

Avoiding the Pitfalls of Distributed Congestion Control in the IEEE WAVE Safety Messaging

Journal:	*IEEE Transactions on Vehicular Technology*
Manuscript ID:	VT-2012-01872
Suggested Category:	Regular Paper
Date Submitted by the Author:	30-Nov-2012
Complete List of Authors:	Kim, Byungjo; Korea University, College of Information and Communication Kim, Hyogon; Korea University, College of Information and Communication
Keywords:	WAVE, V2V, Basic Safety Message (BSM), congestion control, anomaly

SCHOLARONE™
Manuscripts
For Review

Avoiding the Pitfalls of Distributed Congestion Control in the IEEE WAVE Safety Messaging

Byungjo Kim *Student Member, IEEE,* and Hyogon Kim, *Member, IEEE*

Abstract

In the IEEE Wireless Access in Vehicular Environment (WAVE), the periodic broadcast of the Basic Safety Message (BSM) enables neighbor awareness in the vehicle-to-vehicle (V2V) context. In order to maximize the level of awareness and consequently the driving safety, the BSM transmission at the highest allowed rate is desired in principle. A caveat, however, is controlling the BSM traffic within the given channel capacity, because otherwise it can actually lower the delivery probability due to message collisions. In order to avoid such congestion situation, the traditional mode of control is regulating the frequency or the power of the BSM transmission based on the channel load. In this paper, we shed light on the pitfalls that lurk in exercising adaptive control based on an observed global state such as the channel load. Specifically, we show that straightforward threshold-based or hysteresis-based controls can irrevocably deviate the rate assignments from the given vehicle density pattern. As a solution, we show that distributed but coordinated control provably leads to stability and relevance to the given vehicle density pattern.

Index Terms

WAVE, V2V communication, Basic Safety Message (BSM), congestion control, anomaly

For Peer Review

I. Introduction

IN vehicle-to-vehicle (V2V) context in the IEEE Wireless Access in Vehicular Environment (WAVE) [1], [2], [3], periodic broadcast of the Basic Safety Message (BSM) [4] enables neighbor awareness. The BSM from each vehicle reports their position, speed, heading, acceleration, steering wheel angle, brake status and vehicle size among others. In order to maximize the level of awareness and consequently the driving safety, the BSM transmission at the highest allowed rate is desired in principle. Although the WAVE standard suite does not stipulate the rate, most safety applications, *e.g.* as are defined in the VSC-A implementation, require 10Hz [5], [6]. A caveat in periodic BSM transmission, however, is not putting more BSM traffic than the wireless channel can digest. Otherwise,

The authors are with Korea University. E-mail: hyogon@korea.ac.kr. This work was supported by Mid-career Researcher Program through NRF grant funded by the MEST (2011-0028892).

consequent message collisions can actually reduce the number of successfully delivered BSMs than when the rate is lower. In order to avoid the congestion situation, the traditionally considered mode of control has been reducing the frequency [7], [8], [9], [10] or the power [11], [12] of the BSM transmission based on the channel load such as channel utilization or the number of BSM messages heard in the given monitoring interval.

In this paper, we shed light on the pitfalls that lurk in exercising the aforementioned adaptive control based on an observed global state. Specifically, we argue that the rate or power assignment made in each vehicle through a straightforward congestion control can easily become irrelevant to the given vehicle density, which in turn can adversely affect the neighbor awareness and eventually traffic safety. We will show that the pathology stems from positing that neighboring vehicles would see a single, consistent picture of the global congestion state, so their individual reactions to the state will be identical and lead to convergence. But the assumption is only roughly true, and the fine difference in the perceived global state by neighboring vehicles can occasionally trigger opposite adjustments on them. Once such apparently minor event takes place, however, it can start an irrevocable process of polarization even between close neighbors. The reason is that the global state is the sum of the power or rate adjustments made by possibly many vehicles, and the contribution from a single vehicle cannot quickly change it. Moreover, the opposite remediary actions by the neighboring vehicles based on the perceived congestion state can fall in the vicious cycle that amplifies the difference in the rate or power without changing the global state much. Fig. 1 presents a simplified schematic of such a cycle to make the point. The whole point here is that we only need a minor difference in the observed global congestion state to trigger such vicious cycle. For instance, the number of successfully BSMs or channel utilization can be easily affected by the MAC collisions and collision resolution dynamics or by PHY channel fluctuations.

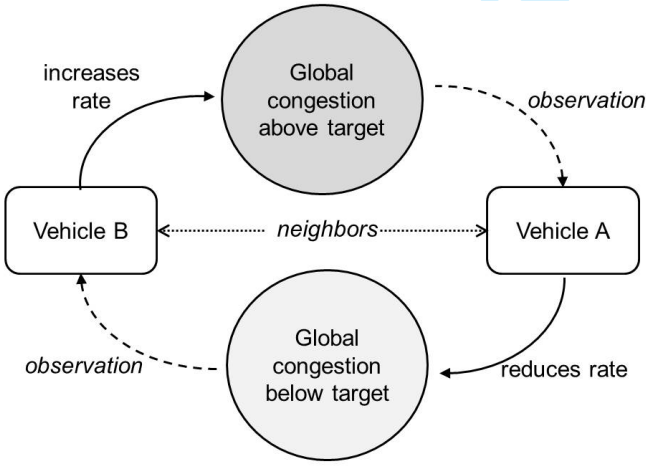


Fig. 1. Undesirable feedback cycle between neighboring vehicles A and B

Since controlling the perceived congestion states between neighbors is impractical, we instead propose to add explicit coordination in the distributed control as a solution. We will demonstrate that by piggybacking the rate or

power information on the BSM, neighbors can coordinate their rate adjustments thus avoiding the anomaly.

The paper is organized as follows. In Section II, we first provide the background on the standards and the technical aspects of V2V safety messaging. Then we briefly discuss the existing congestion control proposals. Note that this paper does not put forth a competing proposal. Rather, we shed light on the less understood aspect of the distributed congestion control and propose a remedy. In Section III, we experiment with two most popular types of adaptive control, *i.e.*, threshold-based and hysteresis-based, to show the anomaly exists regardless of the mechanisms. In addition, we investigate the ETSI-standardized distributed congestion control, as it uses the hysteresis-based control. In order to prevent non-essential implementation details from diffusing the focus, we set up a generic additive-increase multiplicative-decrease (AIMD) adaptive control for the two types of tested control modes. Then we show that the anomaly indeed takes place irrespective of specific control algorithms. Finally, we present a solution called the mean-check control that can augment existing distributed congestion control schemes at a low bandwidth and complexity cost to avoid the pathology. We discuss why the solution is structurally robust. Section IV concludes the paper.

II. Background

A. Safety communication in WAVE

The WAVE standard allocates seven 10MHz channels in 5.9GHz band for communication between WAVE devices such as on-board unit (OBU) and road-side unit (RSU). The frequency band is often called the Dedicated Short Range Communication (DSRC) band. Unlike in wireless LAN (WLAN), WAVE communication can use two channels [3]. The central channel among the seven called the Control Channel (CCH) is exclusively used by vehicular safety messages (*e.g.* BSM) and service announcements, where the services are provided in the rest six channels that are called Service Channels (SCHs). The WAVE standard allows four different modes of the channel usage based on the OBU capacity, but recent trend is where the OBU is equipped with dual-radio and always stay on the CCH using one of them ("continuous mode") [3], while the other radio is used for any WAVE service subscriptions.

B. Related work on congestion control

The congestion control problem for the BSM traffic has been traditionally tackled from two angles. First, vehicles could cooperatively control their transmission power to affect the number of BSMs audible in their neighborhood. Second, they could explicitly control the BSM frequency. In the first category, Torrent-Moreno *et al.* [11] controls the power to limit the BSM traffic to a maximum beaconing load (MBL) that the DSRC channel can allocate for the BSM traffic. This work discusses a power increment scheme that stops as soon as the MBL is reached. Mittag *et*

For Peer Review

al. [12] proposes to explicitly exchange power information with neighboring vehicles because the local adaptation proposed in Torrent-Moreno [11] could render some vehicles to violate the MBL [12]. Yoon and Kim [13] further points out that without such communication, distributed power control based on the observable shared state such as [11] can lead to extreme distortion and unfairness in terms of power setting [13].

In the frequency control category, Xu and Barth [7] decreases the frequency upon channel congestion, observable by transmission attempt failures and low packet reception rates. Khorakhun *et al.* proposes to use the channel utilization to regulate the BSM frequency [8]. If the channel busy time is below a threshold, the frequency is lowered. Drigo *et al.* [9], and Baldessari *et al.* [10] also take similar approaches. Park [14] estimates the number of neighboring vehicles from the number of BSMs to compute the optimal messaging rate.

Power control or rate control, these distributed control approaches commonly assume that the congestion metric they use should compute similarly for neighboring vehicles, and that distributed and independent adaptation by the vehicles naturally lead to power or rates fair and relevant to the given vehicle density pattern. Below, we show that such assumption is naïve, and can lead to unwarranted polarization even between neighbors, undermining fairness and neighbor awareness that will negatively affect driving safety.

III. Evaluation Experiments

In this section, we experiment with two generic versions of commonly used distributed control, namely *threshold-based* and *hysteresis-based*, and expose their instability. We use the Qualnet 4.5 simulator, which we augmented to model capture effect more precisely. In order not to diffuse the focus, we employ the simplest adaptation mechanism in this paper, the additive-increase multiplicative-decrease (AIMD). We do not consider the vehicle mobility in this paper because a recent extensive measurement study shows that 802.11p performance is not too sensitive to Doppler effect when the packet transmission time is shorter than the coherence time as is the case of short BSM transmissions [15]. Also, to show that the instability problem takes place even with the most stable channel environments, we exclude fading from the channel model. The power is set at 19 dBm so that the communication distance is approximately 300m. The simulation configuration is summarized in Table I. Finally, we limit ourselves to the investigation of rate control, where the instability problem under power control is fundamentally similar and documented in our previous work [13].

Table II show the parameters used in the distributed congestion control algorithms. For the initial BSM rate, vehicles use $R_{init} = 10$Hz. They pick a random start time for the BSM transmission, which repeats every 100ms until the rate change takes place. For the channel busy percentage (CBP, denoted by ρ), they perform the congestion metric measurement every second to use it for the rate decision for the next interval. In this paper, we use the CBP as sensed by the Clear Channel Assessment (CCA). Every vehicle obtains the channel busy percentage every T seconds. Unless otherwise mentioned, we assume that the time instance of the channel load report is not necessarily

For Peer Review

Parameter		Value
Channel	Frequency band	5.9GHz
	Path loss	2-ray ground
	Shadowing	Constant / Mean=4.0
	Fading	None
PHY	Protocol	802.11p
	Data rate	3Mbps
	Tx power	19.193 dBm
	ED threshold	-85dBm
	Capture effect	Enabled
MAC	Protocol	802.11p + 1609.4
	Channel usage	Continuous
	CW_{min}	3
Application	BSM size	300B
Vehicular	Road topology	3000m racetrack
	Headroom distance	10m (equal spacing)
	Mobility	None

TABLE I
SIMULATION SETTING FOR COMPARISON

Parameter	Meaning	Default value
R_{init}	Initial BSM rate	10Hz
R_{max}	Max BSM rate	25Hz
R_{min}	Min BSM rate	1Hz
α	Additive constant	1
β	Multiplicative constant	0.8
T	Update interval for ρ	1s
θ, θ_l	Threshold / Lower bound	0.4
θ_h	Upper bound	0.6
ρ	Channel busy percentage	

TABLE II
ALGORITHM PARAMETERS

synchronized among vehicles. For CCA, the energy detection (ED) mode is used, where the channel is marked busy whether or not the ongoing transmission is decodable as long as the energy on the channel is above the ED threshold. For the AIMD algorithms, we use the additive constant of 1 and the multiplicative constant of 0.8. So when the rate increase decision is made, it is incremented by 1 (Hz) per interval T. When the rate decrease is necessary, the current rate is scaled by 0.8. The rate is clamped between $R_{min} = 1$ and $R_{max} = 25$ Hz. For threshold-based control, the threshold θ_l is set to 0.4. Namely, an upward of 40% of channel activity during the last interval triggers the rate decrease. For hysteresis-based control, two thresholds are used in the generic simulation: $\theta_l = 0.4$ and $\theta_h = 0.6$. They are chosen by ETSI according to [16], where it is argued the offered load should be between 40 and 60 percent for efficient channel usage. As long as the congestion metric stays between these two values, the rate is not changed. But when the metric is either over θ_h or below θ_l, the rate is adjusted.

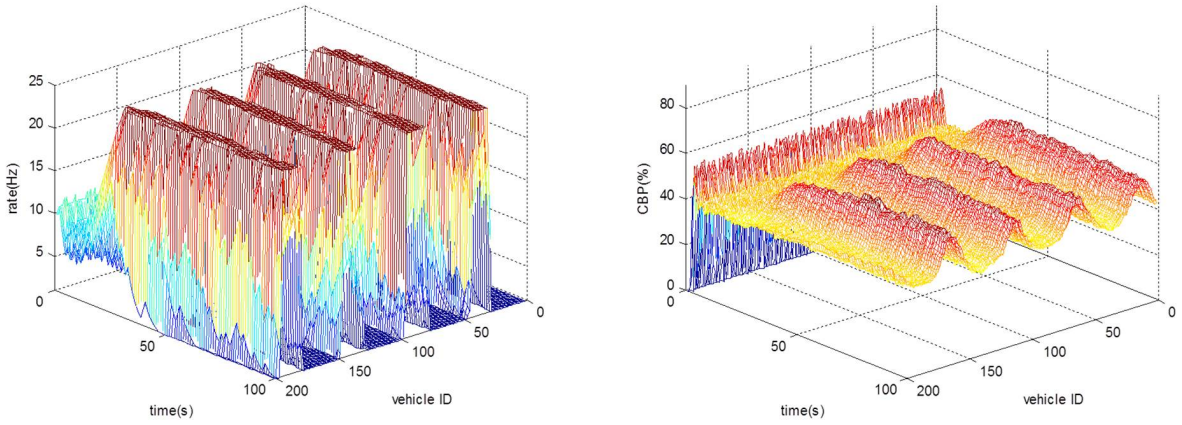
For Peer Review

A. Threshold-based control

Algorithm 1 Threshold-based control

1: $R = R_{init}$
2: **for** each period T **do**
3: **if** $\rho < \theta$ **then**
4: $R = \min(R + \alpha,\ R_{max})$
5: **else if** $\rho \geq \theta$ **then**
6: $R = \max(R \times \beta,\ R_{min})$
7: **end if**
8: Transmit with rate R
9: **end for**

First, we experiment with the threshold-based control. Here, the vehicles monitor the channel through CCA and if the energy detection (ED) reports 40% or more activity, they multiplicatively reduce the BSM frequency. But if the CBP is below the threshold, they additively increase the frequency. In order to show that the straightforward rate assignment based on a common observed congestion state such as the sensed channel busy percentage can fall into a pathology, we perform a simulation according to Algorithm 1. We place 300 vehicles equi-distant on a single lane circular racetrack of 3Km in circumference, and let them start with $R_{init} = 10$Hz. The racetrack topology is used to avoid the edge effect. Due to the uniform vehicle density, one would expect the rate settings converge to comparable levels across all vehicles. However, the experiment result points otherwise. Fig. 2 shows the time evolution of the rate assignment made by vehicles and the resulting channel busy percentage according to Algorithm 1 (for clarity, we only show 200 out of the total 300 vehicles).



(a) Adjusted rate (b) Channel busy percentage (CBP)

Fig. 2. Anomaly in generic threshold-based rate control

We can see that the rate adaptation deviates from the even vehicle pattern almost as soon as the system starts. Note that such distortion inevitably affects the neighbor awareness, as some vehicles get to transmit BSM at the

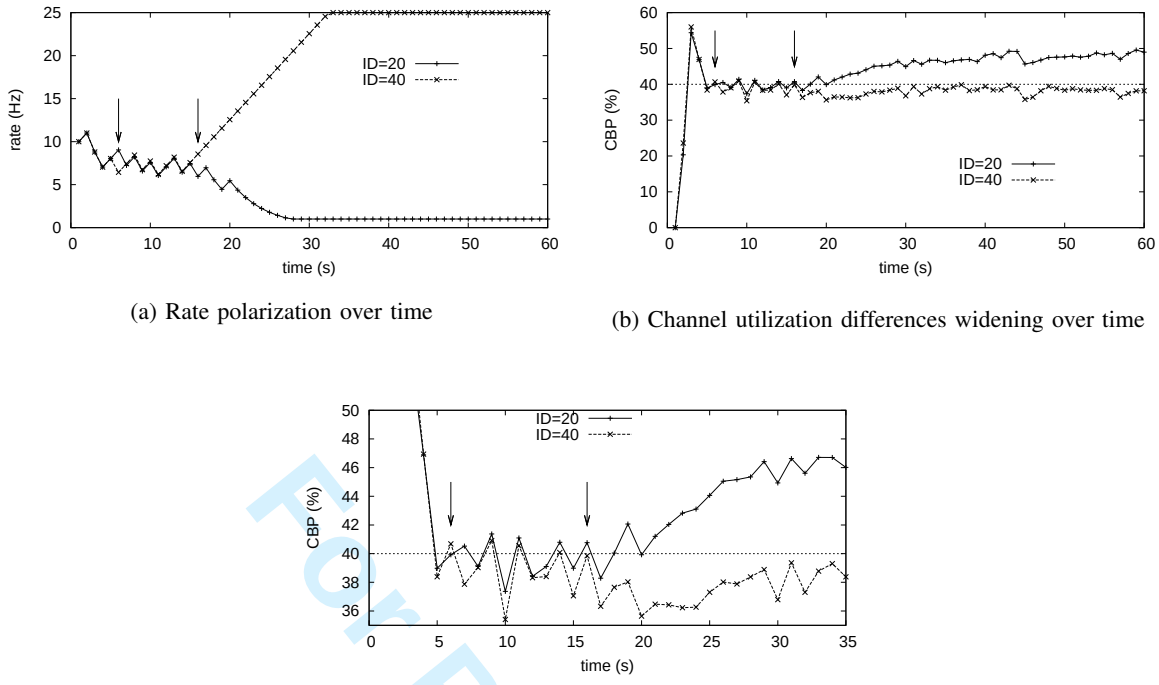
For Peer Re

lowest frequency of 1Hz. We later show that the sustainable BSM rate given the vehicle density is approximately 6Hz under the threshold of $\theta_l = 0.4$. Notice that the CBP is over θ_l where the vehicles use the lowest rate R_{min}, while it is much lower where the rate is higher, just as we predicted in Fig. 1. This is an obscure but crucial dynamics that most distributed control proposals gloss over, which arises from two aspects of distributed control based on a "common" global state:

1) Inevitable microscopic perturbation of the channel utilization around the given target threshold is frequently classified and acted upon differently by neighboring vehicles. The microscopic movements can be caused by many factors: randomness (*e.g.* from MAC contention resolution algorithm) and/or unreliability (*e.g.* from channel fading) in the WAVE communication, the granularity of control parameters, or the time-varying vehicular topology.
2) Once close vehicles take opposite actions, they may end up repeatedly applying the same measure, as the channel utilization does not change much even after the remediary actions. This is firstly because the channel state is a collective state difficult to change single-handedly and secondly because opposing rate adjustments made by neighboring vehicles can cancel out each other's effect on the channel utilization.

As a consequence of these aspects, the vehicles can diverge further away from each other. It can eventually lead to the complete polarization of rate assignments by neighboring vehicles, while looking at almost identical channel utilization, as we see in Fig. 2. In order to shed light on this dynamics in more detail, we pick out two relatively close vehicles from the previous simulation and see how their perception of the channel load and their corresponding action change over time. The two vehicles with ID=20 and 40 in Fig. 2 are approximately 150m away from each other within mutual communication range. Initially up to $t = 5$, they agree on the state of the channel (*i.e.*, CBP over θ or below it) although their numeric readings of channel busy percentage almost always differ (Fig. 3(b)). Occasionally, however, their judgements do disagree, because the threshold θ happens to come between their numeric readings. The first such instance comes at $t = 6$ (Fig. 3(c), marked by the leftmost arrow) and $t = 7$. Based on these individual judgements, these two vehicles take opposite actions: *e.g.* vehicle 20 lowers the rate and 40 raises it at $t = 6$ (Fig. 3(a)). Another disagreement comes at $t = 16$, and every time such incident repeats, their perceived channel busy percentage increasingly deviates (Fig. 3(b)), exactly in the opposite direction of their intended control (Fig. 3(a)). This is what we predicted in Fig. 1: a minor differences in the perceived channel utilization pushing neighbor vehicles towards opposite rate extremes while the channel state largely stays around the target threshold. Almost all existing works pay attention to the latter, but they gloss over the unfairness among vehicles, the cause, and its safety implications.

For Peer Review



(a) Rate polarization over time

(b) Channel utilization differences widening over time

(c) Channel utilization (magnified)

Fig. 3. How the rate polarization is triggered and amplified

B. *Hysteresis-based control*

From the experiments in the previous section, one might draw a conclusion that the threshold should not be so "fine" that it sets apart neighbors with similar perceptions of the channel utilization. Namely, one might propose to enlarge the boundary between the congested and the non-congested regimes. This idea essentially advocates hysteresis. A recent standard from the European Telecommunications Standards Institute (ETSI) indeed employs it in the distributed congestion control for vehicular safety messaging [17]. In this section, we perform a simulation experiment to check whether or not the hysteresis eliminates the problem discovered in the previous section. To tell the conclusion first, unfortunately, it does not, and we demonstrate it below.

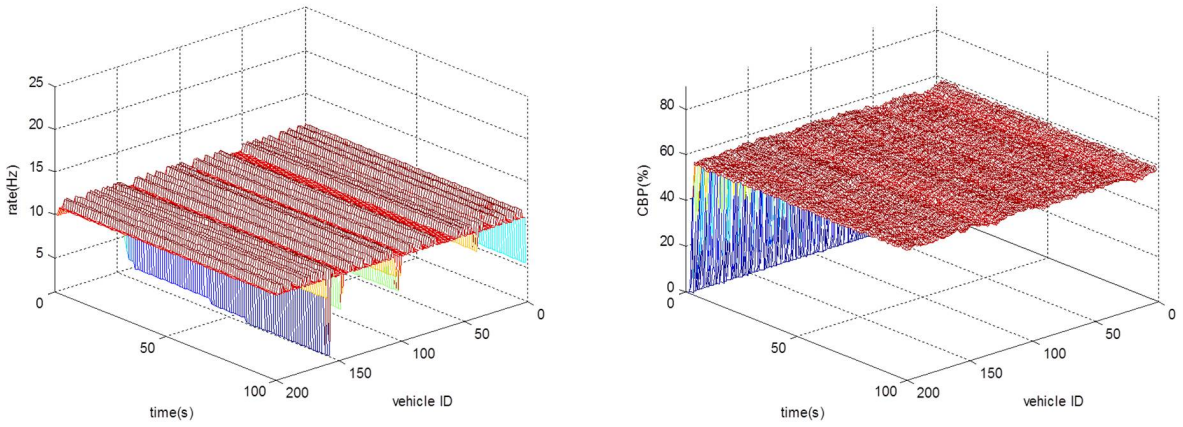
1) Generic hysteresis-based control: A generic algorithm based on the idea of hysteresis is shown in Algorithm 2. Notice that the hysteresis-basis control is the same as the threshold-based control except that two thresholds are used instead of a single threshold. The behavior of the algorithm is the same when the channel busy percentage is over θ_h or below θ_l. But when the perceived channel busy percentage stays between θ_h and θ_l, the rate is left unchanged.

Fig. 4 shows how the rate and the channel busy percentage change according to Algorithm 4. We notice that the hysteresis-based control generally appears more stable. However, some vehicles suffer a drastic rate drop, most conspicuously the vehicle with ID=158, whereas the channel busy percentage is not visibly anomalous. In order to

For Peer Review

Algorithm 2 Hysteresis-based control

1: $R = R_{init}$
2: **for** each period T **do**
3: **if** $\rho < \theta_l$ **then**
4: $R = \min(R + \alpha,\ R_{max})$
5: **else if** $\rho > \theta_h$ **then**
6: $R = \max(R \times \beta,\ R_{min})$
7: **end if**
8: Transmit with rate R
9: **end for**



(a) Adjusted rate (b) CBP

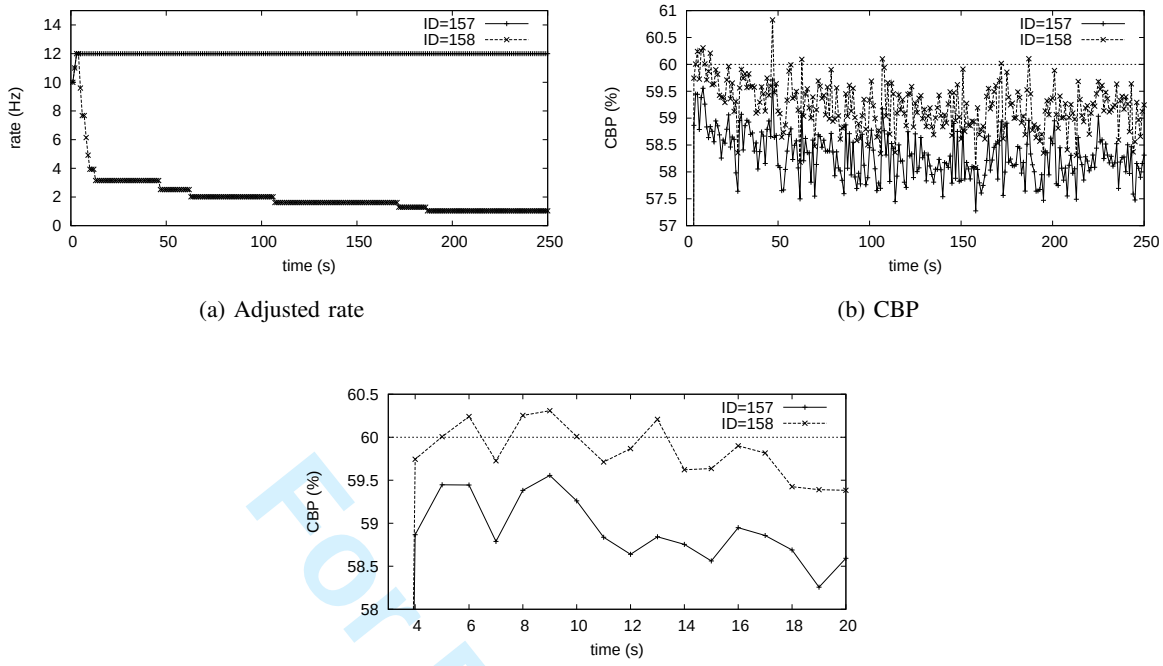
Fig. 4. Anomaly in hysteresis-based rate control

understand how the pathology develops, again we examine the rate and channel busy percentage dynamics for two vehicles in and around the collapse, with ID=157 and 158, respectively.

In Fig. 5, we see that the two vehicles initially start with the same rate for a couple of seconds. But then the rate on vehicle 158 suddenly dips to very low values, never to recover. Around this incident, we notice that the channel busy percentage at vehicle 158 temporarily goes beyond θ_h. For vehicle 157, on the other hand, the channel busy percentage is maintained below θ_h albeit close to it. Although transiently and to a minor degree, the two vehicles opposite each other across θ_h, and a few breaches over θ_h are enough to reduce the rate of vehicle 157 close to 2Hz. A peculiarity of the hysteresis-based control is that the vehicle cannot begin to grow the rate again, unless its perceived channel busy percentage goes below θ_l. Therefore, once lowered, the rate has to stay low until there is a large disruption in the vehicle topology such as a traffic density drop.

The above experiment tells us that putting a buffer zone between congested and non-congested regimes is not the solution to the anomalous rate adaptation behavior among vehicles. Rather, the result reaffirms the analysis of the previous section that minor fluctuations (see the y-axis scale of Fig. 5(b)) and consequent opposing classifications of the channel state are enough to trigger and aggravate the extensive rate adjustment deviations among neighbors.

er Review



(a) Adjusted rate (b) CBP

(c) CBP (magnified)

Fig. 5. Anomaly in generic hysteresis-based rate control

2) ETSI style rate control: Recently, the ETSI standardized a decentralized congestion control (DCC) algorithm for vehicular traffic [17]. To control congestion, it employs a state machine with three states: Relaxed, Active, and Restrictive. These states assign different power, messaging rate, channel sensitivity, and PHY transmission modes. Among them, we limit ourselves to the messaging rate control aspect, to put it in perspective with other control schemes considered in this paper. Fig. 6 shows the rate control state machine derived from the ETSI DCC. In the ETSI specification, the rate increase decision is made when the given condition on channel busy percentage holds for 5 (=`NDL_timeUp`) seconds. Likewise, the rate decrease is made when the required condition holds for 1 (=`NDL_timeDown`) seconds [17].

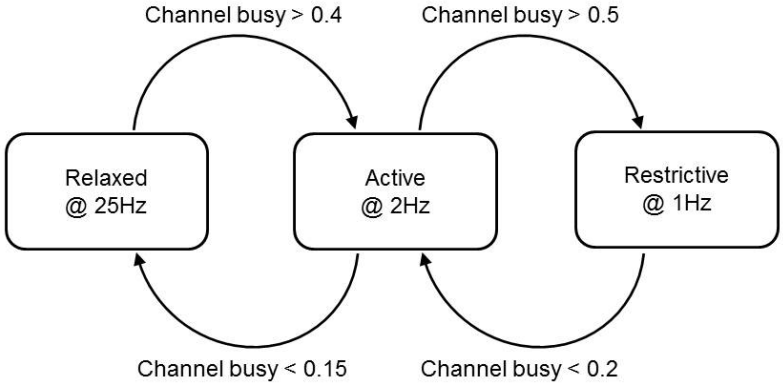
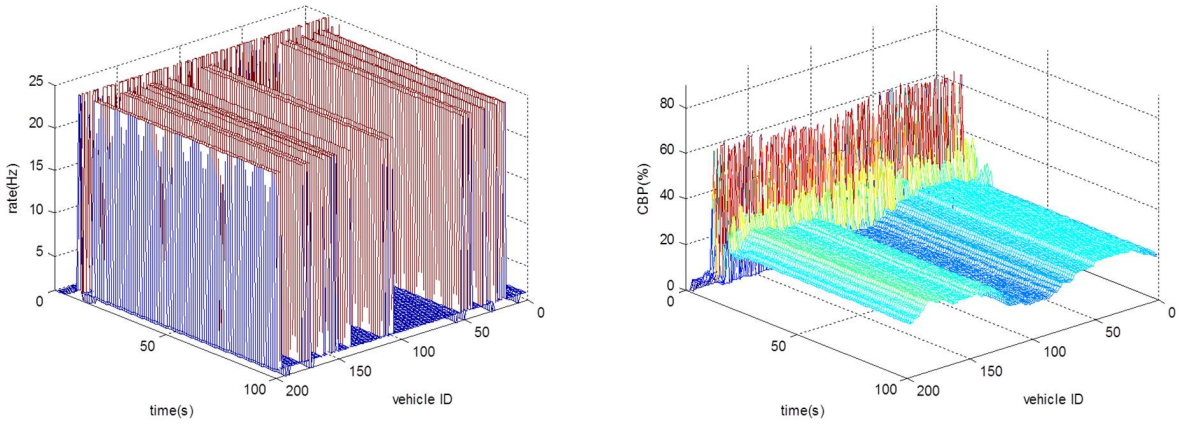


Fig. 6. Rate transition diagram derived from the ETSI DCC access state machine: required conditions marked on arcs

For Peer Review

Fig. 7 shows the simulation results of the state machine, with all parameters set identically with the threshold- and hysteresis-based experiments. In Fig. 7(a), we notice that the unfairness in the rate setting is severe among neighboring vehicles. Still, the channel busy percentage as shown in Fig. 7(b) does not show a correspondingly wild deviations although it shows some differences. In order to understand the dynamics behind the rampant polarization between the rates, we pick and investigate two neighboring vehicles that are pushed to different extremes, vehicles 46 and 47.



(a) Adjusted rate (b) CBP

Fig. 7. Anomaly in ETSI-derived rate control

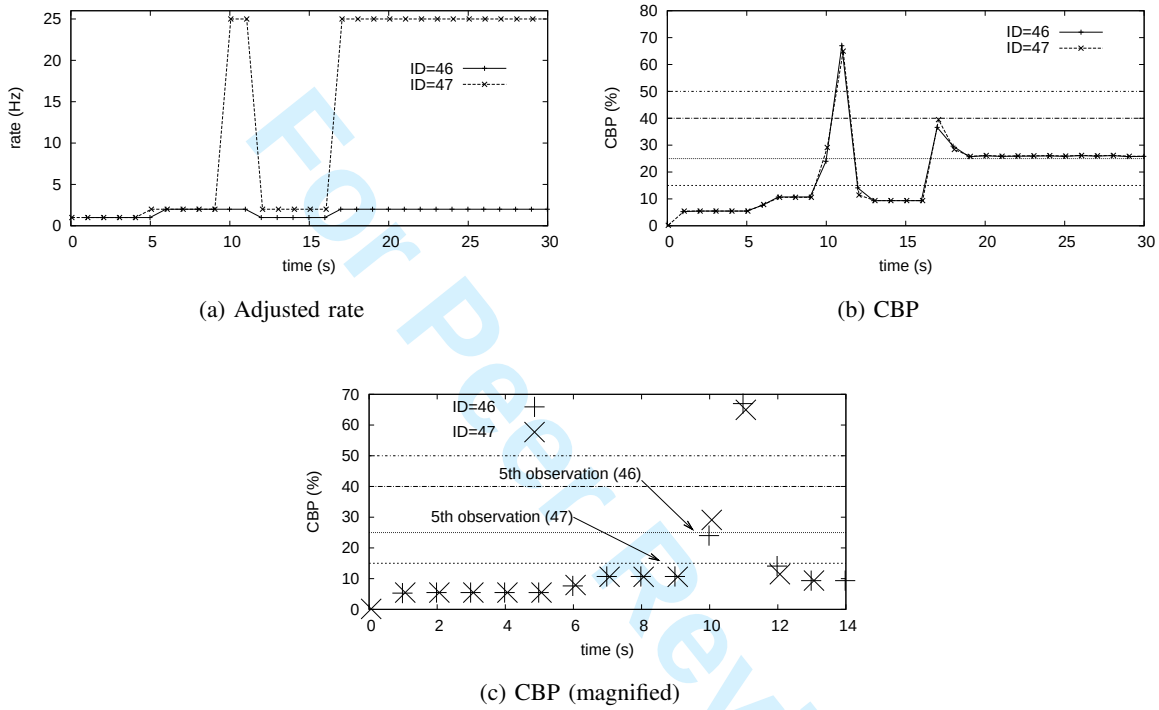
Fig. 8(a) shows that these two vehicles indeed takes radically different rate adjusting steps. However, Fig. 8(b) hardly exhibits any clue that would lead to such pathology. This is one of the important points that this paper makes. If we only focus on the average packet delivery ratio and average channel load, the pathological rate polarization (which has no justification whatsoever) and its potential issues go unnoticed.

Although a macroscopic view of the perceived CBP (Fig. 8(b)) does not show big differences between the two neighboring vehicles just few meters apart, the microscopic version (Fig. 8(c)) explains why the polarization talks place. First, we notice that the channel assessment timing for the vehicle 46 is behind that of the vehicle 47 (Recollect that the CBP measurement time instances are not synchronized among vehicles). According to the ETSI DCC, rate can go up only when the low channel utilization is maintained for at least 5 seconds. Indeed, the two vehicles observe less than 15% CBP initially for 5 consecutive observations. This leads to the rate increase to 2Hz for vehicle 47 and 46 at time $t = 4.06$ and $t = 4.98$, respectively. The low CBP condition persists another 5 seconds for node 47, until when it takes the channel measurement at $t = 9.06$. Accordingly, vehicle 47 reflects the result to its rate adjustment, raising it to 25Hz. But at $t = 9.98$ when it is the turn for vehicle 46 to take the channel assessment, the CBP is no longer below 15%, partly because its neighbor vehicle 47 (and also some others now shown here for clarity) jumped their rates to 25Hz. Therefore, only by a split second difference, the CBP readings

Peer Review

by the two neighbor vehicles can be on different sides of a threshold and lead to exactly opposing reactions.

The same course of action repeats at around $t = 16$, where the fifth consecutive readings of lower than 15% CBP push the rates of both vehicles by one notch to 25Hz and 2Hz, respectively. But as a consequence of the rate increase by the vehicles, the CBP first jumps to near 40% and eventually hovers just above 25%. Interestingly, this condition holds both vehicles in their current state (see Fig. 6), so the neighboring vehicles end up having radically different BSM rates even though their readings of CBP are almost identical.



(a) Adjusted rate

(b) CBP

(c) CBP (magnified)

Fig. 8. Analysis of ETSI-derived hysteresis

From the above experiment, one might think that the polarization between neighbors stems from the time difference between their CBP observations — and that synchronized CBP readings may cure the pathology demonstrated above. Indeed, synchronization may be relatively easily achieved using the GPS time in the IEEE WAVE. However, our further investigation shows that at best it causes another undesirable problem of oscillation, that leads to repeated episodes of high packet loss rates [18] on the condition that all vehicles start at the same time. Under a more realistic assumption that the start times of the vehicles are all different, however, the pathological property resurfaces even if the CBP readings are tightly synchronized. In the next section, we will present a very simple but robust solution that yields neither rate polarization nor oscillation.

For Peer Review

C. Mean-checked threshold-based control

The key idea we put forth in this paper is that the congestion control should be not only distributed but also coordinated. Making the IEEE WAVE system behave so is simple. Specifically, in the BSM message format [4], the rate information can be carried in the Part II as a local content. By doing so, vehicles can inform of their BSM rate through each broadcast BSM so that the neighboring vehicles that receive the BSM can use it to obtain the average BSM rate of the neighborhood. Once the average is obtained, we can use it as an additional check whenever making a rate adjustment decision. We call it the mean-checked rate control in this paper. Note that it is orthogonal to the rate control schemes discussed above, and it can augment them to cure their pathology. Algorithm 3 is a version that augments the thresh-based control as described in Section III-A. Note that the only difference in the algorithm compared with the original threshold-based control is that the rate assignments in one's neighborhood $\mathcal{N}$ is additionally considered before the rate assignment is made (*i.e.*, lines 4 and 8). The neighborhood is defined by the vehicles in one's communication range, namely $\mathcal{N} = \{v_1, v_2, \ldots, v_n\}$. Below, we will call this mean-checked threshold-based control simply by the mean-checked control.

Algorithm 3 Mean-checked rate control

1: $R = R_{init}$
2: **for** each period T **do**
3: **if** $\rho < \theta$ **then**
4: **if** $R < \bar{R}(\mathcal{N})$ **then**
5: $R = \min(R + \alpha,\ R_{max})$
6: **end if**
7: **else if** $\rho \geq \theta$ **then**
8: **if** $R \geq \bar{R}(\mathcal{N})$ **then**
9: $R = \max(R \times \beta,\ R_{min})$
10: **end if**
11: **end if**
12: Transmit with rate R
13: **end for**

Although looking simple, the addition of two checks against the neighborhood average rate $\bar{R}(\mathcal{N})$ leads to a significant change in the structure of the control. We can visualize this using a vector field in Fig. 9 that reveals the dynamics of the mean-checked control. The vector field is easily obtained by consulting the algorithm at a given rate under a given channel utilization level. For the sake of illustration, let us suppose $\mathcal{N} = \{v_1, v_2\}$ and the desirable operating point O^* is reached with $\bar{R}(\mathcal{N}) = 12.5$Hz. Note that O^* is where the two diagonals meet in the vector field. The diagonal from the top-left to the bottom-right corner represents the sum rate that yields the target channel load, whereas the one that runs from the bottom-left to the top-right corner is the fair rate line satisfying $r_1 = r_2$. Let us call the former "sum rate diagonal" denoted by Δ_s and the latter, "equal rate diagonal" denoted by Δ_e. We remark that in the threshold- and hysteresis-based schemes discussed above, there is no constraint that

For Peer Review

relate the BSM rates of neighbors other than that the sum rate should not be so large as to overshoot O^*. For instance, with $r_1 + r_2 \leq 25$Hz as the constraint, r_1 and r_2 can diverge in the opposite directions yet meeting the sum rate constraint. Under the circumstances, no definite vector field can be formed. For instance, with $\mathcal{N} = \{v_1, v_2\}$, one could move almost in any directions as long as the sum rate moves toward Δ_s.

In the mean-checked control, the vector field is clearly defined. Below the sum rate diagonal Δ_s, if $r_1 < \bar{R}(\mathcal{N}) < r_2$ (a condition that corresponds to the operating regime marked "C" in Fig. 9), the vector field pushes only r_1 to increase. This will continue until the operating point reaches either Δ_e or Δ_s. Then the vector field pushes both r_1 and r_2 along Δ_e or Δ_s respectively, towards O^*. Notice that it implies the mean-checked control exercises an equalizing force on the rates instead of polarizing them. Similar convergence occurs for regime "D." Over Δ_s, if $r_1 < \bar{R}(\mathcal{N}) < r_2$ ("B" in Fig. 9), r_2 is pulled down while r_1 holds until the operating point hits either Δ_e or Δ_s, along which the rates are pushed down towards O^*.

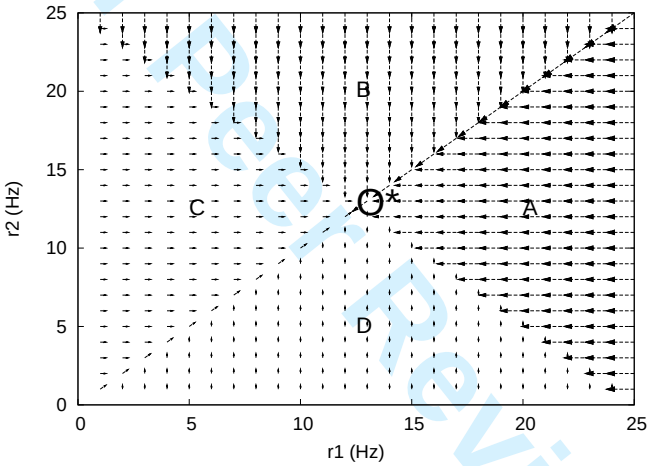
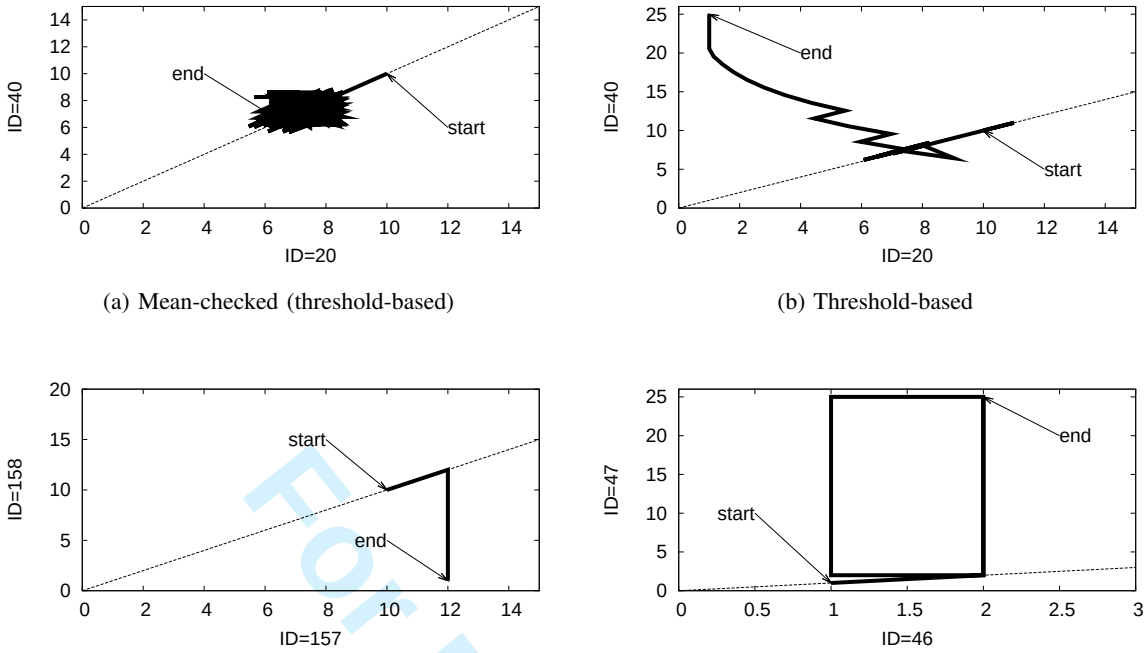


Fig. 9. Vector field for the mean-checked control for two vehicles, as dictated by Algorithm 3

Since the rate decrease is multiplicative, the rate of convergence is proportional to the current rate of the vehicles over Δ_s. It is shown in Fig. 9 as larger arrow sizes for the positions farther from the center. On the other hand, the rate is constant under Δ_s since the rate increase is additive. The rate of convergence can be controlled by setting α and β differently. Actual vector field should be multi-dimensional, where the number of dimensions is determined by the number of vehicles in the communication range. As a matter of fact, Fig. 10(a) plots the rate adjustments made by randomly selected two neighbor vehicles in the mean-checked control as the multi-dimensional vector field forces them. The diagonal line plots Δ_e. The starting point ("start") is apparently not sustainable, and as we discussed above, the rates move down along Δ_e as we predicted in Fig. 9. Although frequently thrown away from Δ_e, the operating point stays close to it, supposedly in the neighborhood of O^*. In contrast, the threshold-based control shown in Fig. 10(b) drifts away from Δ_e (and obviously from O^* as it is on Δ_e). Fig. 10(b) is another representation of Fig. 3, and we can clearly see the difference in the control dynamics of these two schemes. As

For Peer Review



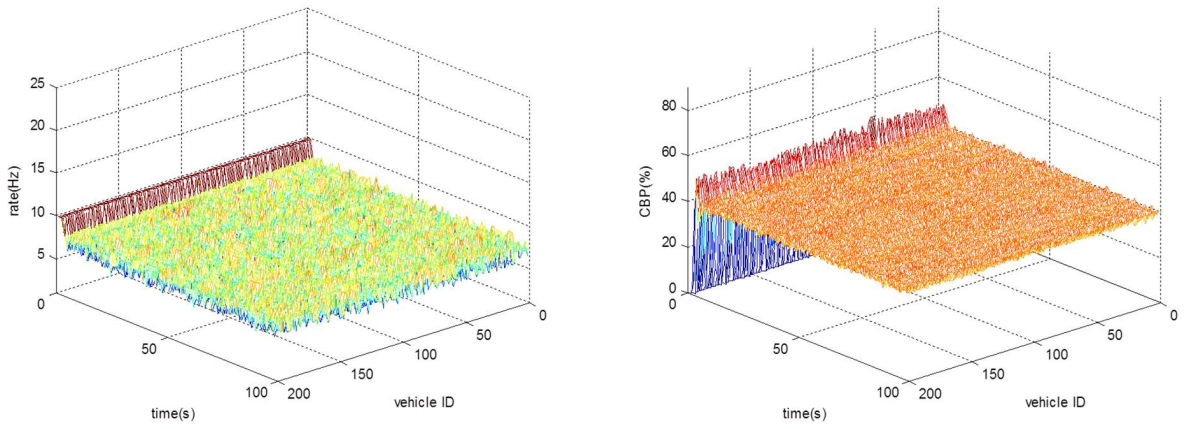
(a) Mean-checked (threshold-based) (b) Threshold-based

(c) Hysteresis-based (d) Hysteresis-based (ETSI-style)

Fig. 10. Rate movements of neighboring nodes in the threshold-based control and mean-checked control

for hysteresis-based controls, we observe that they also are easily knocked out from the equilibrium, and show less dynamics due to the suppressing effects of the hysteresis.

Fig. 11 shows the rate assignments made by the mean-checked control algorithm. With only a minor modification to the threshold-based algorithm, it successfully finds the rate equilibrium for neighboring vehicles, while maintaining the target channel utilization.



(a) Adjusted rate (b) CBP

Fig. 11. Mean-checked (threshold-based) rate control

For Peer Review

Finally, we vary the vehicular traffic density along the road in order to check if the mean-checked algorithm successfully regulates the BSM rate, reflecting the given traffic pattern. We give the varying traffic density at 200 positions on the road track as in Fig. 12, where the density value is defined to be the number of neighboring vehicles in 100m radius. Fig. 13(a) and (b) compare the rate adaptation results based on the threshold-based control and the mean-checked control, respectively. As we notice, the threshold-based control leads to anomalous rate assignments that are largely irrelevant to the given traffic pattern. In contrast, the mean-checked control converges to one that correctly reflects the density. The reason that there is relatively long perturbation at the beginning of the simulation is because we let the vehicles start randomly within 50 seconds from the beginning of simulation.

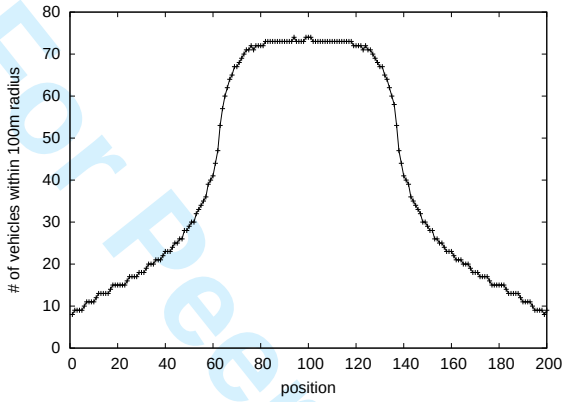
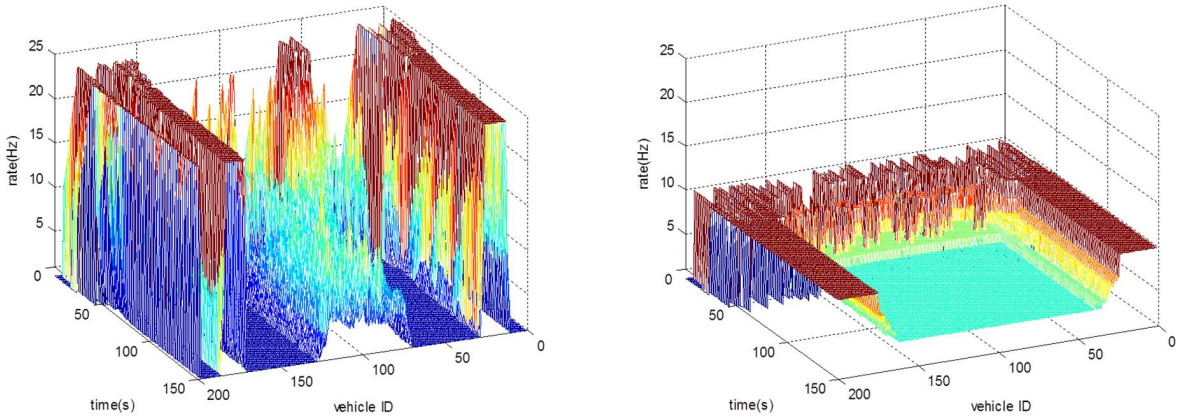


Fig. 12. Traffic density



(a) Rate evolution under threshold-based control (Hz) (b) Rate evolution under mean-checked control (Hz)

Fig. 13. Adaptation under varying vehicle traffic density

Last but not least, we made a few simplifying assumptions to avoid diffusing the focus on the core dynamics of the uncoordinated distributed congestion control of the IEEE WAVE safety messages. We can relax them and investigate if more realistic assumptions significantly change the qualitative conclusion that we draw above. For

instance, we can assume a fading channel, significantly dispersed starting times for vehicles, using a running average of channel busy percentage over a longer period of time instead of one-second interval, different rate adaptation constants α and β, and uneven inter-vehicle distance distributions among others. They are too numerous, so we omit the discussion in this paper. In [18], however, we experiment with all these cases and demonstrate that they do not change the dynamics we observed above, and that coordinated control is the only remedy to the pathology. It is only natural, since the root cause of the pathology we discussed so far comes none of these factors, but from the fact that the inevitable microscopic perturbation of the channel utilization around the given target threshold is frequently classified differently by neighboring vehicles.

IV. Conclusion

Most congestion control proposals in the IEEE WAVE context aim at regulating the BSM transmission rate (or power) to not exceed a certain channel utilization threshold. What these proposals have failed to grasp, however, is that individual rate (or power) settings of even closely neighboring vehicles can dramatically diverge while their perceived channel utilization are apparently identical to a first degree. Such pathology has a negative implication on the levels of awareness they will offer to ambient vehicles, raising a safety issue.

In this paper, we first uncover the root cause of the pathology, that the inevitable microscopic perturbation of the channel utilization around the given target is frequently classified differently by neighboring vehicles. The microscopic movements can be caused by the randomness (*e.g.* from MAC contention resolution algorithm) and/or unreliability (*e.g.* from channel fading) in the WAVE communication, granularity of rate adaptation parameters, or by the time-varying vehicular topology. The problem is that once close vehicles take opposite actions, they can repeatedly apply the same measure, as frequently the channel utilization does not change much first due to it being a collective state and second due to opposing rate adjustments made by the vehicles canceling each other out. As a consequence, the vehicles can diverge further away from each other. It eventually leads to the complete polarization of rate assignments by neighboring vehicles, while having almost identical channel utilization.

We propose a remedy to the pathology, which is composed of two steps. First, each vehicle explicitly professes its current rate assignment in the BSM message in the optional Part II. Second, using the averaged rate as a safeguard, they only change their rates if it does not violate the safeguard. This simple coordinative ingredient yields a provably stable control so that vehicles in the same neighborhood converge their rates. We demonstrate the pathology and the effect of the proposed solution through extensive simulation.

The most important lesson from our study in this paper is that for the IEEE WAVE safety messaging, a cooperative congestion control among vehicles cannot just rely on distributed algorithms independently executed on each vehicle. Rather, explicit coordination is a requirement, and introducing this cautionary measure to existing schemes is not expensive, as the bandwidth and the computation overhead for the additional check are minimal.

For Peer Review

REFERENCES

[1] IEEE 802.11 WG, Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications, Amendment 6: Wireless Access in Vehicular Environments, IEEE Std 802.11p-2010, 2010.

[2] IEEE 1609 WG, IEEE Standard for Wireless Access in Vehicular Environments (WAVE) — Networking Services, IEEE Std 1609.3-2010, 2010.

[3] IEEE 1609 WG, IEEE Standard for Wireless Access in Vehicular Environments (WAVE) — Multi-channel Operation, IEEE Std 1609.4-2010, 2010.

[4] SAE International, Dedicated Short Range Communications (DSRC) Message Set Dictionary, SAE J2735, November 2009.

[5] CAMP Vehicle Safety Communications Consortium, Vehicle safety communications project task 3 final report: Identify intelligent vehicle safety applications enabled by DSRC, Tech. Rep. DOT HS 809 859, March 2005. Available at http://www.its.dot.gov/research_docs/pdf/59vehicle-safety.pdf.

[6] Vehicle Safety Communications – Applications (VSC-A) Final Report, Tech. Rep. DOT HS 811 492A, September 2011.

[7] H. Xu and M. Barth. A Transmission Interval and Power Level Modulation Methodology for Optimizing Inter-Vehicle Communications. In *Proceedings of ACM VANET*, 2004.

[8] C. Khorakhun, H. Busche, and H. Rholing. Congestion Control for VANETs based on Power or Rate Adaptation. In *Proceedings of Workshop on Intelligent Transportation (WIS)*, 2008.

[9] M. Drigo, W. Zhang, R. Baldessari, L. Le, A. Festag, and M. Zorzi. Distributed Rate Control Algorithm for VANETs (DRCV). In *Proceedings of ACM VANET*, 2010.

[10] R. Baldessari, D. Scanferla, L. Le, W. Zhang, and A. Festag. Joining Forces for VANETs: A Combined Transmit Power and Rate Control Algorithm. In *Proceedings of Int. Workshop on Intelligent Transport (WIT)*, 2010.

[11] T. Torrent-Moreno, P. Santi, and H. Hartenstein. Fair Sharing of Bandwidth in VANETs. In *Proceedings of ACM VANET*, 2005.

[12] J. Mittag, F. Schmidt-Eisenlohr, M. Killat, J. Harri, and H. Hartenstein. Analysis and Design of Effective and Low-Overhead Transmission Power Control for VANETs. In *Proceedings of ACM VANET*, 2008.

[13] Y. Yoon and H. Kim. Resolving Distributed Power Control Anomaly in IEEE 802.11p WAVE. *IEICE Transactions on Communications*, January 2011.

[14] Y. Park and H. Kim. Application-Level Frequency Control of Periodic Safety Messages in the IEEE WAVE. *IEEE Transactions on Vehicular Technology*, 61(4):1854–1862, May 2012.

[15] F. Bai, D. D. Stancil, and H. Krishnan. Toward understanding characteristics of dedicated short range communications (DSRC) from a perspective of vehicular network engineers. In *Proceedings of ACM Mobicom*, 2010.

[16] C2C-CC White Paper on Network Design Limits, C2008, 2008.

[17] ETSI, Intelligent Transport Systems (ITS); Decentralized Congestion Control Mechanisms for Intelligent Transport Systems operating in the 5 GHz range; Access layer part, ETSI TS 102 687, July 2011. Available at http://www.etsi.org/deliver/etsi_ts/102600_102699/102687/01.01.01_60/ts_102687v010101p.pdf.

[18] B. Kim and H. Kim. Avoiding the Pitfalls of Distributed Congestion Control in the IEEE WAVE Safety Messaging. Technical report, Korea University, November 2012. Available at http://widen.korea.ac.kr/wiki/index.php?title=Special:Upload&wpDestFile=Anomaly-techreport.pdf.

For Peer Review