

Let's Talk in Groups: A Distributed Bursting Scheme for Cluster-based Vehicular Applications

Michele Segata^{*,a}, Falko Dressler^b, Renato Lo Cigno^a

^a*Department of Information Engineering and Computer Science, University of Trento, Italy*

^b*Heinz Nixdorf Institute and Department of Computer Science, Paderborn University, Germany*

Abstract

Clustering has been an important concept in the scope of vehicular networks. The idea is to reduce channel contention, enable building backbones and might improve spatial reuse. We propose a seemingly simple, yet unexplored idea: extending IEEE 802.11 frame bursting MAC access to multiple stations aggregated into a cluster. We call this approach Distributed EDCA Bursting (DEB). The focus of our work is not on building clusters, but exploring what is the gain that can be achieved by the standard IEEE 802.11p channel access if we introduce the principle of frame bursting (presently not allowed in IEEE 802.11p standard, but the key factor for the efficiency of IEEE 802.11n/ac WiFi channel access). The fundamental idea is to extend the standard frame bursting mechanism so that only the cluster head contends for the channel. Upon winning the contention, the reserved transmission opportunity is shared among the vehicles in the cluster, which transmit a coordinated burst of frames. We describe the idea in detail, particularly analyzing theoretical benefits and limitations of the approach. We implement DEB in the platooning simulation framework PLEXE-Veins and analyze its behavior under different loads (i.e., number of vehicles sharing the channel), transmit power policies, and different MAC and physical layer parameters. We present benefits and drawbacks of our approach, showing that it can overcome channel limits in saturation conditions, reduce channel usage at moderate loads, and improve spatial reuse. Moreover, by coordinating and distributing channel access, DEB reduces the number of collisions at the expenses of a slightly increased average beacon inter-arrival time.

Key words: Inter-Vehicle Communication, Clustering, EDCA Bursting, IEEE 802.11

1. Introduction and Related Work

With the advances in autonomous driving, new demands are being defined for the second generation of such systems: cooperative autonomous driving. Intended already at early stages of Inter-Vehicle Communication (IVC) research, this application poses new requirements in terms of road traffic safety on Intelligent Transportation Systems (ITS) and particularly on the communication protocols. In more than 10 years of IVC research, many applications have been foreseen and new communication primitives have been investigated [1, 2]. Obviously, sharing information among vehicles to implement intelligent cooperation is a fundamental basis for most of them. With the standardization of IVC communication protocols on both the lower (IEEE 802.11p) and the higher layers (IEEE WAVE, ETSI ITS-G5, ARIB T109), the vehicular networking community reconciled on the problems solved and those that need further fundamental research [3]. It turned out that the initial protocols provide a solid basis for many applications relying on short range communication and a fully distributed management of the resources.

Yet, distributed safety applications were identified to be not fully covered. Examples range from emergency braking and intersection collision avoidance [4, 5], where vehicles share their position, speed, and trajectory while approaching an intersection to forecast potential collisions and promptly warn the driver, or simply implement an optimal stopping strategy to platooning [6, 7, 8], which organizes vehicles in groups driven by an automatic system, minimizing inter-vehicle gaps (thus improving traffic flow) as well as the risk of collisions. The spectrum also covers a wide range of other application domains including next generation traffic lights up to completely vitalized versions such as Virtual Traffic Light (VTL) [9, 10, 11], where standard traffic lights are substituted by a cooperative, self-organizing application that enables vehicles to automatically synchronize, reducing useless idle times.

In order to provide the right level of service guarantees, particularly focusing on low latency and high reliability for meeting the hard deadlines of these distributed real-time systems, completely new concepts have been investigated for the underlying communication technologies. Heterogeneous vehicular networks [3] are emerging combining cellular networks, short range radio broadcast, as well as Visible Light Communication (VLC) and millimeter-wave communication using the vehicular radar for communica-

* Corresponding author *Email addresses:*
msegata@disi.unitn.it (Michele Segata), dressler@ccs-labs.org
(Falko Dressler), locigno@disi.unitn.it (Renato Lo Cigno)

tions. For example, VLC [12] turned out to be a perfect candidate for offloading communication from the radio network, thus, reducing the load on the wireless channel [13].

In this paper, we revisit communication issues in the scope of IVC based on IEEE 802.11p. This protocol is, as all the Wireless LAN (WLAN) standards in the IEEE 802.11 family, using a contention based channel access. Thus, even at moderate channel loads, the protocol suffers from packet losses due to congestion. To cope with this problem, the research community proposed several channel congestion algorithms, i.e., protocols that adapt the transmission rate depending on some parameters (e.g., the current channel congestion state) to keep congestion under control and avoid packet losses (see [14, 15, 16, 17] to name a few). The main idea of these concepts is to cope with the problem of channel congestion from an application layer point of view, i.e., without modifying the underlying IEEE 802.11p protocol.

Interestingly, the IEEE 802.11 standard [18] already proposes some mechanisms that can help reducing congestion and improve the overall network throughput. Particularly being developed for multimedia transmissions over WLAN, the 802.11e amendment [19] introduced several new features to support Quality of Service (QoS). Different priorities were defined for different types of traffic in combination with an appropriate scheduling strategy for handling different transmit queues for the respective traffic types. This Enhanced Distributed Channel Access (EDCA) standard also became part of IEEE 802.11p.

Another feature introduced in EDCA but not yet explored in the scope of IEEE 802.11p (and, thus, vehicular networking applications), is the frame bursting feature of the EDCA. The idea of frame bursting is that a station does not contend for the channel for a single transmission, but for a certain amount of time defined as a Transmission Opportunity (TXOP). During this TXOP, a station is allowed to send multiple frames in a row, reducing protocol overhead and improving fairness among stations with different link qualities (see Section 2 for the detailed description of the mechanism).

In fact, frame bursting is currently defined for unicast flows only. In our approach, we exploit bursting for any-to-any communication. In addition, our mechanism is distributed and exploits the knowledge of the communication pattern. In previous work [20], we proposed a distributed implementation of the frame bursting feature, i.e., a station that wins a contention shares its TXOP with a group of vehicles, like a cluster. The idea is based on the observation that several applications explicitly or implicitly employ a clustering mechanism to share data among groups of vehicles. One example is intersection collision avoidance, where clusters of vehicles approaching the intersection naturally form on the road. Vehicles on one road needs to coordinate (e.g., determine which is the closer to the intersection), but also communicate and coordinate with vehicles on the other roads. Organizing the channel access per road-cluster can only improve efficiency and

reduce channel contention. This is especially true when considering heterogeneous communication [21]. Another safety-related application is platooning: a platoon naturally forms a cluster of vehicles which share data for automatic control purposes. However, cluster-based communication in vehicular networks has received so much attention in the past [22, 21, 23] that there is no need for advocacy on its potential role in IVC.

In this paper, we extend our previous work on EDCA bursting in vehicular environments [20]. We not only demonstrate the benefits of a distributed EDCA bursting approach for a specific use-case, i.e., a platooning application, but also explore the potential gain of our proposal in more detail. The choice of platooning relies on the fact that this application creates, by definition, stable cluster of vehicles with a cluster-head, i.e., the leader. Without loss of generality, our concepts can be used in all cluster-based vehicular communication scenarios.

Our main contribution can be summarized as follows:

- We propose the distributed EDCA bursting mechanism, explaining in detail how standard EDCA bursting works and what is needed to implement that in a distributed fashion (Section 2);
- we study the potential resource gains in depth to explore the theoretical benefits and limitations of the system (Section 3);
- we analyze the performance of our approach for a platooning application, showing its benefits and its limitations compared to the classic channel access mechanism (Section 4);
- we implement a pre-scheduling mechanism and investigate its benefits in the platooning scenario (Section 5);
- we test the behavior of DEB under external interference by non-clustered vehicles (Section 6); and
- we list some research questions that still needs to be addressed before adopting this approach in real life applications (Section 7).

2. IEEE 802.11 Bursting

2.1. Standard Protocol Behavior

The IEEE 802.11e amendment introduced several features to differentiate service, enhance throughput and improve fairness in WLANs. The Enhanced Distributed Channel Access (EDCA) distinguishes traffic mapping it to four different MAC queues, as opposed to the prior access scheme (named Distributed Coordination Function (DCF)), where all frames are managed by a single queue. Each EDCA queue has different MAC parameters, such as the amount of time spent in carrier sensing before transmitting or backing-off, the Contention Window (CW) size,




Figure 1: EDCA bursting (as per 802.11e standard) and proposed Distributed EDCA Bursting (DEB).

maximum transmission duration, etc. With EDCA, regardless of the number of logical queues, a station contends the channel for a TXOP, and can send multiple frames as long as it holds the right to transmit. In general this feature is called *frame bursting*. Chipsets pre-802.11e, instead, contend for the channel to send a single frame: N frames require N contentions. This channel access mechanism results in an enormous amount of overhead due to the backoff procedure, and causes unfairness between stations using different link speeds. Imagine two stations with a link speed of 6 Mbit/s and 54 Mbit/s, respectively, that need to send MPDUs of the same size. If they access the channel one after the other repeatedly, the measured application layer throughput will be the same for both, but it will be lower than 6 Mbit/s. By assigning the channel to the stations for a certain amount of time (a TXOP), instead, the two stations will fairly share the channel time, and their application layer throughput will depend only on their link quality and the amount of stations concurrently trying to access the channel.

The 802.11e standard proposes two ways for obtaining a TXOP. The first one is through standard EDCA contention: a station that wins a backoff contention obtains the TXOP. The second one is through the HCF Controlled Channel Access (HCCA), i.e., when the Access Point (AP) implements the Hybrid Coordination Function (HCF). In this mode, the AP divides the time into Contention Periods (CPs) (where stations use standard EDCA) and Contention Free Periods (CFPs) (where the AP assigns the channel to stations that requested for it). During the CFP, the AP polls single stations assigning TXOPs, which can be used by the stations to send multiple frames in a burst.

In both cases, the standard allows stations to perform bursting only in managed mode, i.e., when associated to an AP. Indeed, it is a duty of the AP to inform associated stations whether bursting is enabled and, in case, how long the TXOP is. Thus, strictly sticking to the standard, frame bursting cannot be used in vehicular networks, where there is no AP to coordinate the access. Figure 1a shows a graphical representation of a frame burst as per 802.11e standard.

2.2. Distributed EDCA Bursting (DEB)

The idea we propose is non-standard compliant, but perfectly fits clustering scenarios. Take as an example a platooning application [17]. Such an application autonomously drives a group of vehicles (i.e., a platoon) in a “road-train” configuration. To maintain the desired gap between them,

vehicles within the same platoon share position, speed, and acceleration by means of periodic broadcast beacons. In this particular configuration, the leader can act as a cluster head: when gaining access to the channel, it sends a beacon that carries application information, but also reserves the channel for the amount of time required to send all the beacons of its followers by setting the Network Allocation Vector (NAV) of the MAC frame. Inside the beacon the cluster head includes, together with application layer data, the identifiers of the vehicles that should participate in the bursting procedure during the reserved TXOP: this implements a modified version of the CF-Poll frame of 802.11e. Upon reception of the cluster head’s beacon, the cluster members compute the identifier of the vehicle that should transmit immediately before each of them. The first vehicle in the list will schedule a transmission one Short Inter-Frame Space (SIFS) after the end of the cluster head’s frame. When receiving the frame of the first member, the second in the list does the same, and the process continues until the last vehicle in the list. Figure 1b shows the working principle. Each vehicle properly sets the NAV to protect the remaining part of the reserved TXOP: this mechanism is useful to inform other platooning leaders if they failed to decode one or more beacons in the burst. In the case of a failed reception (e.g., the third vehicle is unable to decode second vehicle’s frame), the distributed bursting stops.

Alternatively, cluster members can schedule their transmission directly after decoding cluster head’s frame. This avoids the protocol to stall in case of losses. We refer to this mechanism as “pre-scheduling”. In this work, we analyze the behavior of both protocol versions (Sections 4 and 5).

Distributed EDCA bursting can reduce protocol overhead times caused by the backoff procedure, but it has some specific requirements. First of all, it requires a cross-layer approach: the application and the data-link layer must share information. The application is in charge of forming the cluster and decide which node is the cluster head. This information must be shared with the MAC. Moreover, the MAC needs to fetch data from the application in an “on-demand” fashion. When a node receives the frame of the previous vehicle in the list, it knows that it will need to send a beacon within a SIFS. The MAC queue, however, might either be empty or include an outdated packet. Thus a more effective, cross-layer technique would be to have the MAC fetch the most recent information just before sending the beacon.

3. Evaluating Potential Resource Gain

In this section we estimate the potential resource gain when using DEB. We consider a saturated channel with no collisions, i.e., there is always a frame being sent in the channel and only one station winning the contention. Under these conditions, we compute the channel utilization for DEB with respect to standard DCF. The result clearly depends on several parameters, which include cluster size, application layer payload, Access Category (AC), and physical layer bit rate.

In IEEE 802.11 we have four ACs: each of those has a dedicated MAC queue and customized MAC layer parameters. The aim is to assign different priorities to frames. In particular, the standard defines the background, best-effort, video, and voice ACs (AC_BK, AC_BE, AC_VI, and AC_VO, respectively). The AC-specific parameters include different minimum and maximum contention window sizes, as well as a different idle waiting time (Arbitration Interframe Space (AIFS)), which clearly has an impact on protocol overhead.

Following the IEEE 802.11 standard [18], for a given MPDU and physical layer bit rate, we can compute frame duration as

$$T_{\text{frame}} = T_{\text{preamble}} + T_{\text{signal}} + T_{\text{sym}} \cdot N_{\text{sym}} \quad (1)$$

where T_{preamble} , T_{signal} , and T_{sym} represent the duration of the preamble, of the physical header, and of a single OFDM symbol, respectively. N_{sym} is the number of OFDM symbols needed to encode the MPDU, which is computed as

$$N_{\text{sym}} = \left\lceil \frac{16 + 8 \cdot \text{MPDU} + 6}{N_{\text{DBPS}}} \right\rceil \quad (2)$$

with N_{DBPS} being the number of data bits per OFDM symbol, which clearly depends on the chosen modulation and coding scheme.

For standard DCF, we can define the fraction of effectively utilized channel time as

$$u_{\text{DCF}} = \frac{T_{\text{frame}}}{T_{\text{AIFS}} + T_{\text{backoff}} + T_{\text{frame}}} \quad (3)$$

with T_{AIFS} and T_{backoff} being the overhead times of the channel access mechanism of IEEE 802.11 (AIFS and backoff), which depend on the chosen AC. The backoff is computed as the average backoff time, i.e., one half of the contention window multiplied by the slot time. Notice that we are considering broadcast frames, so the contention window never grows.

For DEB, instead, the fraction of effective channel utilization is computed as

$$u_{\text{DEB}} = \frac{N_c \cdot T_{\text{frame}}}{T_{\text{AIFS}} + T_{\text{backoff}} + N_c \cdot T_{\text{frame}} + (N_c - 1) \cdot T_{\text{SIFS}}}, \quad (4)$$

where N_c is the cluster size. For $N_c = 1$, $u_{\text{DCF}} = u_{\text{DEB}}$.

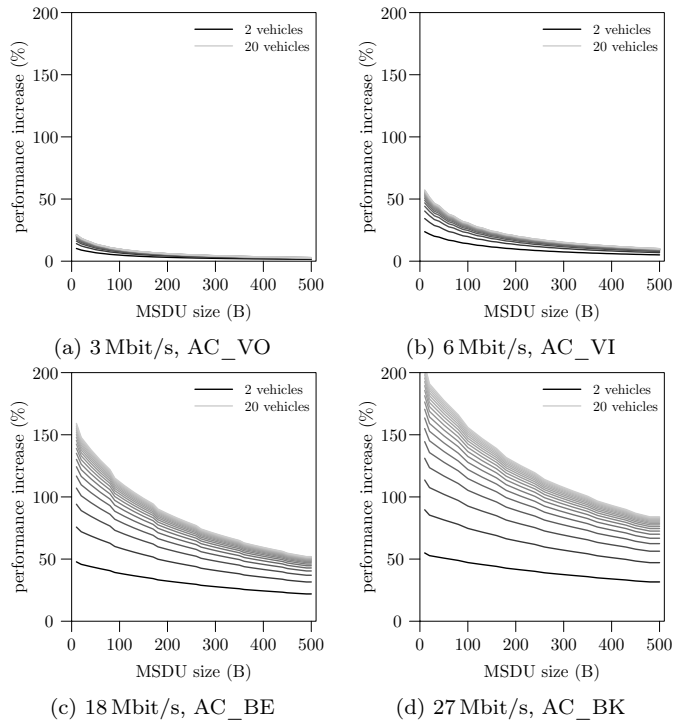


Figure 2: Performance increase ρ_{DEB} as function of payload, for different cluster sizes.

Finally, we can compute the increase in performance ρ_{DEB} of DEB over standard DCF as

$$\rho_{\text{DEB}} = \frac{u_{\text{DEB}} - u_{\text{DCF}}}{u_{\text{DCF}}}. \quad (5)$$

We compute ρ_{DEB} for different MAC Service Data Unit (MSDU) and cluster sizes (10 B to 500 B and 2 to 20, respectively), and for all ACs and physical layer bit rates.

Figure 2 plots ρ_{DEB} as function of the MSDU size, for different cluster sizes and a subset of bit rates and ACs. The main takeaway is that the benefits of DEB are larger where the communication is dominated by protocol overhead, i.e., high bit rates, small MSDUs, and low-priority EDCA queues. Figure 2a shows the performance for the worst-case scenario, i.e., the lowest bit rate and the AC with the highest priority. The increase in performance with respect to standard DCF is marginal, and decreases with the MSDU size. Moreover, the impact of the cluster size is barely noticeable. In contrast, Figure 2d shows ρ_{DEB} for the best-case scenario, i.e., the highest bit rate and the lowest priority queue. In this situation, a cluster size as low as 2 vehicles already provides a performance gain around 50%. For large clusters (in this specific case, 20 vehicles), the performance gain is between 200% and 100% depending on the MSDU size, meaning that we are tripling the performance for small MSDUs and doubling it for larger ones.

Figure 3 shows ρ_{DEB} from another perspective, i.e., as function of the cluster size N_c for different MSDU sizes.




Figure 3: Performance increase ρ_{DEB} as function of cluster size, for different MSDUs.




Figure 4: Screen shot of the simulation scenario.

The figures show that the largest benefit is obtained by using a large cluster size, as expected. The performance gain, however, does not grow indefinitely, but it has a horizontal asymptote which can be computed as

$$\lim_{N_c \rightarrow \infty} \rho_{\text{deb}} = \frac{T_{\text{AIFS}} + T_{\text{backoff}} - T_{\text{SIFS}}}{T_{\text{frame}} + T_{\text{SIFS}}}. \quad (6)$$

4. Performance Evaluation for a Platooning Application

We test the performance of the proposed bursting mechanism (referred to as DEB) by means of simulations. We implement the protocol as an extension of the standard MAC layer included in the Veins framework [24], and we compare it against standard access mechanism (referred to as DCF). As anticipated, we consider a platooning scenario implemented in PLEXE [25], thus we use the words platoon leader and cluster head, as well as platoon follower and cluster member, interchangeably. The choice of this application relies on the fact that it forms stable clusters by design. We assume clusters to be already formed, as in this

Table 1: Network and road traffic simulation parameters.

	Parameter	Value
communication	Path loss model	Free space ($\alpha = 2.0$)
	Frequency	5.89 GHz
	Bit Rate	6 Mbit/s and 18 Mbit/s
	Transmit power	20 dBm and 0 dBm
	CCA threshold	-65 dBm
	Noise floor	-95 dBm
	Minimum sensitivity	-94 dBm
	PHY model	IEEE 802.11p
	MAC model	1609.4 single channel (CCH)
	Access category	AC_VI, AC_BK
	MSDU size	200 B
	Beacon rate	10 Hz
mobility	Speed	100 km/h
	Cluster size (N_c)	8 cars
	Car's length	4 m
	Number of cars	64, 128, 192, 256, 320, 384, 448, 512, 576, 640
	Intra-cluster distance	5 m
	Inter-cluster distance	42 m

work we want to highlight the benefits and the drawbacks of DEB without considering the problem of cluster formation and stability. This is clearly a simplifying assumption, but it is needed for the detailed analysis of results and phenomena, which requires to go down to frame level.

For both protocol versions, we consider a beacon rate of 10 Hz, a standard value considered for a platooning application [26]. When using standard DCF, all vehicles send 10 beacons per second. With DEB, instead, only the cluster heads schedule period beacons, while the members send their beacons according to the bursting scheme. The scenario reproduces a 4-lane freeway where 8-car platoons travel with a constant speed of 100 km/h. We consider a total number of cars going from 64 to 640 to investigate the behavior of the protocol under different network loads. Moreover, we consider two different transmit power settings for both DCF and DEB. In the first one (No TXPC) all vehicles use the same transmit power (20 dBm). In the second (TXPC) only the cluster heads (i.e., the leaders) use a transmit power of 20 dBm, while cluster members (i.e., the followers) use a reduced transmit power of 0 dBm. The following vehicles maintain a gap of 5 m to the one in front by using the California PATH's Cooperative Adaptive Cruise Control (CACC) [27], while leaders maintain an inter-platoon headway time of 1.5 s using a standard Adaptive Cruise Control (ACC). Figure 4 shows a screen shot of the simulation scenario, while Table 1 lists other specific simulation parameters. To obtain a higher statistical confidence, we repeat each simulation setup 10 times.

4.1. General Networking Perspective

We begin the analysis by considering a general networking perspective, in particular by observing the channel busy ratio and the experienced frame collisions. The channel busy ratio indicates the amount of time the PHY layer senses the channel as busy, only in terms of energy level. In the simulation, each vehicle samples its own channel busy ratio over one second intervals. Veins signals a collision on

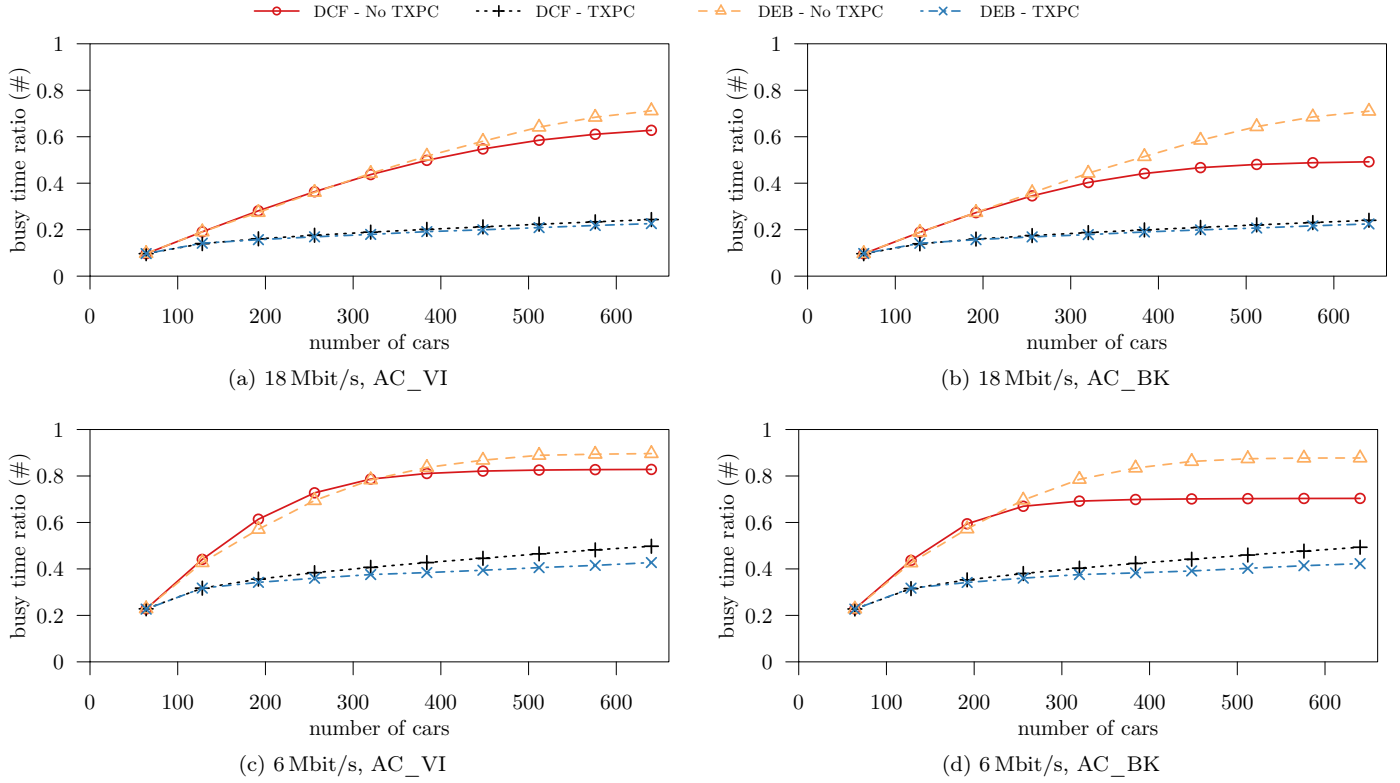


Figure 5: Comparison of DEB against standard DCF in terms of experienced channel load.

a frame when this can not be decoded due to interference. Each vehicle logs the number of frame collisions once per second.

Figure 5 compares DEB and DCF in terms of channel busy ratio, for the different bit rates and ACs considered. Each data point represents the average busy ratio over all cars and all repetitions for the four considered approaches. This metric describes the effective channel utilization but it needs to be carefully interpreted, as a high channel load does not necessarily indicate a negative performance: on the contrary, a higher channel busy ratio might indicate a better resource utilization.

Figure 5a compares DEB and DCF for 18 Mbit/s and the video access category. First we can see that, for the majority of the cases, the two approaches lead to the same channel busy ratio. The reason is that, given the high bit rate and the low priority queue, the channel is not overloaded. This means that the channel is capable of handling all the frames being sent. As soon as the channel starts being overloaded, channel busy ratio for the two approaches diverge. DEB with no transmit power control results in a higher busy ratio with respect to DCF because, by reducing overhead times, it raises the maximum achievable load.

For a lower priority AC (Figure 5b) this phenomenon is even more evident. The large overhead times caused by long AIFS and backoff procedure makes it impossible for DCF to reach a busy ratio larger than 50%, while there is no notable difference for DEB when changing the access

category.

The results for 6 Mbit/s (Figures 5c and 5d) and no transmit power control are qualitatively similar, but the channel busy ratio is much higher due to longer frame duration. Again, by reducing overhead times, the effective channel utilization of DEB is higher, overcoming the limits of standard DCF.

When enabling transmit power control, however, the behavior is inverted, i.e., the busy ratio for DEB is lower than for DCF. This can be explained by analyzing the simulations at frame-level detail. Figure 6 shows a subset of events taken from a small simulation developed to identify protocol behavior in detail. The simulation includes two 5-car clusters, one cluster having even vehicle ids (0 to 8) and the other having odd vehicle ids (1 to 9), which are at a distance of roughly 1 km. The four subplots show the physical layer events for the two cluster heads and the two second vehicles. Each box represents a frame, and the color indicates the kind of event (blue, green, and red for a transmitted, received, or lost frame, respectively). The figure shows the physical layer events for vehicles 0, 1, 2, and 3, i.e., the two cluster heads and the two closest cluster members.

In this snapshot, the two cluster heads try to access the channel roughly at the same moment, with vehicle 1 being the winner of the channel contention. The frame is received by the cluster members which schedule their transmission in order, one after the other. Cluster head

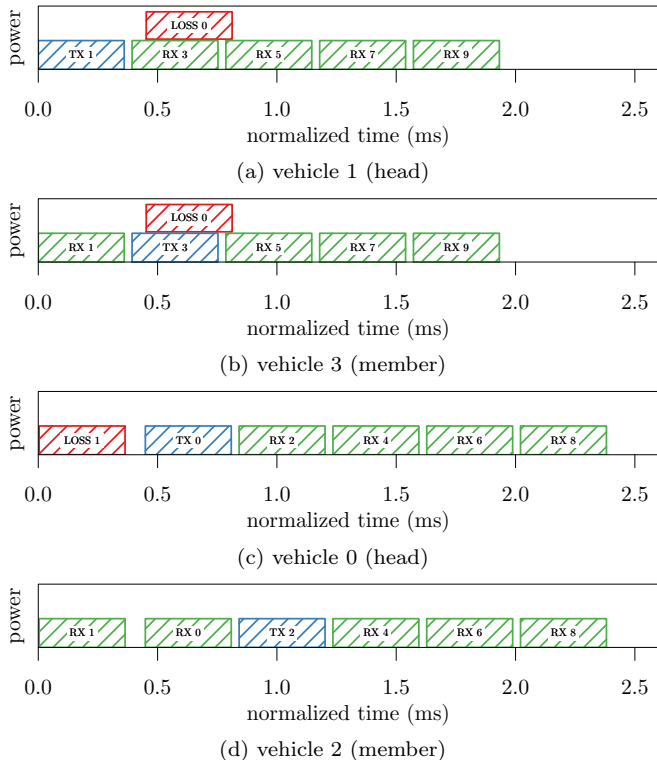


Figure 6: Burst overlap causing lower busy channel ratios when using transmit power control. Blue, green, and red boxes represent transmitted, received, and lost frames, respectively. The height of the frame does not represent actual power.

0, however, is unable to decode the frame from the other cluster head, thus it does not set its NAV and, after waiting for one AIFS and performing the backoff, starts to transmit because it has no indication about the channel being currently busy. Moreover, it is not able to perform carrier sensing upon frames of the other cluster members, as they are using a reduced transmit power. DEB thus starts for the “even” cluster as well, and the members successfully complete the burst without interfering with the other cluster. If we consider the perspective of another vehicle driving between the two clusters, its perceived channel busy ratio will be lower than when using DCF, because the two bursts are almost completely overlapped. With standard DCF, such a situation is much less likely to occur. DEB actively coordinates in-cluster transmission, so every time two cluster-heads initiate a burst phase, the two bursts will be mostly overlapped. With DCF, instead, given that each frame is scheduled independently of the others, a complete overlap for all frames is very unlikely to occur. For this reason, DCF causes a higher channel busy ratio with respect to DEB. In this regard, DEB can improve spatial reuse.

Figure 7 shows the number of collisions per seconds, averaged over all vehicles and repetitions. The plots show that, in all situations, the channel reservation mechanism of DEB avoids random channel contention, thus reduc-

ing the overall amount of collided frames. By taking a closer look to the case in which transmit power control is disabled, we obtain the largest advantage when using a high physical layer bit rate, as predicted by the analysis in Section 3. Given the lower frame duration (and thus the higher amount of available resources), distributing channel utilization with DEB drastically reduces frame collisions with respect to pure random contention. When decreasing the bit rate, the longer frame duration results in lower resource availability, thus the reduction in the number of collisions is less pronounced, but still fairly large.

When employing transmit power control, instead, the difference is larger for the lower bit rate, but this is simply due to the fact that, when using a high bit rate and a low transmit power, the channel load is minimal and the number of collisions close to zero for both DCF and DEB. In addition, we have a lower amount of collisions because, by employing a lower transmit power, each vehicle will “sense” (and thus try to decode) a smaller amount of frames.

4.2. Intra-cluster Performance

The analysis in Section 4.1 only looks at the performance from a generic network perspective, without considering cluster existence at all. In this section we perform an intra-cluster analysis considering inter-arrival times for beacons within each cluster. The sample application we consider here (i.e., a CACC) requires the data sent by the cluster head and the cluster member directly in front (leader- and predecessor-following platooning control). During the simulation, vehicles record the inter-arrival time for each beacon received from the platoon leader and the vehicle in front. We then compute the average inter-arrival time and the standard deviation for each experiment. More formally, let \mathcal{A} be the set of all inter-message arrival times recorded by a vehicle. We compute the weighted average inter-arrival time as

$$\mu = \frac{\sum_{a \in \mathcal{A}} a^2}{\sum_{a \in \mathcal{A}} a}. \quad (7)$$

Equation (7) uses the inter-arrivals as weights as, for example, an inter-arrival of 200 ms must count twice as much of a 100 ms one. A standard average would treat a long and a short inter-arrival the same. In practice, in a time span of 10 s, we can record 100 0.1 ms time samples or a single 10 s sample, and the latter would be considered statistically irrelevant. In fact, its relevance is the same as it spans over the same time period, so we need to weight each sample by its length. Similarly, we compute the weighted standard deviation as

$$\sigma = \sqrt{\frac{\sum_{a \in \mathcal{A}} a \cdot (a - \mu)^2}{\sum_{a \in \mathcal{A}} a}}. \quad (8)$$

In Figure 8 we plot μ and σ for leader packets, for the two bit rates we consider and the AC_BK access category. Each data point is the average μ and σ over all vehicles

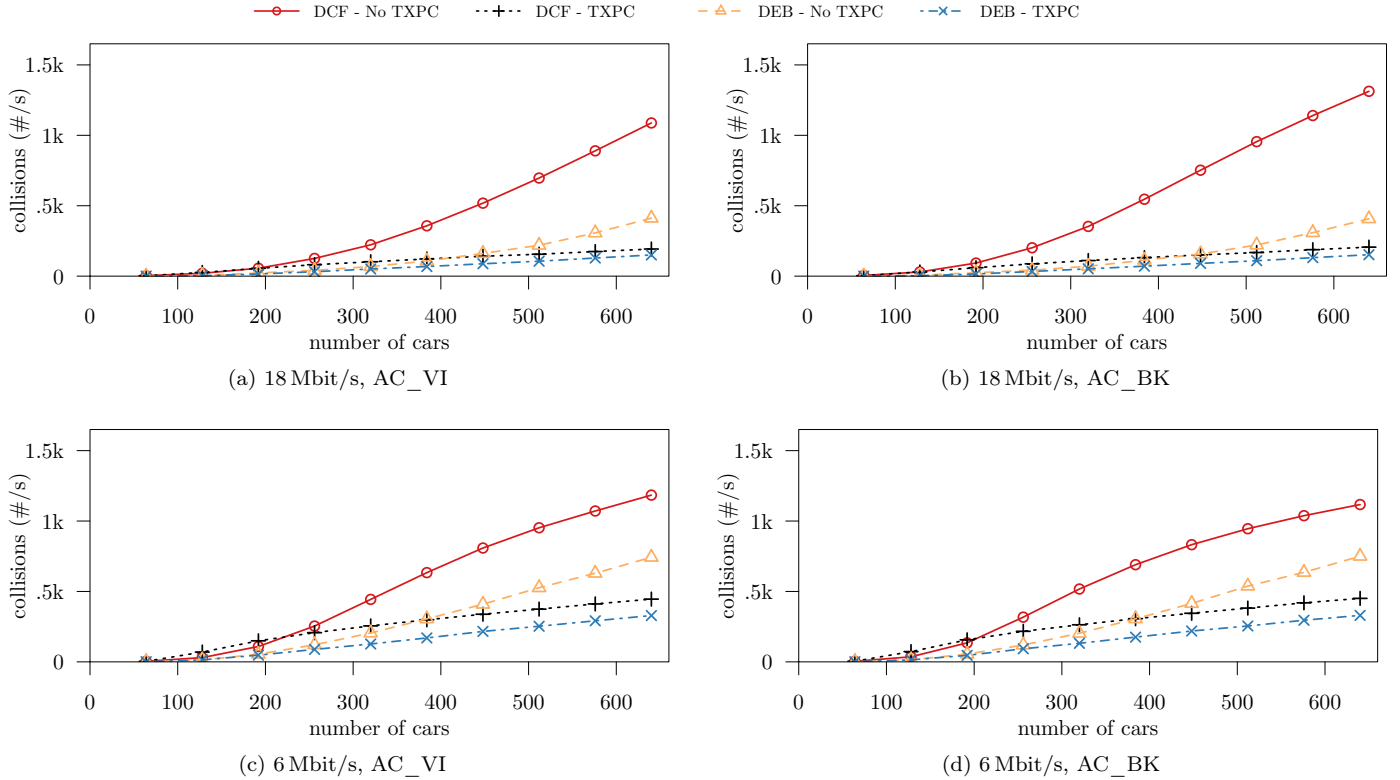


Figure 7: Comparison of DEB against standard DCF in terms of experienced frame collisions.

and simulation repetition. First we can observe that DEB does not always result in the lowest possible inter-arrival. Reserving the channel for a burst indeed adds additional channel access delay. A single frame for an MSDU of 200 B at 6 Mbit/s lasts roughly 360 μ s, while an 8-car burst for the AC_BK access category lasts around 3.4 ms, i.e., almost one order of magnitude longer. Deferring transmission for multiple bursts can thus easily lead to larger inter-arrival times. This is also shown by the larger standard-deviation for DEB. In addition to the defer-induced delay, collisions can also affect the inter-arrival time as they lead to frame losses.

For a bit rate of 18 Mbit/s the average inter-arrival is much smaller, as the burst duration is shorter. The protocol with the smallest average inter-arrival is DCF with transmit power control. In [28, 6] we indeed show that tuning the transmit power to specific application requirements is extremely beneficial but, as shown in Figure 7, the random contention can increase the number of collisions. In addition, DCF with transmit power control reaches a good performance because of the particular scenario we consider, where vehicles in the same cluster are very close each other. Having very localized communication increases frame reception probability even in the case of strong interference. As shown by the collisions analysis in Figure 7, for a larger communication domain DEB results in better performance.

Figure 9 shows μ and σ for front-vehicle packets aver-

aged over all cars and simulation repetitions. The major difference is the reduction in inter-arrival time for DCF with no transmit power control, which is due to the smaller inter-vehicle distance. DEB performance is similar to the one shown in Figure 8 as, given its bursting mechanism, front packet inter-arrivals are correlated to the ones of the leader. Another phenomenon is the large σ for DEB when using transmit power control in Figure 9b, as well as a larger μ than when transmit power control is not enabled. DEB indeed suffers a “burst interruption” weakness when using transmit power control. Figure 10 depicts a sketch of the problem we discovered while analyzing DEB behavior in detail.

Imagine to have a cluster made of vehicles with identifiers going from 1 to 8 (1 being the cluster head), and a farther cluster head with id, for example, 20. The figure shows the channel-level perspective for vehicles 4 and 5. If cluster head 20 misses the frame of cluster head 1, the former might try to access the channel at any time because, given that cluster members are using transmit power control, it is not able to perform carrier sensing (as for the case in Figure 6). In this particular example, vehicle 20 starts to transmit during the SIFS between the frames of vehicles 3 and 4. Vehicle 4 transmits its frame in any case, as it schedules its transmission after the successful reception of frame 3. As frame 20 comes first, however, vehicle 5 can synchronize to that frame and try to receive it, blocking the reception of frame 4. This interrupts the burst for vehicles

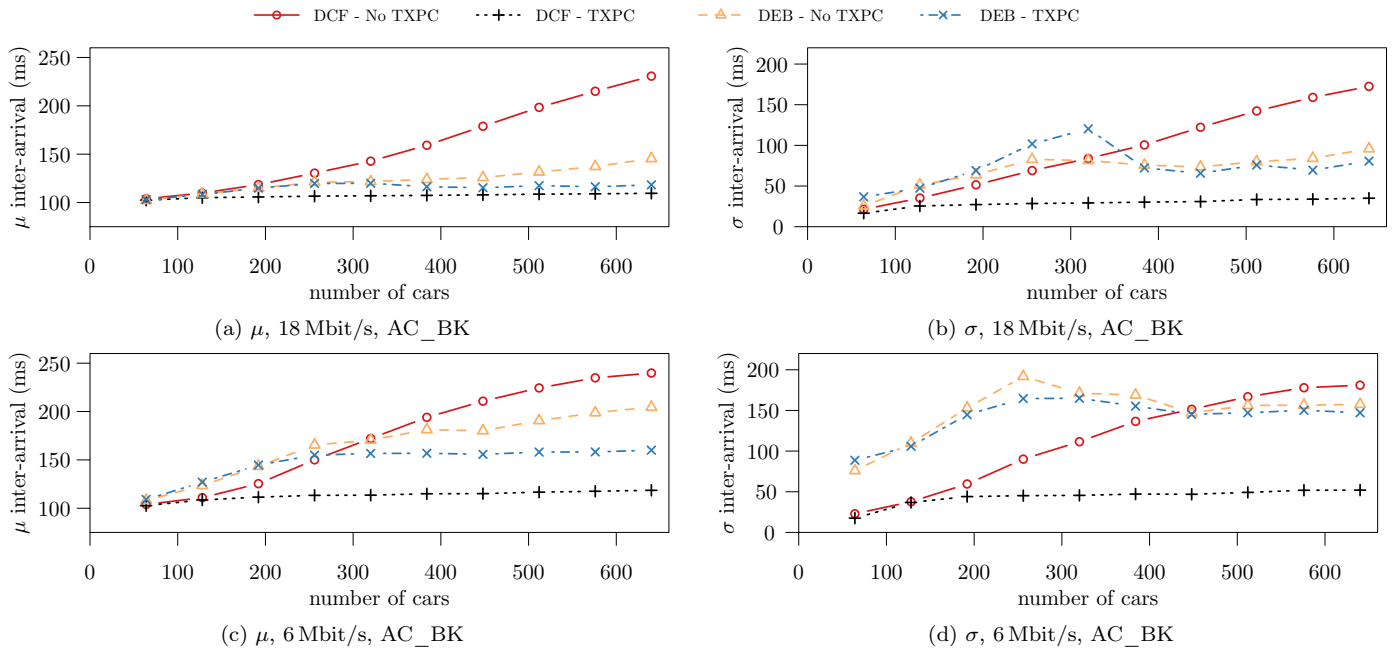


Figure 8: Weighted mean and standard deviation for leader packets' inter-arrival time.

5, 6, 7, and 8, so for vehicles which are at the tail of the polling list it is highly probable not to receive a frame.

This weakness can lead to a tremendous performance degradation and high inefficiency, as the cluster head reserves the channel for the duration of a burst, but a large portion of the reserved transmission time might be unused due to the interruption. The solution is to employ a pre-scheduling mechanism, i.e., each cluster member schedules its transmission synchronizing directly onto cluster head's frame instead of waiting for the previous cluster member's frame. We further see some peaks around the density of 300 cars. We were able to reproduce these results even for a very large number of simulations. Obviously, these peaks result from the specific simulation setup and configuration (transmit power, number of cars, interference range). Next section analyzes the performance of this mechanism.

5. Enabling Burst Pre-scheduling

In this section, we analyze the performance of DEB with the addition of pre-scheduling. As briefly anticipated, when this feature is enabled, each cluster member schedules its transmission when receiving cluster head's frame, instead of waiting the frame of the previous vehicle in the poll list. This avoids the complete interruption of a burst in case of a lost frame.

Figure 11 shows the weighted average and standard deviation (as per Equations (7) and (8)) for front packets and a bit rate of 18 Mbit/s. We omit the performance for leaders' packets as pre-scheduling has nearly no impact on them, as expected. First, the results highlight that pre-scheduling gives no benefit when transmit power is not

employed. This is due to the fact that cluster members' frames are "protected" by their higher transmit power, i.e., their reception probability is higher and they trigger carrier sensing on farther vehicles, avoiding the situation depicted in Figure 10.

Pre-scheduling is instead beneficial for low-power frames. This is clearly witnessed first, by the reduction in average inter-arrival time and second, by the lower standard deviation. In particular, both the average and the standard deviation are comparable to the results with no transmit power control. However, the average inter-arrival time when transmit power control is disabled starts to rise for the scenarios with the largest number of vehicles. Finally, the impact for a lower bit rate (i.e., 6 Mbit/s, not shown here) is less relevant, as such modulation is much more robust to noise and interference. In conclusion, pre-scheduling is fundamental for "fragile" frames, i.e., the ones transmitted using a lower power and a higher bit rate. Pre-scheduling does not incur into additional drawbacks with respect to the standard scheduling, so this mechanism should always be enabled, regardless of the transmit power or the bit rate.

6. Performance under External Interference

In this section we relax the dedicated-channel assumption. We run an additional set of simulations where two lanes are occupied by clusters of platooning vehicles, while the other two are filled by human-driven vehicles sending 10 beacons per second. Cluster vehicles are configured in groups of eight as in the other experiments. This time the total amount of cluster vehicles ranges from 32 to 320 as we only occupy two lanes. Human-driven vehicles

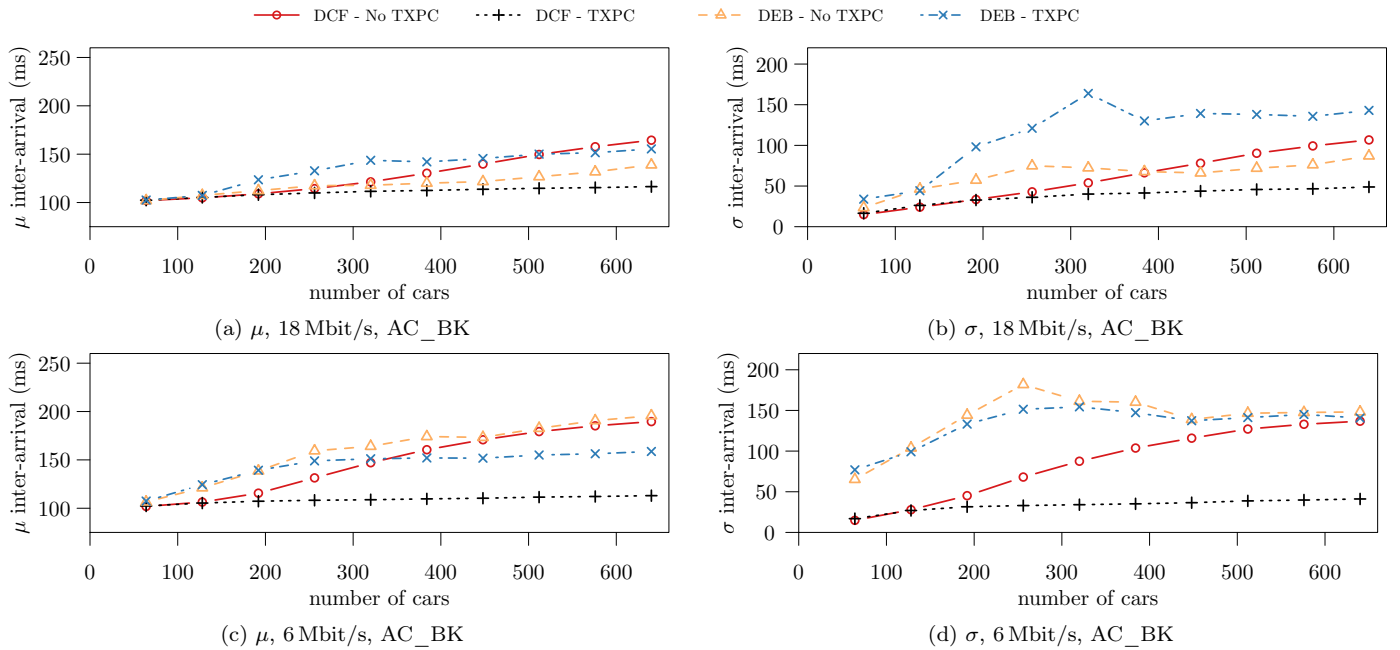


Figure 9: Weighted mean and standard deviation for front packets' inter-arrival time.

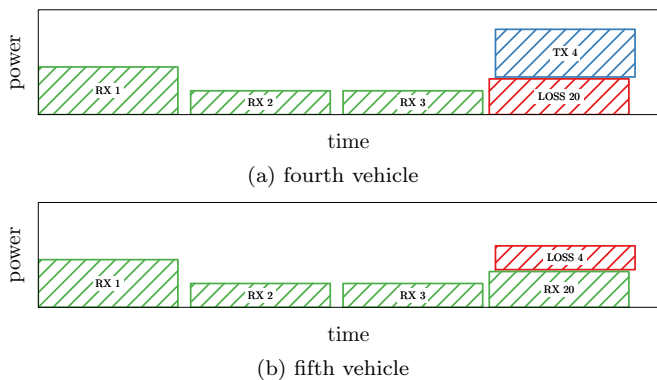


Figure 10: Distributed EDCA bursting problem in case of packet losses. Channel perspective for fourth and the fifth vehicle of cluster. Blue, green, and red boxes represent transmitted, received, and lost frames, respectively. The height of the frames gives an indication of received power but does not represent actual values.

are not grouped and thus do not use DEB, but they respect the NAV allocated by cluster heads. The number of human-driven vehicle is variable, as we inject them in the simulation to cover the same highway portion occupied by the platoons. Clearly we have less human-driven vehicles than platooning ones, as their inter-vehicle distance is larger.

Given the different setup with respect to the previous experiments, the results would not be directly comparable. For this reason, we run the same scenarios using DCF only as well, comparing the performance against DEB.

Figure 12 shows the results for a 18 Mbit/s bit rate

and the AC_BK access category with the pre-scheduling mechanism enabled. We plot channel busy ratio and collisions, as well as μ and σ (Equations (7) and (8)) for both leader and front vehicle packets. The plots only include the results for clustering vehicles.

Concerning channel busy ratio and collisions, the results are coherent with the experiments where only clustering vehicles are present (Figures 5b and 7b). The busy ratio and the number of collisions decrease with respect to previous experiments as the number of vehicles and the overall density is lower, but the result are qualitatively comparable, meaning that DEB works even in case of external interference.

The same holds for the μ and σ statistics, and the results are comparable with the ones shown in Figures 8a, 8b, 9a and 9b. The main difference is a degradation in performance for standard DCF when using transmit power control. Human-driven vehicles indeed transmit their frames at full power (20 dBm), so the frames sent by platoon followers are corrupted by the interference. DEB, instead, by reserving the channel and setting the NAV protects low-power frames, while still improving spatial reuse.

The results in this final section thus show that DEB is robust to external interference, provided that interfering vehicles respect the NAV set by cluster head in their frames. This, however, does not require any change to the standard, as the NAV is present in IEEE 802.11 since its first version.

7. Conclusion and Discussion

In this paper, we presented a distributed EDCA bursting mechanism to improve cluster-based communication in

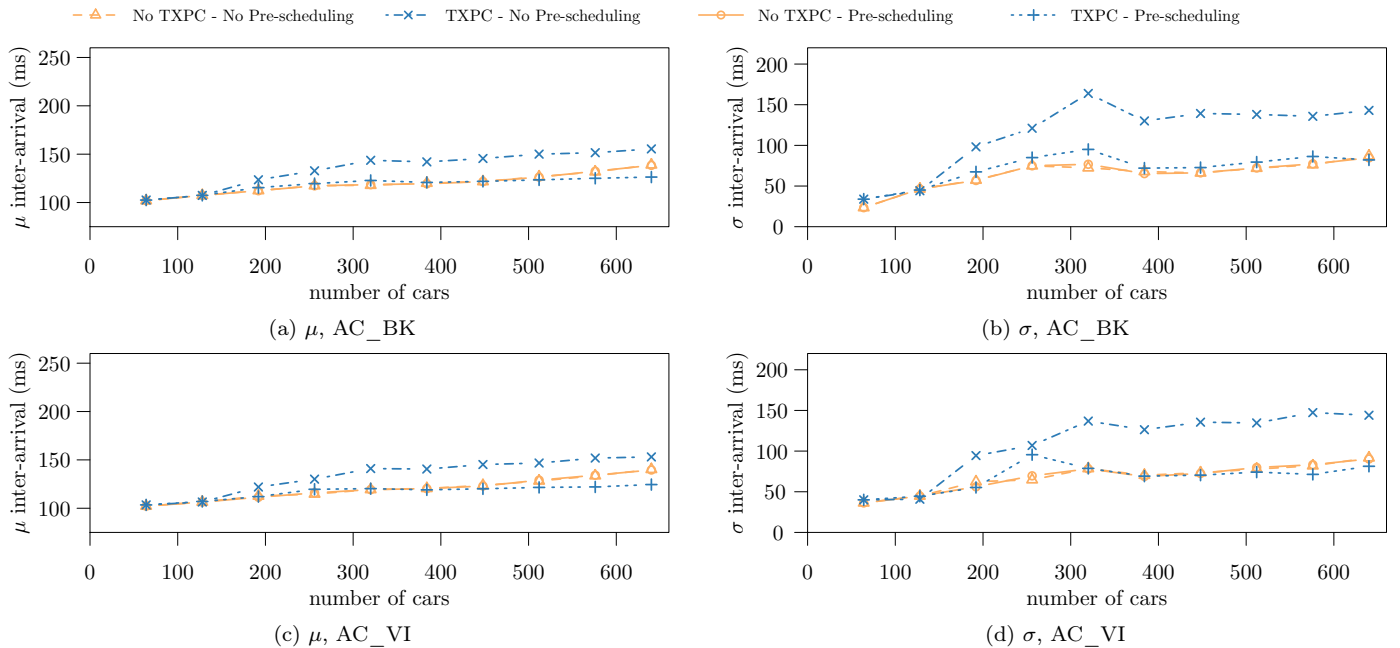


Figure 11: Weighted mean and standard deviation for front packets’ inter-arrival time when using pre-scheduling, 18 Mbit/s.

IVC. The idea is to modify the standard 802.11e bursting mechanism, which usually works for a unicast communication between a station and the AP, and have a cluster head send a beacon to reserve the channel for the duration of a TXOP. The cluster head, with its beacon, polls cluster members, which send their data one after the other in a burst, i.e., having each of their frames separated by a SIFS.

We made a first implementation of the proposed mechanism in PLEXE and tested its performance against the standard DCF-controlled channel access mechanism. We compared the two approaches from a generic and an intra-cluster perspective. In the generic perspective we analyzed the performance in terms of channel busy ratio and experienced collisions. Concerning busy ratio, we have shown that the performance of DEB depends on the chosen transmit power policy. When all vehicles use the same transmit power, DEB is capable of overcoming channel limits due to overhead times. When cluster members use a reduced transmit power, instead, DEB reduces the actual utilization by improving spatial channel reuse. In terms of collisions, instead, DEB has superior performance with respect to standard DCF, independently of the transmit power policy.

With respect to the intra-cluster perspective we analyzed the weighted average and standard deviation of beacon inter-arrival times. We have shown that DEB performance depends on the actual channel load. The reservation mechanism of DEB causes additional delay, and the higher the load, the higher the delay. For example, average inter-arrival for DEB is larger for a 6 Mbit/s bit rate, as each 8-car burst lasts roughly 3.4 ms. For a higher bit rate (i.e., 18 Mbit/s) the experienced average inter-arrival time

is much smaller.

Finally we have shown that, despite being beneficial for the “network as a whole”, the use of transmit power control is potentially harmful for intra-cluster performance, as the loss of a frame can stop the bursting mechanism. We can prevent this problem by changing the scheduling policy, i.e., scheduling beacon transmission directly on cluster head’s beacon rather than waiting for the frame of the previous vehicle in the burst list.

In conclusion, the analysis shows that DEB definitely has the potential of improving cluster communication in vehicular environments. This is also witnessed by the performance of DEB under external interference. To show this potential and to describe the phenomena in detail, we had to focus our analysis on a simplified scenario and on a single application, and we believe there are other interesting research questions to be investigated.

First of all, cluster management will have a great impact on DEB. Part of the communication overhead will be dedicated to the creation and the maintenance of the clusters. Under which conditions is DEB still beneficial? How does cluster stability affect the performance?

Second, the dynamics of the vehicles can also play a major role. In our analysis all vehicles travel in the same direction, but we might have clusters crossing each-other, for example at an intersection or simply in the reverse driving direction.

Clusters will interfere each other depending on the amount of time they stay within interference range. Finally, there is an issue related to the duration of the TXOP. In our setup, a cluster head could potentially reserve the channel

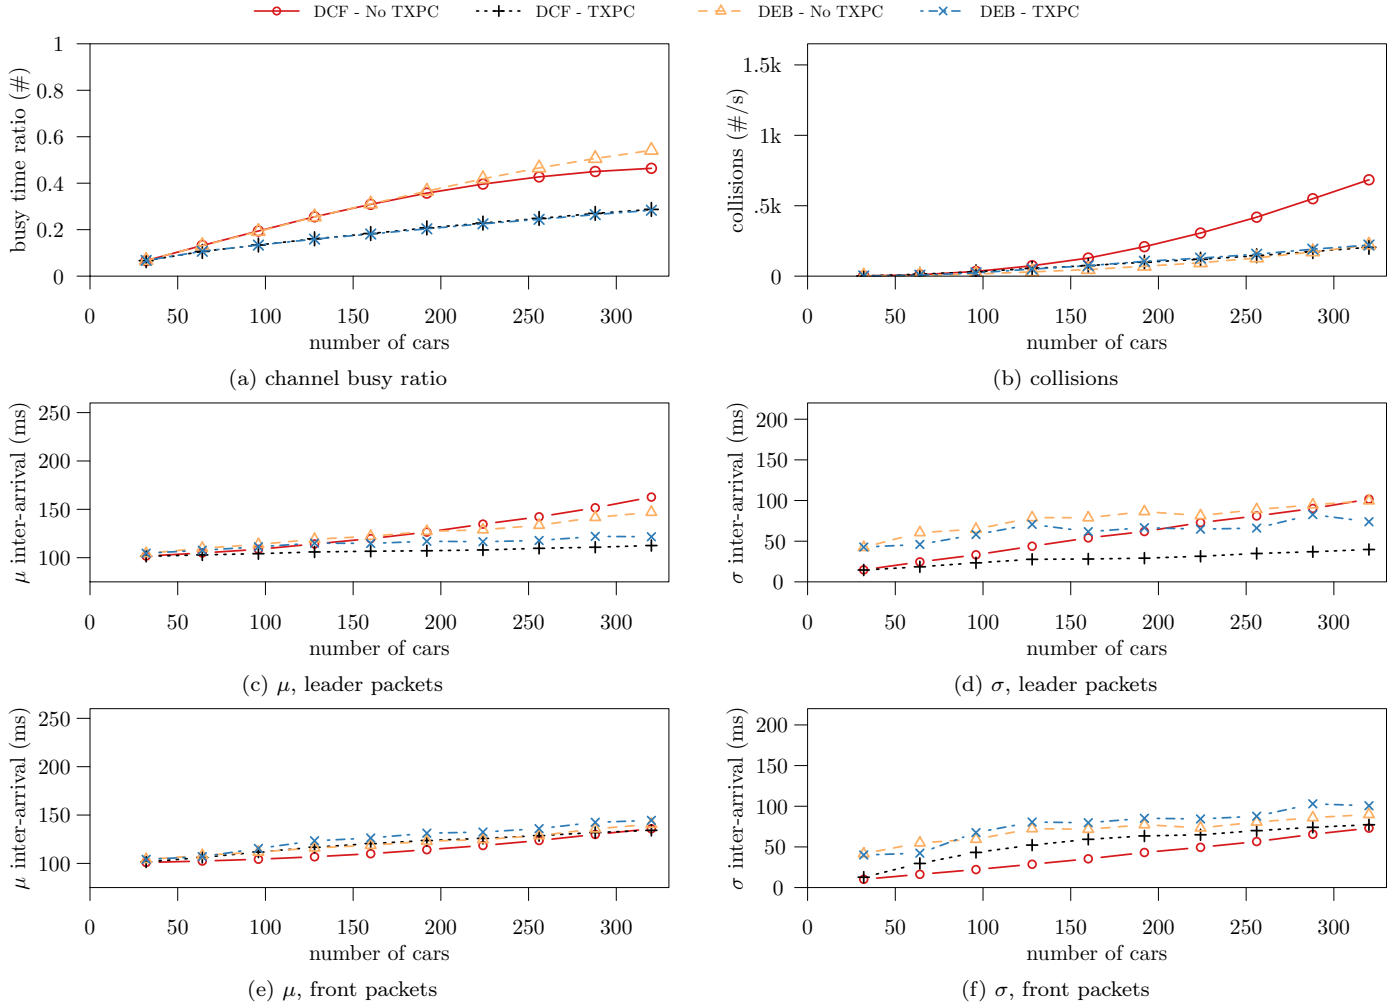


Figure 12: Performance evaluation for DEB under interference. Results computed for the scenario using 18 Mbit/s, AC_BK, with pre-scheduling.

for an indefinite amount of time, causing unacceptable access delays to other clusters. The IEEE 802.11e standard defines the maximum TXOP duration that each station is allowed to use. In this regard, a similar approach should be taken by DEB as well, so we need to study the impact of different maximum TXOP lengths on the performance. On the one hand, a short TXOP reduces access delays but limits the benefits of DEB. On the other hand, a long TXOP maximizes the benefits while increasing channel access time. We believe this work can thus foster further research on this promising cluster-based communication scheme.

References

- [1] National Highway Traffic Safety Administration, Vehicle Safety Communications Project Task 3 Final Report: Identify Intelligent Vehicle Safety Applications Enabled by DSRC, Project Report DOT HS 809 859, NHTSA (Mar. 2005).
- [2] C. Sommer, F. Dressler, Vehicular Networking, Cambridge University Press, 2014. doi:10.1017/CBO9781107110649.
- [3] F. Dressler, H. Hartenstein, O. Altintas, O. K. Tonguz, Inter-Vehicle Communication - Quo Vadis, IEEE Communications Magazine 52 (6) (2014) 170–177. doi:10.1109/MCOM.2014.6829960.
- [4] M. Segata, R. Lo Cigno, Automatic Emergency Braking - Realistic Analysis of Car Dynamics and Network Performance, IEEE Transactions on Vehicular Technology 62 (9) (2013) 4150–4161. doi:10.1109/TVT.2013.2277802.
- [5] S. Joerer, M. Segata, B. Bloessl, R. Lo Cigno, C. Sommer, F. Dressler, A Vehicular Networking Perspective on Estimating Vehicle Collision Probability at Intersections, IEEE Transactions on Vehicular Technology 63 (4) (2014) 1802–1812. doi:10.1109/TVT.2013.2287343.
- [6] M. Segata, Safe and Efficient Communication Protocols for Platooning Control, Phd thesis (dissertation), University of Innsbruck (Feb. 2016).
- [7] V. Turri, Fuel-efficient and Safe Heavy-duty Vehicle Platooning Through Look-ahead Control, Phd thesis, KTH Royal Institute of Technology (Oct. 2015).
- [8] P. Fernandes, Platooning of IVC-Enabled Autonomous Vehicles: Information and Positioning Management Algorithms, for High Traffic Capacity and Urban Mobility Improvement, Phd thesis, University of Coimbra, Portugal (Apr. 2013).
- [9] T. Neudecker, N. An, O. K. Tonguz, T. Gaugel, J. Mittag, Feasibility of Virtual Traffic Lights in Non-Line-of-Sight Envi-

- ronments, in: 9th ACM International Workshop on Vehicular InterNetworking (VANET 2012), ACM, Low Wood Bay, Lake District, UK, 2012, pp. 103–106. doi:10.1145/2307888.2307907.
- [10] M. Ferreira, R. Fernandes, H. Conceição, W. Viriyasitavat, O. K. Tonguz, Self-organized Traffic Control, in: 7th ACM International Workshop on Vehicular InterNetworking (VANET 2010), ACM, Chicago, IL, 2010, pp. 85–90. doi:10.1145/1860058.1860077.
- [11] F. Hagenauer, P. Baldemaier, F. Dressler, C. Sommer, Advanced Leader Election for Virtual Traffic Lights, ZTE Communications, Special Issue on VANET 12 (1) (2014) 11–16. doi:10.3969/j.issn.1673-5188.2014.01.002.
- [12] W. Viriyasitavat, S.-H. Yu, H.-M. Tsai, Channel Model for Visible Light Communications Using Off-the-shelf Scooter Taillight, in: 5th IEEE Vehicular Networking Conference (VNC 2013), IEEE, Boston, MA, 2013, pp. 170–173. doi:10.1109/VNC.2013.6737605.
- [13] M. Segata, R. Lo Cigno, H.-M. Tsai, F. Dressler, On Platooning Control using IEEE 802.11p in Conjunction with Visible Light Communications, in: 12th IEEE/IFIP Conference on Wireless On demand Network Systems and Services (WONS 2016), IEEE, Cortina d’Ampezzo, Italy, 2016, pp. 124–127.
- [14] G. Bansal, J. Kenney, C. Rohrs, LIMERIC: A Linear Adaptive Message Rate Algorithm for DSRC Congestion Control, IEEE Transactions on Vehicular Technology 62 (9) (2013) 4182–4197. doi:10.1109/TVT.2013.2275014.
- [15] European Telecommunications Standards Institute, Intelligent Transport Systems (ITS); Decentralized Congestion Control Mechanisms for Intelligent Transport Systems operating in the 5 GHz range; Access layer part, TS 102 687 V1.1.1, ETSI (Jul. 2011).
- [16] C. Sommer, S. Joerer, M. Segata, O. K. Tonguz, R. Lo Cigno, F. Dressler, How Shadowing Hurts Vehicular Communications and How Dynamic Beaconing Can Help, IEEE Transactions on Mobile Computing 14 (7) (2015) 1411–1421. doi:10.1109/TMC.2014.2362752.
- [17] M. Segata, F. Dressler, R. Lo Cigno, Jerk Beaconing: A Dynamic Approach to Platooning, in: 7th IEEE Vehicular Networking Conference (VNC 2015), IEEE, Kyoto, Japan, 2015, pp. 135–142. doi:10.1109/VNC.2015.7385560.
- [18] IEEE, Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, Std 802.11-2012, IEEE (2012).
- [19] Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications - Amendment 8: Medium Access Control (MAC) Quality of Service Enhancements, Std 802.11e-2005, IEEE (2005).
- [20] M. Segata, D. Goss, R. Lo Cigno, Distributed EDCA Bursting: Improving Cluster-based Communication in IVC, in: 17th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc 2016), 1st International Workshop on Internet of Vehicles and Vehicles of Internet (IoV-Vol 2016), ACM, Paderborn, Germany, 2016, pp. 13–18.
- [21] A. Memedi, F. Hagenauer, F. Dressler, C. Sommer, Cluster-based Transmit Power Control in Heterogeneous Vehicular Networks, in: 7th IEEE Vehicular Networking Conference (VNC 2015), IEEE, Kyoto, Japan, 2015, pp. 60–63. doi:10.1109/VNC.2015.7385547.
- [22] F. Li, Y. Wang, Routing in Vehicular Ad hoc Networks: A Survey, IEEE Vehicular Technology Magazine 2 (2) (2007) 12–22. doi:10.1109/MVT.2007.912927.
- [23] H. R. Arkian, R. E. Atani, A. Pourkhalili, S. Kamali, Cluster-based traffic information generalization in Vehicular Ad-hoc Networks, Elsevier Vehicular Communications 1 (4) (2014) 197–207. doi:10.1016/j.vehcom.2014.08.003.
- [24] C. Sommer, R. German, F. Dressler, Bidirectionally Coupled Network and Road Traffic Simulation for Improved IVC Analysis, IEEE Transactions on Mobile Computing 10 (1) (2011) 3–15. doi:10.1109/TMC.2010.133.
- [25] M. Segata, S. Joerer, B. Bloessl, C. Sommer, F. Dressler, R. Lo Cigno, PLEXE: A Platooning Extension for Veins, in: 6th IEEE Vehicular Networking Conference (VNC 2014), IEEE, Paderborn, Germany, 2014, pp. 53–60. doi:10.1109/VNC.2014.7013309.
- [26] J. Ploeg, B. Scheepers, E. van Nunen, N. van de Wouw, H. Nijmeijer, Design and Experimental Evaluation of Cooperative Adaptive Cruise Control, in: IEEE International Conference on Intelligent Transportation Systems (ITSC 2011), IEEE, Washington, DC, 2011, pp. 260–265. doi:10.1109/ITSC.2011.6082981.
- [27] R. Rajamani, H.-S. Tan, B. K. Law, W.-B. Zhang, Demonstration of Integrated Longitudinal and Lateral Control for the Operation of Automated Vehicles in Platoons, IEEE Transactions on Control Systems Technology 8 (4) (2000) 695–708. doi:10.1109/87.852914.
- [28] M. Segata, B. Bloessl, S. Joerer, C. Sommer, M. Gerla, R. Lo Cigno, F. Dressler, Towards Communication Strategies for Platooning: Simulative and Experimental Evaluation, IEEE Transactions on Vehicular Technology 64 (12) (2015) 5411–5423. doi:10.1109/TVT.2015.2489459.