its length should be minimized. However, due to fading, CBR measurements are subject to a probability distribution whose variance is inversely proportional to the

length of the CMDI. Our analysis has shown that a reasonable trade-off between variance and reaction time is between 200 and 400 ms. In this work, we use a CMDI of 250 ms.

In order to reduce the probability of reacting to a “false alarm”, PULSAR additionally uses a first-order digital low-pass filter to reduce measurement noise. The smoothed out value  $CBR_{avg}$  is calculated from a new measurement  $CBR_{new}$  as  $CBR_{avg} = 0.5 * CBR_{avg} + 0.5 * CBR_{new}$ .

### B. Rate Adaptation

When a new CBR measurement arrives at the end of each CMDI, PULSAR compares the measured value against the target value. In a binary decision, a node’s Tx rate is either decreased if there is excess channel load or increased otherwise. Thereby the Tx rate is adapted within the min. and max. Tx rate limits provided by the application layer. The actual rate increment or decrement is calculated using Additive Increase Multiplicative Decrease (AIMD) and the target rate mechanism described below. The combination of both mechanisms ensures local and global fairness as well as sufficient convergence time when vehicles are moving fast.

In [24], the authors show that, for wired networks, AIMD is favorable over the other three combinations of additive and multiplicative increase and decrease since it ensures max-min fairness. An analysis by the authors of [25] has shown that this also holds for VSC. However, in some cases, AIMD may not converge to fairness if nodes are not synchronized in their measurements [26]. PULSAR therefore assumes that all nodes are (loosely) synchronized in their CMDI, e.g., using GPS. We are currently evaluating LIMERIC [25], which does not have this restriction, for usage in PULSAR.

In terms of controlling a system, it is desirable that control measures take effect immediately. However, when adapting the Tx rate at the end of each CMDI, nodes typically have already scheduled their next transmission based on the old Tx rate. In addition, a node’s current Tx interval may be longer than the CMDI. PULSAR therefore uses a *rescheduling mechanism* for already scheduled transmissions. Let  $i_{orig}$ ,  $i_{new}$  and  $i_{rest}$  denote the original and newly calculated Tx interval and the remaining part of  $i_{orig}$ , respectively. The pending transmission is canceled and rescheduled using  $time_{rescheduled} = time_{now} + \frac{i_{rest}}{i_{orig}} * i_{new}$ .

Vehicular networks are characterized by a high dynamic in node topology. In order to ensure especially local fairness, the rate adaptation mechanism is required to have a sufficiently short *convergence time*. Given a fixed-length CMDI, one way would be to tune AIMD parameters aggressively, i.e. to take into account large fluctuations of Tx rate and CBR. Another way is to make use of the average Tx rate of a node’s neighbors. PULSAR uses a node’s *target rate*, i.e., an average of the Tx rates of its neighbors, as a *gravitation pull* when calculating Tx rate adjustments. That is, when increasing, the increment is doubled if the vehicle’s current rate is below target rate. Else, if the current rate is already above target rate, the increment is halved. In opposite direction, the same principle applies to the decrement. The target rate itself is

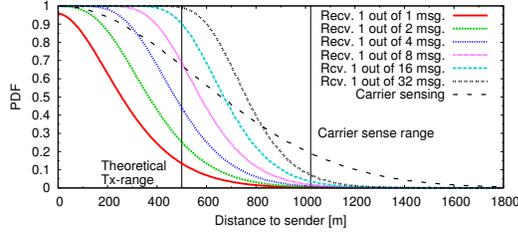


Fig. 4. Piggybacking efficiency at 0.6 CBR, Tx = 500 m

calculated as an exponentially weighted moving average of the Tx rate information  $r_c$  contained in received BSMs (e.g., using  $\delta=0.1$ ):  $r_{target} = (1 - \delta) * r_{target} + \delta * r_c$ . Compared to an average derived, e.g., from a neighbor table, fading results in an implicit weighting within the target rate, favoring close neighbors over distanced ones and large groups of neighbors over small groups. In our analysis, the target rate mechanism has shown a smoother convergence behavior than a simple neighbor average derived from a neighbor table. This mechanism is going to be further evaluated and improved in future work.

### C. Information Sharing

In the related work, there is no consensus on whether or not information on channel conditions needs to be exchanged among vehicles. While some approaches share channel state information [16][17], others rely on local measurements only [5][6][9][13]. The latter approach features less overhead. However, it may lead to a violation of the local and global fairness principle. Due to fading, CBR measurements are subject to a certain randomness. Therefore, even neighboring nodes can come to different assessments of the channel state. For example, consider a situation in which node A detects congestion, while its neighbor B does not. In this case, A would increase its rate, while B would decrease, making room for A to increase again and so on. Over time, this effect can result in imbalance and unfairness. In addition, nodes can contribute to congestion at a distant location without measuring congestion themselves. Thus, while from a system point of view the objective of controlling channel load can be met with local measurements only, the participation and fairness principles require the exchange of channel state information.

The participation rule suggests that all nodes within a certain range, e.g., CS range, from a congested location need to be informed of their contribution in order to participate in congestion control. A CCA threshold of -95 dBm results in a CS range which has approximately twice the length of the (theoretical) Tx range. Therefore, for uniform Tx range settings, it is not possible to notify all nodes within CS range of an occurring congestion by using 1-hop piggybacking. One way would be to send a high-power message containing the required information, another one is to piggyback information over two hops. Unless an increment of the Tx power is required anyway, the latter approach is likely to introduce less additional load to the channel.

Figure 4 illustrates the feasibility of constructing a 2-hop information dissemination mechanism for congestion control

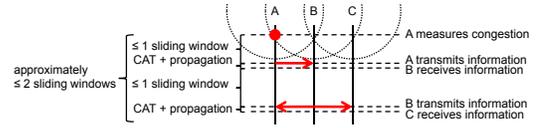


Fig. 5. Max. age of 2-hop piggybacking information

over a Rayleigh channel. The vertical bars in the figure illustrate the underlying Tx range (500m) and CS range (1021 m). The figure also shows that, with Rayleigh fading, there is a non-trivial chance that a generated message is carrier sensible beyond the nominal Tx range and even CS range. For fairness, it is therefore necessary to propagate congestion information at least up to CS range. While the probability of reception is relatively low for a single transmission under channel load (the red/leftmost curve in the figure), generally more than one vehicle measures and reports congestion. It is sufficient to receive just one out of the messages transmitted by all of these vehicles in order to receive the contained congestion information. Therefore we can assume that, with high probability, the nominal Tx and CS range can be covered in one and two *piggybacking hops*, respectively.

When constructing a 2-hop piggybacking scheme as outlined above, the most straightforward approach is to react to the congestion indication received from local measurement, 1-hop and 2-hop information within the last CMDI. However, in doing so, we would not take into account the resulting *information dissemination delay*: Nodes measuring congestion themselves would decrease their Tx rates first, while others contributing to congestion might not yet be aware of doing so, increasing their Tx rates even further. To avoid this violation of the global fairness principle, the propagation delay has to be taken into account when reacting to local and reported congestion. The general objective is for all nodes to react to the same state of the system at the same time.

In PULSAR, each node attaches its own congestion state based on its last CBR measurement as well as the max. congestion state it has received from nodes within range R and time frame T to each beacon. Assuming non-deterministic propagation, R is an adjustable parameter which determines the size of one piggybacking hop and therefore also the *participation range*, i.e., the distance from a congested location at which nodes participate in congestion control. T is a sliding window of the same length as the CMDI in order to prevent the relayed information from being older than one CMDI.

Figure 5 sketches a scenario of 3 nodes (A, B and C) in a unit disc graph model where B and C are not congested but contributing to the congestion at A. However, A cannot reach C directly. The figure illustrates that the max. age of the 2-hop information received at C is approximately two sliding windows, i.e., two CMDIs. In order to have all nodes react to the same system state, PULSAR delays a node's reaction by a max. of two CMDIs. To be precise, at the end of CMDI t, C reacts to the max. congestion state of its local measurement of t-2, its 1-hop information received between t-1 and t-2 and its 2-hop information received between t and t-2, but generated at t-2. In order to determine in which CMDI the reported 2-

TABLE I  
SIMULATION PARAMETERS

Parameter	Value
Noise floor	-99 dBm
CCA threshold	-95 dBm
Data rate	6 mbps
RF fading model	Nakagami, $m = 1$
Preamble and frame header capture SINR	5 dB
Frame body reception capture SINR	8 dB
Safety message size (with security overhead)	250 Bytes
Min./ max. Tx rate	1 Hz/ 10 Hz
(Theoretical) Tx range	500 m (10.21 dBm)
Additive increase and multiplicative decrease	0.1 Hz, 0.03
Min./ max. CBR threshold	0.5/ 0.6
Length of CBR channel monitoring period	250 ms
Size of one piggybacking hop	750 m ( $1.5 * \text{Tx-range}$ )
$\beta$ [6][17]	0.02 s

hop information has been generated, each packet additionally includes a time stamp of the last reception of 1-hop congestion information.

At first glance, it may seem problematic that the congestion information originated by A is relayed from B to C, but also from B back to A. However, the reaction delay introduced above takes care of this issue. Since, at time  $t$ , A reacts to the max. congestion state of its own measurement at  $t-2$  and the received information originated at  $t-2$ , the bounce back has no negative effect.

## VI. EVALUATION

So far we have already discussed many subtle design decisions in VSC congestion control. In this section, we demonstrate the necessity of two key mechanisms in PULSAR, i.e., target rate and 2-hop piggybacking. For this purpose, we evaluate PULSAR in two contrasting scenarios: A static highway crossing and a dynamic highway scenario consisting of two groups of vehicles passing each other. We show that PULSAR satisfies the design principles defined in Section II while comparing its performance against two algorithms which at first glance resemble it closely [6][17]. However, we are going to show that some of the design decisions in these two protocols can lead to undesired behavior in terms of fairness and convergence. Note that, in order to facilitate a fair comparison with the related work which does not support non-uniform settings, we use uniform Tx rate intervals and Tx ranges for all nodes. An evaluation of PULSAR with non-uniform settings is going to be presented in future work.

1) *Description, configuration and analysis of [6]:* [6] introduces a scheme for Tx power or rate adaptation. Since our work focuses on rate control, in the following we use the latter algorithm for a comparison. After each transmission, a node compares its current CBR measurement (evaluated over the last Tx interval) against a threshold. If the measured value is lower than the threshold, it increases its Tx interval by a fixed value  $\beta$  and vice versa. For fairness, a node is prevented from decreasing its Tx interval if it is already below the average value of its neighbors within Tx range.

One key difference between the design of [6] and PULSAR is that, instead of adapting Tx rate using AIMD, [6] adapts the Tx interval by a fixed value of  $\beta$  which is configured as 0.1 s. Note

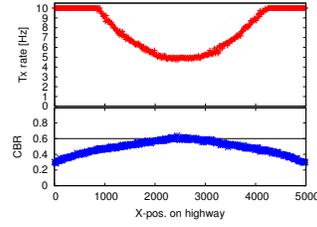


Fig. 6. PULSAR w.r.t. position

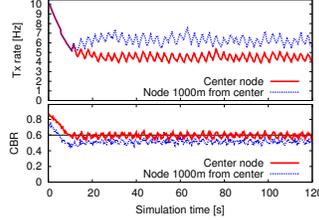


Fig. 8. PULSAR w.r.t. time

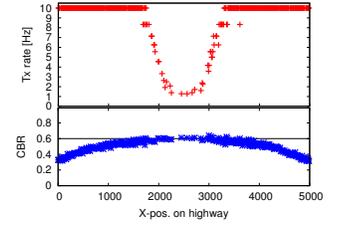


Fig. 7. [6] w.r.t. position

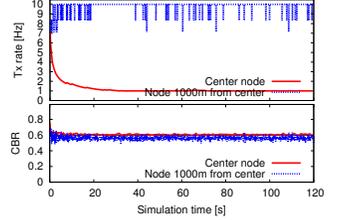


Fig. 9. [6] w.r.t. time

that this adaptation is equivalent to a non-linear (exponential) adaptation of the Tx rate. For example, if a node transmits at 9 Hz, i.e., with a Tx interval of 0.111 s, a decrement of 0.1s results in a Tx rate of 90 Hz. On the other hand, a node transmitting at 1 Hz would increase its rate to 1.1 Hz in the same configuration. It is important to note that [6] has been designed for non-safety applications assuming a max. Tx rate of 1 Hz. While the fundamental problem is still there, in this configuration, the consequences of the non-linear adaptation are less severe. In order to provide a fairer comparison, we set  $\beta$  to 0.02 s in this work. Further, we use a min. and max. Tx rate of 1 and 10 Hz, respectively. Note that the lower value of  $\beta$  comes at the price of a very long convergence time when starting the simulation at the min. Tx rate. In the crossing scenario described below, the algorithm of [6] needs more than 60 s to converge to equilibrium.

2) *Comparison of PULSAR and [6] in crossed roads scenario:* Figures 6 and 7 illustrate a snapshot each of the converged rate allocation and resulting channel load observed on the longer one of two roads crossing each other for PULSAR and [6], respectively. The setup is intended as one highway (4 km in length) crossing another one (5 km) over a bridge. Therefore, vehicles are not in conflict in terms of right of way at the intersection. The result has been derived from ns-2 simulations (version 2.34) using the parameters given in Table I. The scenario was set up to facilitate a Tx rate of 5 Hz at the center of the crossing (x-position 2.5 km) without violating the CBR threshold of 0.6 if all nodes use the same Tx rate throughout the scenario. In terms of global fairness, the objective in this scenario is therefore to keep the Tx rate of the nodes at the intersection near 5 Hz. We observe in Figure 7 that [6] does not meet this objective, since the nodes at the most congested location get pushed down to the min. rate of 1 Hz. The reason for this behavior is that in [6], nodes do not exchange information on channel conditions. Thus, not all nodes within CS range of the occurrence of congestion are aware of contributing. Note that, while not shown for space constraints, running PULSAR with 1-hop piggybacking

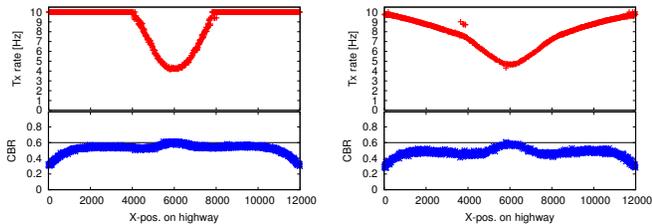


Fig. 10. PULSAR w.r.t. position

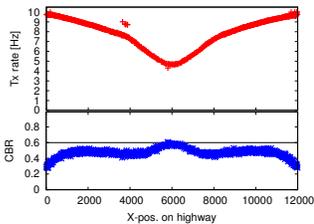


Fig. 11. SOURC [17] w.r.t. position

instead of the 2-hop scheme results in a similar rate allocation as [6]. Therefore, 1-hop piggybacking is not sufficient for achieving global fairness. The result of PULSAR with 2-hop piggybacking is shown in Figure 6. We can see that nodes further away from the intersection participate in congestion control, allowing those at the most congested location to transmit at about 4.5 Hz which is reasonably close to the optimal value. Figures 8 and 9 show the rate allocation and resulting CBR with respect to simulation time for PULSAR and [6], respectively. We observe that PULSAR’s Tx rate is moving up and down as if it were pulsing (due to AIMD), while the allocation of [6] stays more constant, interrupted by outlier values which result from the exponential rate adaptation.

3) *Description of SOURC [17]*: SOURC is an extension to [6] which features 2-hop piggybacking as well as a min. and max. CBR threshold. Each node attaches its local CBR measurement as well as the max. CBR value received from its neighbors within Tx range (“local CBR”) to each packet. A node then calculates the “global CBR” as the max. local CBR received during the last Tx interval. Like in [6], after each transmission, a node increases its Tx interval by  $\beta$  if its local CBR measurement is above the max. threshold. If the latter is not the case, a node checks if its global CBR is below the *min.* threshold. If this is the case and additionally its Tx interval is above neighbor average, it decreases its Tx interval by  $\beta$ . Otherwise, i.e., if its global CBR is above the min. threshold or its Tx interval is above neighbor average, it sets its Tx interval to the neighbor average. In a nutshell, only nodes measuring congestion themselves decrease their Tx rate actively and only those with the lowest Tx rate may actively go up. Everyone else’s Tx rate is adjusted implicitly by adapting to the neighbor average.

4) *Comparison of PULSAR and SOURC in crossed roads scenario*: Figures 10 and 11 illustrate the converged rate allocations and CBR measurements for PULSAR and SOURC, respectively, in a static highway crossing scenario with the same setup as above except for a longer road in x-direction (12 km). We can see that both protocols meet the objective of keeping the Tx rates at the intersection (x-pos. 6000 m) near the optimal value of 5 Hz. However, we observe that, while PULSAR limits participation in congestion control to approximately CS-range, SOURC reduces the Tx rates of nodes located up to 6 km away from the congested location. This behavior results from the implicit Tx rate adjustments in SOURC. By having nodes adapt to neighbor average, participation in congestion control literally propagates forever. Since nodes at such large distances as 6 km from the congested

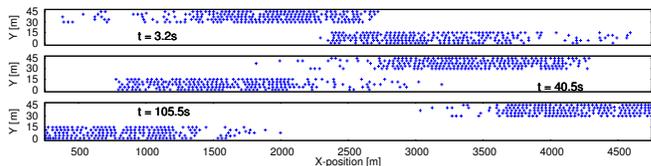


Fig. 12. Node positions in dynamic highway scenario

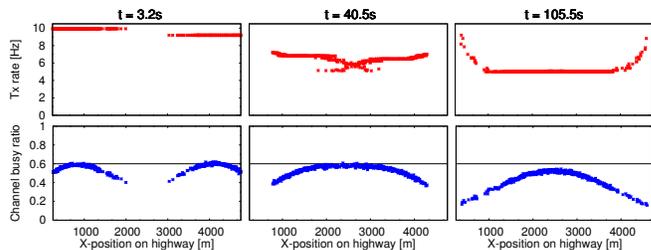


Fig. 13. PULSAR without target rate in dynamic highway scenario

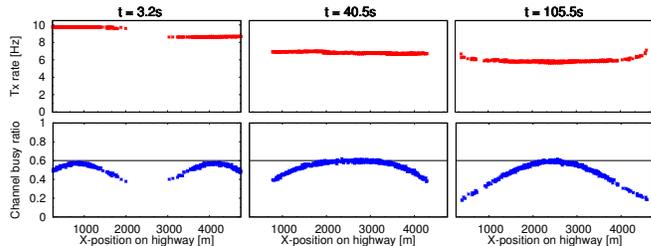


Fig. 14. PULSAR with target rate in dynamic highway scenario

location have a negligible contribution to congestion, if any, the rate allocation of SOURC is not (globally) fair. Note that the implicit rate adaptation in SOURC is another way of dealing with the information dissemination delay in 2-hop piggybacking. However, as we have just seen, it comes at the cost of unbounded participation in congestion control.

5) *Evaluation of PULSAR and SOURC in dynamic highway scenario*: Figures 13 and 14 illustrate PULSAR’s Tx rate allocation and resulting channel load without and with target rate, respectively, for the 3 snapshots of a dynamic highway scenario shown in Figure 12. In this scenario, two groups of vehicles pass each other in opposite directions. At simulation time  $t = 40.5$  s, we observe that, without target rate, PULSAR’s rate allocation resembles the letter X near x-position 2500 m. Therefore, the local fairness principle is not fulfilled. In other words, the protocol does not converge fast enough to adapt to the changes in vehicle density. With target rate, on the other hand, we can see that PULSAR’s rate allocation is a smooth curve, i.e., the local fairness principle is fulfilled. Note that, since SOURC also makes use of the average Tx rate when taking adaptation decisions, its convergence behavior in this scenario is very similar to the one of PULSAR with target rate. Therefore, the result is not shown here. However, throughout the scenario, SOURC makes slightly less use of the channel due to its two CBR thresholds.

## VII. SUMMARY AND CONCLUSIONS

VSC is going to be dominated by periodic 1-hop messages from all vehicles to their neighbors in order to create a mutual awareness. Without regulation, these beacons can easily lead to channel congestion. In this work, we have reviewed and

consolidated the state of the art into general design principles which aim at controlling channel load while satisfying safety applications' awareness requirements. We have suggested a design methodology for reactive congestion control for VSC which adapts a vehicle's Tx rate based on guidance by the application layer on Tx range as well as min. and max. Tx rate. Following the design principles and the suggested methodology, we have derived required protocol building blocks and introduced a resulting protocol named PULSAR. In the evaluation, we have demonstrated that PULSAR fulfills the design principles: It is a distributed algorithm which achieves local fairness by taking into account a node's neighbors' Tx rates. By means of 2-hop piggybacking, all nodes within CS range of a congested location participate in congestion control. This also leads to the fulfillment of the global fairness principle which requires to maximize the Tx rate of the nodes which are throttled the most by congestion control while not unnecessarily constraining other vehicles which do not contribute to congestion. Finally, PULSAR is designed to accommodate different min. and max. Tx rate intervals and Tx ranges. Therefore, it is able to operate within the space available for adjustments defined by safety applications. For the eventual standardization of VSC congestion control, it is important to understand in detail the implications of different protocol design decisions and mechanisms. Therefore, this work puts emphasis on a step-by-step design and a thorough discussion of many subtle implications of design decisions and mechanisms. We have not only evaluated PULSAR's behavior in contrasting scenarios designed to stress the underlying protocol mechanisms but have also compared it against two other state of the art approaches, studying in detail their behavior and discussing their advantages and disadvantages. An important lesson learned by the comparison is the insight that protocols that seem to differ only in some "details" may show significant differences in their temporal and spatial behavior. Thus, it is also an interesting issue for future research to study the robustness of congestion control in cases where some vehicles will intentionally or unintentionally differ from the congestion control policy. Further, while in this work, the description and evaluation of PULSAR has been restricted to uniform settings of Tx range and min./ max. Tx rate for reasons of space constraints, presentation clarity and fairness of comparison against the related work, we are going to present an extended description and evaluation of PULSAR with non-uniform settings in future work.

#### ACKNOWLEDGMENT

This work was partially funded by the German Research Foundation (DFG) Research Training Group GRK 1194 "Self-organizing Sensor-Actuator Networks".

#### REFERENCES

- [1] A. Weinfeld, "Methods to Reduce DSRC Channel Congestion and Improve V2V Communication Reliability," in *Proc. of the ITS World Congress*, 2010.
- [2] R. K. Schmidt, A. Brakemeier, T. Leinmüller, B. Böddeker, and G. Schäfer, "Architecture for Decentralized Mitigation of Local Congestion in VANETs," in *Proc. of the 10th International Conference on ITS Telecommunications (ITST)*, 2010.

- [3] M. Torrent-Moreno, J. Mittag, P. Santi, and H. Hartenstein, "Vehicle-to-Vehicle Communication: Fair Transmit Power Control for Safety-Critical Information," *IEEE Trans. on Vehicular Technology*, vol. 58, no. 7, pp. 3684–3703, 2009.
- [4] J. Mittag, F. Schmidt-Eisenlohr, M. Killat, J. Härri, and H. Hartenstein, "Analysis and Design of Effective and Low-Overhead Transmission Power Control for VANETs," in *Proc. of the 5th ACM international workshop on VehiculAr Inter-NETworking (VANET)*, 2008.
- [5] R. Baldessari, D. Scanferla, L. Le, W. Zhang, and A. Festag, "Joining Forces for VANETs: A Combined Transmit Power and Rate Control Algorithm," in *Proc. of the 7th International Workshop on Intelligent Transportation (WIT)*, 2010.
- [6] C. Khorakhun, H. Busche, and H. Rohling, "Congestion Control for VANETs based on Power or Rate Adaptation," in *Proc. of the 5th International Workshop on Intelligent Transportation (WIT)*, 2008.
- [7] M. Torrent-Moreno, P. Santi, and H. Hartenstein, "Fair sharing of bandwidth in VANETs," *Proc. of the 2nd ACM International Workshop on Vehicular Ad-hoc Networks (VANET)*, 2005.
- [8] R. K. Schmidt, T. Leinmüller, E. Schoch, F. Kargl, and G. Schäfer, "Exploration of Adaptive Beaconing for Efficient Intervehicle Safety Communication," *IEEE Network*, vol. 24, pp. 14–19, Jan. 2010.
- [9] C. Huang, Y. Fallah, R. Sengupta, and H. Krishnan, "Adaptive Intervehicle Communication Control for Cooperative Safety Systems," *IEEE Network*, vol. 24, no. 1, pp. 6–13, 2010.
- [10] M. Sepulcre, J. Gozalvez, J. Härri, and H. Hartenstein, "Contextual Communications Congestion Control for Cooperative Vehicular Networks," *IEEE Trans. on Wireless Commun.*, vol. 10, no. 2, pp. 385–389, 2011.
- [11] M. Sepulcre, J. Mittag, P. Santi, H. Hartenstein, and J. Gozalvez, "Congestion and awareness control in cooperative vehicular systems," *Proc. of the IEEE*, vol. 99, no. 7, pp. 1260–1279, Jul. 2011.
- [12] M. Bouassida and M. Shawky, "A Cooperative Congestion Control Approach within VANETs: Formal Verification and Performance Evaluation," *EURASIP Journal on Wireless Comm. and Networking*, vol. 2010, pp. 11:1–11:12, Apr. 2010.
- [13] C. Sommer, O. K. Tonguz, and F. Dressler, "Traffic information systems: Efficient message dissemination via adaptive beaconing," *IEEE Communications Magazine*, vol. 49, no. 5, pp. 173–179, May 2011.
- [14] D. Jiang, Q. Chen, and L. Delgrossi, "Optimal Data Rate Selection for Vehicle Safety Communications," in *Proc. of the 5th ACM international workshop on VehiculAr Inter-NETworking (VANET)*, 2008.
- [15] Y. Zang, L. Stibor, X. Cheng, H. Reumertman, A. Paruzel, and A. Barroso, "Congestion Control in Wireless Networks for Vehicular Safety Applications," in *Proc. of the 8th European Wireless Conference*, 2007.
- [16] J. He, H. Chen, T. Chen, and W. Cheng, "Adaptive Congestion Control for DSRC Vehicle Networks," *IEEE Communications Letters*, vol. 14, no. 2, pp. 127–129, 2010.
- [17] H. Busche, C. Khorakhun, and H. Rohling, "Self-Organized Update Rate Control for Inter-Vehicle Networks," in *Proc. of the 10th International Workshop on Intelligent Transportation (WIT)*, 2010.
- [18] Y. P. Fallah, C. Huang, R. Sengupta, and H. Krishnan, "Analysis of Information Dissemination in Vehicular Ad-Hoc Networks With Application to Cooperative Vehicle Safety Systems," *IEEE Trans. on Vehicular Technology*, vol. 60, no. 1, pp. 233–247, Jan. 2011.
- [19] X. Guan, R. Sengupta, H. Krishnan, and F. Bai, "A Feedback-Based Power Control Algorithm Design for VANET," in *Mobile Networking for Vehicular Environments*, 2007.
- [20] N. An, T. Gaugel, and H. Hartenstein, "VANET: Is 95% of Packet Reception Safe?" in *Proc. of the 11th International Conference on ITS Telecommunications (ITST)*, 2011.
- [21] "Vehicle Safety Communications Project Task 3 Final Report: Identify Intelligent Vehicle Safety Applications Enabled by DSRC," DOT HS 809 859, 2005.
- [22] T. ElBatt, S. K. Goel, G. Holland, H. Krishnan, and J. Parikh, "Cooperative Collision Warning using Dedicated Short Range Wireless Communications," in *Proc. of the 3rd international workshop on Vehicular Ad hoc Networks (VANET)*, 2006.
- [23] R. K. Schmidt, T. Leinmüller, B. Böddeker, and G. Schäfer, "Adapting the wireless carrier sensing for VANETs," in *Proc. of the 7th International Workshop on Intelligent Transportation (WIT)*, 2010.
- [24] D. Chiu and R. Jain, "Analysis of the Increase and Decrease Algorithms for Congestion Avoidance in Computer Networks," *Computer Networks and ISDN Systems*, vol. 17, no. 1, pp. 1–14, Jun. 1989.
- [25] J. Kenney, G. Bansal, and C. Rohrs, "LIMERIC: A Linear Message Rate Control Algorithm for Vehicular DSRC Systems," in *Proc. of the 8th ACM int. workshop on VehiculAr Inter-NETworking (VANET)*, 2011.
- [26] S. Gorinsky and H. Vin, "Additive Increase Appears Inferior," 2000, TR 2000-18, Department of CS, University of Texas at Austin.