Congestion Control for Vehicular Safety: Synchronous and Asynchronous MAC Algorithms

Sundar Subramanian
Qualcomm R&D NJRC
Bridgewater, NJ, USA
sundars@qualcomm.com

Marc Werner
Qualcomm CDMA Technologies, GmbH
Nuremberg, Germany
mwerner@qualcomm.com

Shihuan Liu
Dept. of Electrical Engg.
Iowa State University
Ames, Iowa, USA
liush@iastate.edu

Jubin Jose
Qualcomm R&D NJRC
Bridgewater, NJ, USA
jjose@qualcomm.com

Radu Lupoaie
Qualcomm CDMA Technologies,GmbH
Frankfurt, Germany
rlupoaie@qualcomm.com

Xinzhou Wu
Qualcomm R&D NJRC
Bridgewater, NJ, USA
xinzhouw@qualcomm.com

ABSTRACT

The IEEE 802.11p standard specifies the PHY and MAC layer operations for transmitting and receiving periodic broadcast messages for vehicular safety. Many studies have identified issues with the CSMA based IEEE 802.11p MAC at high densities of devices, mainly reflected by low packet reception rate. In this paper, we make an interesting observation that with increasing density, the IEEE 802.11p MAC tends towards an ALOHA-type behavior where concurrent transmissions by close-by devices are not prevented. This behavior can lead to poor packet reception rate even for vehicles in close neighborhood.

Many efforts have been made to address the IEEE 802.11p MAC issues to provide better performance for DSRC safety applications, including the introduction of Decentralized Congestion Control (DCC) algorithm to ETSI standards in Europe. In this paper, we evaluate the performance of the proposed DCC algorithm and observe that the nominal parameters in DCC are unsuitable in many scenarios. Using transmit power control as an example, we develop a simple rule within the DCC framework that can significantly improve the safety packet reception performance with increasing densities. The DCC algorithms are fully compatible with the IEEE 802.11p standards and *asynchronous* in nature.

A parallel approach to handle high device densities is a slotted *synchronous* MAC, where time is slotted based on GPS synchronization and each transmitter contends for a set of recurring time slots (or channels) with periodicity matching the required safety message periodicity. As compared to the per-packet based contention scheme as in CSMA defined in IEEE 802.11, such a scheme is much better suited for periodic safety broadcast. In this paper, we design a *standard compliant* TDM overlay on top of the MAC layer that can significantly improve the packet reception performance. Combined with a distributed resource selection protocol, the synchronous MAC can discover even more neighboring devices than the improved asynchronous approach, making DSRC safety applications more reliable.

VANET'12, June 25, 2012, Low Wood Bay, Lake District, UK.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Wireless Communication

General Terms

Algorithms; Performance

Keywords

Congestion control; Power control; DSRC; Slotted overlay

1. INTRODUCTION

Vehicle-to-vehicle (V2V) and vehicle-to-infrastrcutre (V2I) communications can greatly improve the safety and convenience of the road users [1]. The pervasive appeal of vehicular communications has led to the allocation of 70 MHz of spectrum in the 5.9 GHz band by the FCC dedicated for V2V and V2I communications in the United States. The standards for communication between vehicles are specified by the IEEE WAVE (Wireless Access in Vehicular Environments) standards [15] in United States. In particular, the physical (PHY) and the media access (MAC) layer of the WAVE standards are specified by the IEEE 802.11p standard [3]. In Europe, the V2V and V2I communication standards are specified by ETSI (European Telecommunication Standards Institute) standards [4] which are also based on the IEEE 802.11p standard. The IEEE 802.11p standard is an amendment to the existing IEEE 802.11a-2007 or Wi-Fi [2] standard with an operating bandwidth of 10 MHz, achieved by operating the IEEE 802.11a chips in a 'half-clocked' operation.

One of the main driving forces behind vehicular communications is the application of vehicular safety. It is envisioned that vehicles equipped with DSRC (Dedicated Short Range Communications) radios shall periodically transmit safety broadcasts containing their position and velocity, allowing vehicles to hear the broadcasts of the neighboring vehicles and predict their trajectories. The vehicles may then compute their own positions to identify possible collision threats.

Many studies have observed that the IEEE802.11p MAC, based on CSMA/CA, has serious issues with congestion at high device densities. Authors in [8, 10, 24, 12] have observed undesirable packet collisions and long delays between successful packet receptions at high densities. The congestion issues can have a significant impact on the effectiveness of the safety applications.

Many efforts have been made to address the congestion issue in the DSRC research community, the industry and the government. In particular, in the US, some of the OEMs have conducted independent studies of congestion control algorithms [16, 14]. The US Department of Transportation (U.S. DOT) has also partnered with the Crash Avoidance Metrics Partnership (CAMP) and Vehicle Safety Communications (VSC3) Consortium to provide an emerging standard to handle congestion related issues [5]. In Europe, the ETSI organization has standardized a framework [6] to mitigate the MAC layer congestion issues at high vehicle densities, usually referred to as the *Distributed Congestion Control* (DCC) algorithm. In both the congestion control algorithms being currently tested in US and the DCC algorithms proposed in Europe, each DSRC device adjusts its local transmission parameters, including transmit power, transmit periodicity and modulation scheme, based on the measured *channel load*. Such schemes are fully compatible with the *asynchronous* IEEE 802.11p PHY/MAC and do not take advantage of the periodicity of safety broadcasts. Meanwhile, to the best of our knowledge, the proposed DCC algorithm was not fully tested and evaluated in the literature.

As an alternative to the inherently asynchronous CSMA mechanism, the approach proposed by [9] is to operate a medium access in a *synchronous* manner based on a self-organizing time-division multiplexing approach (STDMA) where fixed periodic time-slots are allocated and vehicles choose a resource that is least occupied. Such a resource allocation algorithm is shown to have a more graceful degradation in performance with increasing device densities. However, their approach is not compliant with the existing IEEE 802.11p standard, and thus quite unlikely to be incorporated into the IEEE WAVE framework.

Motivated by these issues, we attempt to address the following questions in this paper:

1. How does the behavior of CSMA/CA for broadcast scale with increasing number of vehicles and what is the core reason for the congestion issues?

2. Are the proposed DCC algorithms effective in mitigating the congestion issues? If not, are there approaches within their framework to improve the behavior?

3. Is it possible to employ a synchronous TDM based MAC scheme over the existing IEEE 802.11p standard without violating it?

The main contributions of the paper are:

1. First, through simulations we clarify the principal reason behind the congestion issues of IEEE 802.11p MAC at high densities which are widely reported in the literature. We observe that the carrier sense mechanism in 802.11p to prevent nearby devices from transmitting simultaneously breaks down at high densities. Many devices within the carrier sense range synchronously count down their back-off timers. In fact, 802.11p MAC's performance is observed to tend towards the behavior of an ALOHA-like MAC[1] indicating that the CSMA mechanism is ineffective in providing a 'guard zone' around the transmitters.

2. Next, we evaluate the performance of the proposed DCC algorithm [6]. The DCC approach is to change a set of transmit parameters based on measured channel load. We identify that the proposed algorithm is too conservative in response to increasing density. Using transmit power control (TPC) as an example, we develop a systematic method to design the state-machine and observe that the modified TPC algorithm can improve on 802.11p operating with or without the DCC algorithm.

3. Finally, we construct a synchronous time-division multiplexing (TDM) overlay on the IEEE 802.11p MAC that scales elegantly with increasing density. We show that it is possible to bypass the IEEE 802.11p back-off mechanism (*without* violating the standard) by carefully injecting packets into the MAC layer at globally synchronized slots. By choosing a good slot, the devices can ensure that the concurrent transmitters are as far away from each other as the density will allow. A good configuration of simultaneous transmitters can improve the packet reception rate for close-by nodes significantly.

1.1 Related Work

The behavior of IEEE 802.11p when used to transmit periodic broadcast messages has been extensively studied in the literature [13, 21, 23, 10, 12, 24, 22, 18]. The inability of CSMA/CA based MAC systems to deterministically provide a good packet reception rate vs. distance or a short inter message reception delay have been described in [10, 24, 22].

A typical approach for improving the discovery performance is a distributed congestion control mechanism as utilized in [6, 16, 14, 22, 18]. The basic idea is to utilize an underlying state machine which updates its transmit parameters (including power, periodicity and carrier sense range) as a response to the observed channel load. Typically, the response is to reduce the power or periodicity when a larger channel load (based on channel busy time) is observed. In [16], the authors design a congestion control mechanism where the packet injection rate is controlled to attain a given target channel load. However, the issue of discovery performance as well as the choice of the optimal channel load are not explicitly considered. In [14], the authors design a cross-layer control system where the objective is to not improve the efficiency of the MAC but to improve the vehicle tracking accuracy. The authors consider a lossy shared channel where increased message frequency can increase the channel congestion and effectively cause a loss in accuracy of other vehicles' positions. The proposed algorithm is a method to adapt the periodicity of transmission to attain the optimal

[1]In an ALOHA-MAC, the transmitting devices form a spatial Poisson process and the probability of two devices transmitting simultaneously is independent of the distance between them.

accuracy. In [18], the authors control the power based on the relative position of the weakest decodable packet received. In [22], the authors consider a distributed scheme to maximize the minimum transmit power among devices under a maximum channel load criterion.

The DCC algorithm introduced by ETSI for mitigating packet collisions also adapts the transmission parameters based on observed channel load [6]. However, we are unaware of any literature that provides clear guarantees on the stablity or the efficiency of the algorithm with DCC recommendations. In Section 3, we study the behavior of nominal parameters and algorithms provided in [6].

Alternatively, synchronous approaches to scheduling periodic transmissions have also been considered. In [8, 9] the authors propose a globally synchronous slotted system where devices choose a good slot (a resource block) based on the observed energy within the slot. They compare such approaches with the CSMA based mechanisms and demonstrate that synchronous schemes can outperform CSMA in discovering neighboring devices. Such schemes have also been used in device-to-device systems such as Flashlinq [25].

A typical way to introduce a similar slotted structure in CSMA/CA is to introduce another layer above the MAC and create a slotted TDM overlay [20]. If the slot size is large ($\sim$ 10 broadcast packets long) as proposed in [17], a device has the flexibility to choose the best sub-slot, but the concurrent transmitters within a slot can still be close to each other. In this paper, we propose a much more fine-grained slotted overlay with specific slot durations to ensure a 'TDM' like behavior even with IEEE 802.11p MAC. Furthermore, the approach does not violate the current standard and is compliant with the IEEE 802.11p MAC specifications. This is the first such approach to the best of our knowledge.

2. PERFORMANCE OF IEEE 802.11P MAC

The IEEE 802.11 MAC mechanism has been observed to have congestion issues at high densities by many studies. In this section, we analyze the core reason behind the deterioration of the IEEE 802.11p MAC with increasing density. We observe that at high densities, many devices synchronously count their timers down and collide with transmitters within their sensing range. We provide simulations that show that IEEE 802.11p MAC does not guarantee any 'guard zones' around a transmitter (region where no other transmitters are allowed to transmit simultaneously) and observe that the limitations of IEEE 802.11p result in the packet reception performance tending towards an undesirable ALOHA-like behavior.

To test the performance of broadcasting using an IEEE 802.11p MAC, we use the ns2 platform [7] (a discrete-event simulation platform) to simulate the asynchronous IEEE 802.11p channel access algorithm and associated lower layer functions. The IEEE 802.11p functionalities were simulated by using the 80211MacExt package in ns2 [11].

Throughput this paper, we consider a 6-lane scenario as the simulation field. Each lane is 4 meters wide and 2000 meters long. The number of cars on each lane varies, resulting in different car densities. All cars are placed regularly as a grid, i.e., the distance between any two adjacent cars on each lane is the same. The simulation field is illustrated in Figure 1. We use a wrap-around model of a network along the length of the road. Important parameters used in the simulation are listed in Table 1.

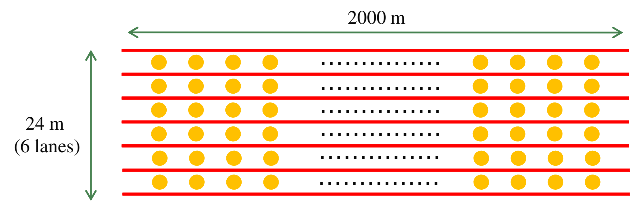

Figure 1: Setup: 6-lane road

Table 1: Simulation parameters

packet size	200 Byte
carrier frequency	5.9 GHz
carrier sense threshold	-76 dBm
noise floor	-96 dBm
transmit power	20 dBm
antenna height	1.5 m
broadcasting CW	15
periodicity	100 ms

We first provide the packet reception performance of the IEEE 802.11p MAC at three different vehicular densities (see Figure 2) and plot the probability of a successful reception as a function of distance. The results shown are in agreement with other reported results in the literature [24, 22].

First, we observe that almost all packets are transmitted within a 100 ms delay at all densities. The back-off counters are discrete, small (CW = 15) and are decremented after each time the channel returns to an idle state after an ongoing transmission and the contention for the next packet begins. Since the packet durations are small (about 0.55 ms), within 100 ms, almost all packets are thus transmitted. Thus, packet collisions rather than delay become the main concern.

For a fixed periodicity of broadcast messages and packet size, distance between concurrent transmitters will reduce at higher densities. This is due to limited time resources, but the number of transmissions becomes larger at high density. Thus, the distance within which packets can be received with high probability also reduces. Ideally, the discovery range should scale such that the number of devices that can be discovered remains a constant at all densities.[2] However, we observe that IEEE 802.11p scales even worse than this expected degradation with density due to two main factors.

1. The carrier sense range (determined by CS threshold and the transmit power) is fixed and the number of devices within the CS range increases with density.

2. The probability of two devices within carrier sense counting down to zero even though they may be very close increases significantly.

To illustrate this point, we plot the probability of successful reception vs. distance normalized for the density. The discovery probability versus normalized distance curves for these densities are given in Figure 3. Normalized distance is obtained by dividing the actual distance by the inter-car

[2]For interference dominated scenarios in regular networks, a change in density is merely scaling the network, and the number of devices that can be discovered in an ideal scenario should be similar in all densities.

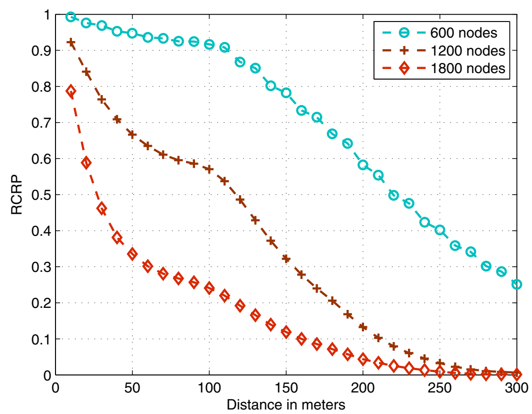

Figure 2: Packet reception performance of plain IEEE 802.11p - the x-axis is tx-rx distance and the y-axis is the ratio of correctly received packets (RCRP).

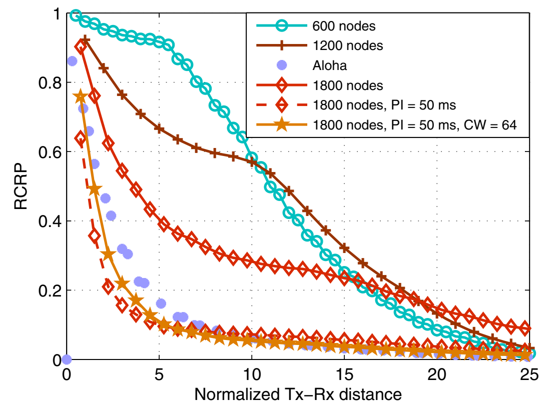

Figure 3: Discoverability performance of 802.11p: the x-axis is distance/inter-car-spacing (normalized distance) and the y-axis is the ratio of correctly received packets (RCRP). PI = 50 ms is a scenario where only 50 ms out of every 100 ms is allowed for broadcast.

spacing corresponding to the given density. From Figure 3, it is seen that 802.11p's discovery decreases significantly with density, especially at shorter distances, even after correcting for the increase in density. The performance also tends towards an ALOHA-like behavior. In slotted ALOHA, each transmitter merely chooses to transmit with a probability independent of other transmitters around it. Note that in ALOHA, the concurrent transmitters form a Poisson point process and there is no mechanism to provide any guard zone around a transmission. In 802.11p, the guard zones are expected to be provided by the carrier sense mechanism. However, if two devices within carrier sense range count their back-off timers to zero simultaneously, the positions of these transmitters can be random within the sensing range, and behave similar to an ALOHA process.

The breakdown of carrier sensing is demonstrated in Figure 4. We obtain the statistics of the distance of the closest transmitter to a vehicle when it is transmitting. The empirical pmf (probability mass function) of this random variable at three different densities are plotted in Figure 4. For higher densities, we observe that CSMA is incapable of avoiding simultaneous transmissions within the carrier sense (CS) range, which is 297 m for the parameters in Table 1. More importantly, not even a smaller guard zone is observed to exist for higher densities.

The simulations suggest that the vanilla version of the 11p MAC degrades significantly and tends towards an ALOHA-like behavior and some solutions are required to improve the discovery of nearby vehicles.

While it may appear that chosing a larger CW may reduce the probability of collision, the approach does not significantly alter the performance (see Figure 3, CW = 64). This may be attributed to the fact that as the CWmin becomes larger, the time taken for a device to countdown to zero also increases linearly. The concurrent transmission probability is related to both the number of devices attempting to transmit when a device is waiting as well as the probability of choosing the same counter. By increasing the CWmin, we reduce the probability of two devices choosing the same counter (roughly by a factor of $\frac{1}{\mathrm{CWmin}}$), however we increase

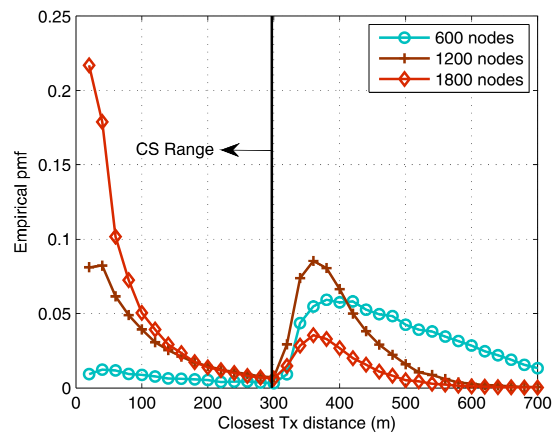

Figure 4: Empirical pmf of the distance to the closest concurrent transmitter.

the number of devices that arrive in the waiting time of a device (which is proportional to CWmin), leaving the collision probability unchanged. Thus, we require alternate approaches to reducing the packet collisions. In the next few sections we explore asynchronous and synchronous algorithms to mitigate the congestion issues.

3. DCC ALGORITHM EVALUATION AND IMPROVEMENT

In this section, we evaluate the Decentralized Congestion Control (DCC) scheme proposed by the European Telecommunications Standards Institute (ETSI) [6] to mitigate the MAC layer congestion issues at high vehicle densities. The DCC mechanism is based on an underlying state machine where the transmit parameters are chosen based the observed channel load, which does not require changes in the underlying PHY/MAC standards as defined in IEEE 802.11p. We study the performance of the proposed DCC parameter setting in the 6-lane highway setup at different vehicle densities. To the best of our knowledge, this is the first result on evaluation of the proposed DCC algorithms in published literature. We observe that the default state parameters defined in the current ETSI standards are very conservative in the use of the channel and also do not discriminate between different vehicle densities effectively. Taking transmit power control as an example, we describe a method to choose appropriate transmit powers at different loads using DCC's state machine based approach, which can significantly enhance the performance of DCC algorithms.

3.1 ETSI algorithm for decentralized congestion control

The state machine proposed for congestion control is depicted in Figure 5. The transmission parameters associated with a state may include transmit power (P), packet transmission interval (PI) and carrier sense threshold (CST), coding scheme (MCS) among other parameters. The channel load (CL) is defined to be the fraction of time the received power was greater than the CST. A state transition to a higher congestion state occurs when all measured CLs for the past *second* are larger than CL_{up}. The transitions towards lower congestion state occur if the CLs measured during the past five seconds are all lower than CL_{down}.

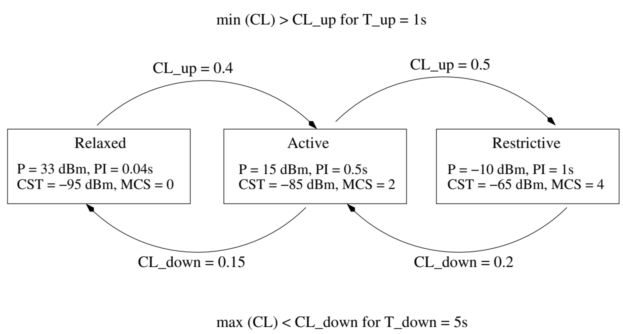

Figure 5: The state machine proposed by ETSI DCC framework. It consists of three states and associated transmit parameters and state transition rules. When the CL is too high, the DCC algorithm tends to change all three parameters simultaneously to ease the congestion.

3.2 Performance evaluation of the ETSI DCC algorithm

To study the above described state machine, we simulate the periodic transmission of safety messages of size 200 bytes in a 6-lane highway, as described in Section 2, at different vehicle densities.

In Figure 6, we plot the fraction of successfully received packets vs. the transmitter-receiver distance using the parameter setting in the ETSI DCC standards. From this figure it appears that the ETSI DCC based system exhibits a good performance in terms of a packet reception rate at distances up to 300 m. However, far less packets are actually transmitted per second compared to a default IEEE 802.11p system. Towards this end, we also present the number of packets *successfully* received per second in ETSI DCC framework, as compared to the default IEEE 802.11p setting, in Figure 7.

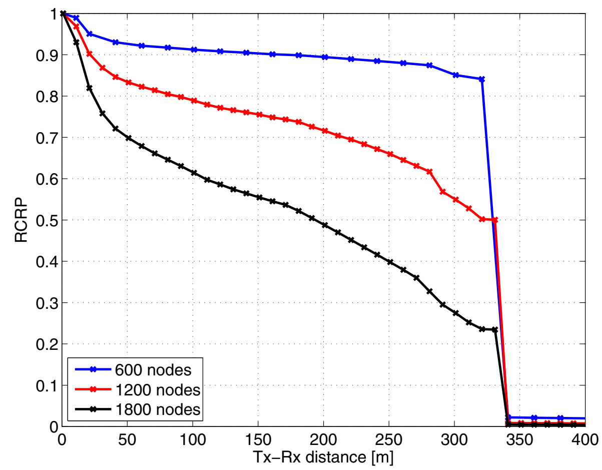

Figure 6: Packet reception rate versus distance for ETSI DCC

States/Number of Vehicles	600	1200	1800
RELAXED	0.01	0.01	0.01
ACTIVE	0.91	0.90	0.90
RESTRICTIVE	0.08	0.09	0.09

Table 2: State distribution at different densities with G5CC state machine.

As we can see in Figure 7, although the ETSI DCC framework does improve the packet reception probability on a *per packet* basis, the total number of successfully received packets can be much *worse* than the plain IEEE 802.11p protocol, especially for vehicles which are less than 100 m away. In other words, the improvement of the packet reception probability is achieved at a huge cost of the number of packet transmissions and it is not clear if such an improvement in worthy at the application level. In particular, for positioning tracking algorithm, which is one of the key elements to support safety applications, the accuracy relies very much on the number of GPS coordinate updates one can receive from neighboring vehicles.

Such an undesirable behavior is caused by the conservative setting in the state parameters and also the transition conditions. For example, the transition from the RELAXED to ACTIVE state reduces the transmit power by a factor of 18 dB and reduces the message transmission rate by a factor of 12. Table 2 shows the fraction of time in each state the ETSI DCC protocol resides in, at different vehicle densities. In fact, the measured channel load with 600 nodes was 18% and even with 1800 nodes, the channel load was measured to be only 30%. From here, we can also see that in majority (> 90%) of the time, independent of the density levels, the DCC protocol stays at the ACTIVE state, which transmits at a moderate transmit power and very conservative periodicity (PI), which significantly reduce the number of transmissions in the network. Thus, the ETSI DCC parameters do not differentiate their behavior at different levels of congestion.

3.3 Improving vehicular discovery under DCC

From the earlier sections, we observe that while packet collisions can reduce the channel efficiency at high densities,

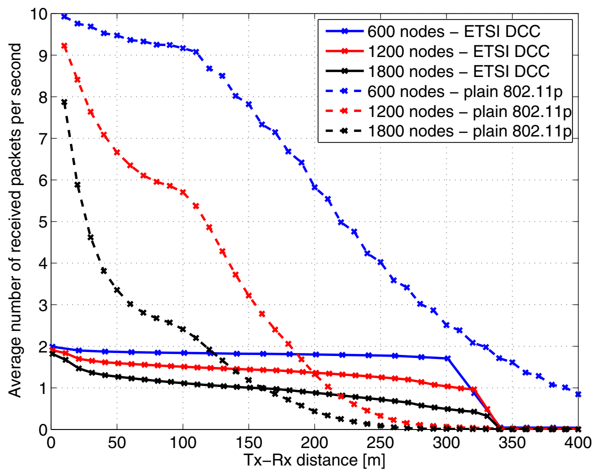

Figure 7: Average absolute number of received packets per second as a function of the transmitter-receiver distance

a very conservative congestion control as defined in current ETSI DCC framework can be as undesirable which can indeed leads to performance losses due the loss of packet transmissions. As we observed from the simulation results, there are two main issues in the ETSI DCC framework:

(1) The targeted CL is too low as defined in Figure 5 (determined by the CL choices in the state transition conditions).

(2) The number of states is too coarse, which leads to very conservative settings in both ACTIVE and RESTRICTIVE settings.

In this section, we attempt to fix these two issues and demonstrate that significant performance gain can be obtained, even with power control itself alone. In other words, to focus on the performance of the state machine, we limit ourselves to only controlling the transmit power (TPC - transmit power control) while keeping the broadcast periodicity, carrier sense threshold and modulation schemes constant. We set the broadcast periodicity to one every 100 ms, the CST = -76 dBm and the modulation scheme to be rate-$\frac{1}{2}$ QPSK. The values are the same as the ones employed in Section 2.

To address the first item above, we ask the following question: what is a *good* targeted channel load (CL)? Intuitively, the channel load (fraction of time the received power is above the CST) can be interpreted as a measure of the number of transmitting devices within a node's carrier sensing (CS) range. If a measured channel load is too small, then the channel is not necessarily fully utilized whereas if the channel load is too high, too many concurrent transmissions will occur within the CS range of the selected transmitter,causing packet collisions which leads to a behavior similar to ALOHA networks. Thus, the tradeoff is to find a good channel load that is large enough to allow for many simultaneous transmissions, while still keeping the packet collisions low. Thus, our approach to transmit power control is two-fold. (i) We identify the channel load that is a good balance between improving channel utilization and packet collisions. We identify this value through simulations

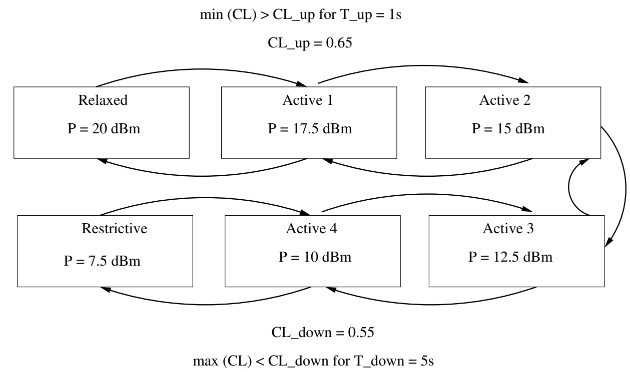

Figure 8: State machine with six states for adapted TPC scheme

at different vehicular densities. (ii) Choose the state transition parameters so that the state machine to operate close to the optimal channel load.

In the 6-lane, 2km-highway scenario, we analyze the broadcast communication performance for several node densities from 25 cars/km/lane to 150 cars/km/ lane. For each density, we consider the average number of correctly received packets by all nodes around the receiver via simulations and study the optimal channel load which maximizes this. Towards this end, we vary the transmit power (the same power for all nodes since the network is symmetric) and compare the simulated result on average number of vehicles discovered, we determine the optimal channel load leading to the best average of received packets in each density. We observed that in all densities, the total number of packets received increased with increasing channel load up to about an occupancy of about 0.65 and as the channel load approached 0.8, the total number of received packets decreased. The optimum channel load at densities of 900, 1200, 1500 and 1800 vehicles all were within a factor of 0.05 from a channel load of 0.65. Thus, simulation results show that 0.65 (or equivalently, 65% of the time above CST) is a good choice of channel load regardless of the vehicle density.

Regarding the second issue of limited number of states in ETSI DCC, we can simply add more states to the state machine with smaller stepsize in power choices. Thus, we propose the following state machine for a dedicated TPC scheme as shown in Figure 8. In order to track the node density with better granularity, 6 states are employed instead of 3 as proposed in the original ETSI DCC proposal. By selecting 55% and 65% as the CL levels that trigger state transitions, for all states in the state machine, it is ensured that the system will exhibit an increased performance by operating at a higher channel load. In line with the observations for optimum transmit powers, the powers used in the state machine vary from 7.5 dBm to 20 dBm with a 2.5 dB stepsize.

In Figure 9, we compare the performance of the proposed TPC scheme and the original IEEE802.11p settings without congestion control. We can make the following observations on the proposed TPC:

(1) At high density, the proposed TPC has a significant improvement for vehicles close by. For example, at 1800 nodes, the plain IEEE 802.11p MAC leads to a packet reception rate of 35% for vehicles 50m away,

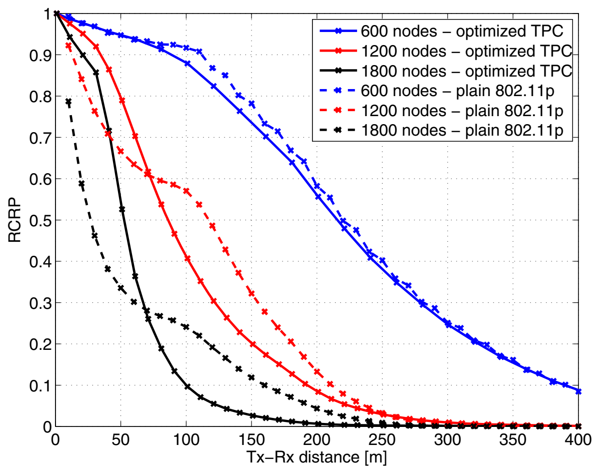

Figure 9: Performance comparison between plain 802.11p and optimized TPC

States/Number of Vehicles	600	1200	1800
RELAXED	0.61	0.01	0.01
ACTIVE1	0.17	0.03	0.01
ACTIVE2	0.12	0.1	0.01
ACTIVE3	0.05	0.2	0.02
ACTIVE4	0.03	0.2	0.03
RESTRICTIVE	0.02	0.46	0.92

Table 3: State Distribution with Transmit power control.

while the proposed TPC leads to 60% packet reception rate, which is clearly much better.

(2) The undesirable ALOHA type of behavior in plain IEEE 802.11p is completely removed under the proposed TPC protocol. Now PDR against normalized range varies very little at different densities.

(3) Since the proposed TPC protocol does not involve transmission periodicity (TP) change, the number of transmitted packets *per second* does not change between the plain IEEE 802.11p protocol or the proposed TPC. Thus, unlike the ETSI DCC protocol, the gain by proposed TPC does not come with a loss in the number of received packets per second.

In Table 3, the probability distribution of the employed states is illustrated for three scenarios with 600, 1200 and 1800 nodes respectively. The tendency is for the lowest state to have the highest probability at low densities and for the highest state to occur more often at high densities, demonstrating that the TPC scheme follows the vehicular density better than DCC (compared to Table 2).

We note that not all nodes within one scenario converge to the exact same transmit power/state even though they were uniformly distributed. This may be due to the random initial states and the hysteresis implied by the state transitions. This variation in power may lead to an unfairness among transmitters in their ability to reach many other receivers. This issue requires further investigation and may be an inherent issue with state machine based approaches.

While the TPC scheme provides improvements over the default IEEE 802.11p MAC and the default DCC parameters, we propose a scheme that can further improve upon the TPC algorithm.

4. STANDARD COMPLIANT SYNCHRONOUS MAC

In Section 3, we considered a power control approach within the DCC framework to improve the behavior, which is fully compatible with IEEE 802.11p standards and thus *asynchronous* in nature. In this section, we describe a mechanism to achieve *synchronous* TDM operation *without* violating the IEEE 802.11p specifications.

The synchronous MAC design provides two main advantages over asynchronous mechanisms. (i) The resource allocated to each device does not change from one broadcast interval to other. Instead, it stays the same for longer durations, and changes only with deviations in density or topology. Hence, it is possible to develop algorithms that converge to an overall efficient resource allocation and remain in that allocation until the situation is changed. (ii) The transmission attempts of each vehicle are periodic. Hence, unlike asynchronous mechanism, packet delays are deterministic and do not incur large variations. These improvements can be provided by means of introducing an additional TDM *SYNC* layer above the existing MAC layer to create fixed resources and by a distributed resource selection (and reselection) algorithm to allocate the TDM resources. The results in this section indicate that a TDM based approach can further improve on gains from power control based approaches.

4.1 Obtaining synchronous transmissions using IEEE 802.11p MAC

The basic idea of an overlay MAC is to provide a slotted frame structure above the MAC layer that injects packets into the IEEE 802.11p MAC in a deterministic synchronous manner. More specifically, it involves dividing the 100 ms broadcast interval (assuming 100 ms to be periodicity of safety message broadcast) into a number of small and non-overlapping time slots (resources), and selecting a particular slot-boundary for each node to inject its broadcast packet into the MAC. As a result, the devices with the same chosen resource transmit concurrently.

For this approach to provide significant gains in device discovery, the following issues need to be addressed.

1. CSMA is an asynchronous mechanism where the packet injection times and packet transmission times may not coincide. The actual packet transmission time depends on the device's backoff process. Thus, the overlay mechanism is required to ensure that the packet transmission times of all devices in the same resource are simultaneous without violating the IEEE 802.11p MAC standard.

2. Devices need to obtain accurate timing (< 2 μs) for synchronous operation. This is to ensure that two devices transmitting in the same resource do not carrier sense each other and yield to each other.

Next, we show that these two issues can be addressed, and the performance of the resulting MAC is significantly better.

4.1.1 *Controlling the channel access time*

In our current setup, the 100 ms broadcast interval consists of a short guard time (1 ms) and 180 resources, each of length 0.55 ms. The time domain structure is illustrated in Figure 10. The slot duration of 0.55 ms is chosen to accommodate a 200 byte packet (typical safety broadcast size) transmitted at 6 Mbps (typical safety broadcast rate) and a silence period of duration which is slightly longer (e.g. $1, \mu$s) than EIFS. Note that the duration of the slots and number of slots can be changed without affecting the fundamental slotted behavior.

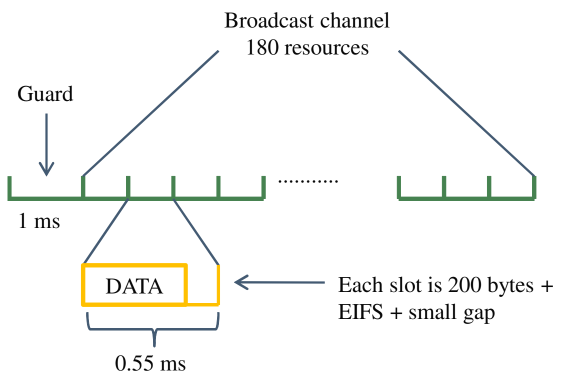

Figure 10: Time domain structure of the channel

Each node selects a resource in the broadcast interval. At the beginning of the selected resource, the SYNC layer pushes a packet into the MAC layer to begin channel contention. By following the IEEE 802.11 MAC standard, once the MAC receives a packet from upper layer, it will wait to transmit until its current backoff counter procedure is completed. If the backoff has completed but it senses the channel to be busy at the moment the packet was injected, it will begin a new backoff process. After a packet has been transmitted, a device goes into a post-transmit back-off. Our slot structure design ensures that when the MAC layer gets a packet, it will have the permission to transmit immediately. The following observations explain this behavior.

1. A device can transmit a packet without waiting if it senses the channel to have been empty for a duration greater than or equal to EIFS and its backoff counter is zero.

2. Since the resource duration is larger than the sum of the packet transmission time and EIFS duration + ε, the channel would be sensed to be idle for at least an EIFS duration at the beginning of any given slot.

3. When the MAC gets a packet at the beginning of a slot, its backoff counter would have run down to 0. Note that any post-transmission backoff procedure would have been initiated at least one broadcast interval (100 ms) ago. Furthermore, the slot structure has enough silence periods (≥ 180) during a broadcast interval for a device to count its back-off timer down to zero even for the largest possible value (which in our setting is 15). Thus, each device would have a back-off counter that is zero when a new packet arrives and also sense the channel to be idle.

4.1.2 *Strict synchronization through GPS*

In vehicular networks, access to a Global Positioning Service (GPS) is necessary for obtaining precise location information. Since the same GPS receiver can provide synchronization of sub-microsecond level accuracy, it is sufficient for the successful operation of our algorithm.

Based on the above observations, we can conclude that the MAC layer essentially bypasses the random back-off in the CSMA mechanism and allows for synchronous transmissions without violating the IEEE 802.11p standard. We note that the above observations have also been confirmed through ns2 simulations.

4.2 Resource allocation algorithm

Given that the IEEE 802.11p MAC can be manipulated to operate in a slotted, synchronous manner, the next aspect is an algorithm to decide which transmitters can occupy the same resource simultaneously. Here, the key idea is to allocate resources such that devices that occupy the same resource are as far apart as possible. This naturally leads to a uniform packing of users sharing the same resource. However, to obtain this behavior in a distributed manner, each node executes the following algorithm. We remark that the algorithm provided below describes the crucial aspects of the method and is not meant to be comprehensive.

1. Every device observes the average energy in each resource in the past K broadcast intervals and ranks these slots (resources) in increasing order of observed average energy. Let the total number of slots be N.

2. Initially, it picks a slot uniformly at random from the first M out of the N slots. Note that these M slots are the ones with the lowest observed energy and typically $M << N$.

3. Periodically, a device transmits a slightly smaller packet in its resource so that it can listen to other devices in the remaining time within its resource. If the current slot is observed to not belong to the top M slots due to changes in topology or arrival of other users, it picks a new resource uniformly at random from the top M slots. The broadcast interval in which the smaller packet is transmitted is chosen independently and randomly by the devices so that not all devices transmit the shorter packet simultaneously. The devices transmit one shorter packet in L broadcast intervals on an average.

In essence, each device greedily chooses a low energy resource as it arrives into the system and re-selects its resource when it identifies that its current resource is no longer 'efficient'. The smaller duration transmissions can be achieved either by sending shorter packets or using a slightly higher rate code. The behavior of a greedy resource selection has been studied in [19] where the authors demonstrate that the greedy resource performs quite close to an optimal packing.

4.3 Simulation Setup and Results

For our simulations, we use a modified ns-2 code. The channel model and node deployment are identical to that in Section 2. The MAC layer protocol also remains unchanged. The main difference is the introduction of a TDM overlay layer. The packets are injected into the IEEE 802.11p MAC

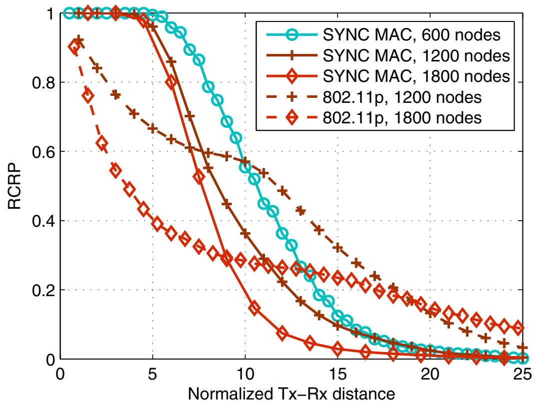

Figure 11: Packet reception performance of SYNC MAC

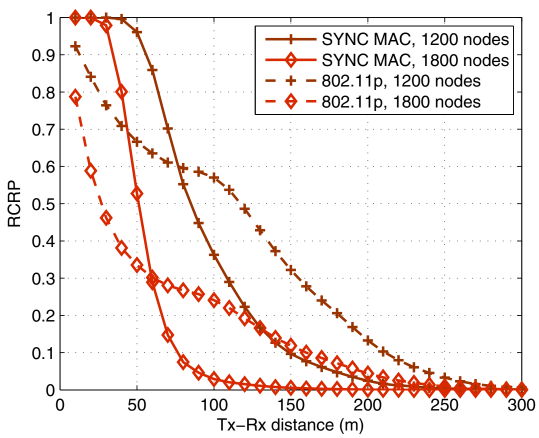

Figure 12: Performance comparison between 802.11p and SYNC MAC

at the chosen resource boundary as described in 4.1. The number of resources $N = 180$ is chosen such that the duration of each slot is 0.55 ms long. We assume that all the broadcast packets are transmitted with the same modulation and coding, which is rate-1/2 QPSK. The shorter-length packets are generated by using rate-3/4 QPSK. We use a history of $K = 2$ broadcast intervals and $M = 20$ corresponding to selecting a slot at random from the top 11.11% of the least energy slots. First, we consider the packet reception performance of the synchronous MAC at varying densities. We plot a figure (similar to Figure 3) that is the packet reception probability versus normalized distance (normalized by inter-car spacing). In Figure 11, we observe that the normalized performance of the SYNC MAC does not significantly change with density. The same plots without the normalization are reproduced in Figure 12, and compared with the corresponding performance obtained by the baseline CSMA scheme without any power control. We observe that the discovery probability for synchronous MAC is better than CSMA at close distances, and is very close to probability 1. While CSMA's discovery performance is slightly better at longer distances, the absolute value of the discovery probability is still low. For example, with 1200 nodes, the 90% discovery probability distance is 55 m in the slotted system whereas CSMA's 90% discovery probability distance is 13 m. At a density of 1800 nodes, the 90% discovery probability distance in the slotted system is reduced to 35 m while even the closest nodes are not guaranteed to reach 90% discovery probability with CSMA. Thus, the guard zone is much more pronounced with the slotted system providing high probability reception for close-by devices.

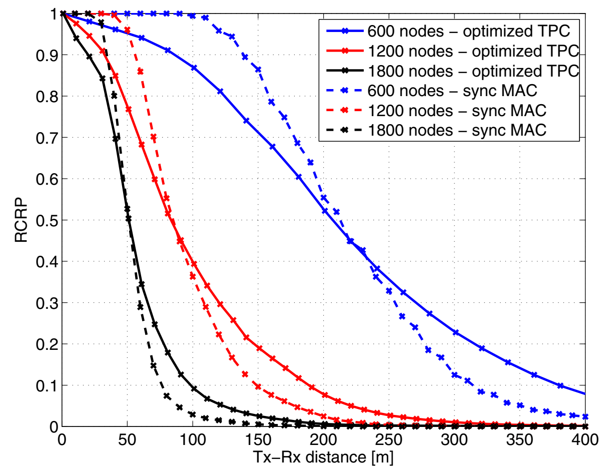

Figure 13: Comparison between the optimized TPC scheme and the sync MAC. The sync MAC performs better than the optimized TPC especially at close ranges.

The convergence time of the resource reselection was also studied through simulations (figures omitted due to space constraints).The devices are initialized with a resource randomly chosen from the 180 slots. We study the time taken for devices to settle down to their steady state resource allocation. In our simulations we observe that almost all devices ($> 99\%$) converge within nearly 15 broadcast intervals (in about 1.5 seconds) starting from a purely random scenario. The algorithm was observed to converge fast from a completely random initial state (i.e. converge fast and provide a good discovery behavior) even under both fading and mobility. The results are omitted due to space constraints.

Finally, the curves comparing the performance of the transmit power control (TPC) with the synchronous MAC scheme are presented in Figure 13. With TPC, we observe that while it improves on unoptimized parameters of plain IEEE 802.11p, its packet reception performance at close distances is still inferior to the synchronous MAC. Although TPC performs marginally better at far away distances, it performs better only in regimes where the absolute performance is already low.

5. CONCLUSIONS

In conclusion, we observe that even after optimization, the distributed power control schemes still do not perform as well as synchronous schemes. Furthermore, the issue of unfairness among different users remains for distributed power control. The synchronous schemes are observed to be stable and perform reliably better than the asynchronous algorithm. However, two important issues remain: (i) very good synchronization is required among all devices to converge to a stable slot structure, and (ii) inter-operability with asynchronous and legacy devices. In terms of the feasibility of implementation, we note that some existing Atheros chipsets have the ability to delay the packet injection into the MAC layer based on an external input (like GPS) which can be used to generate the slotted time behavior. Thus, the promising results in the previous section suggest that it may be worthwhile to study a hardware implementation of the synchronous algorithm using off-the-shelf Wi-Fi chips.

Among asynchronous approaches, there are other methods such as periodicity control and carrier sense threshold control which also need to be investigated in greater detail.

6. REFERENCES

[1] Vehicular Safety Communications Project Task 3. 2005.

[2] *IEEE 802.11-2007 Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*. 2007.

[3] *IEEE 802.11p Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, Amendment 6: Wireless Access in Vehicular Environments*. 2007.

[4] *ETSI ES 202 663: European profile standard for the physical and medium access control layer of Intelligent Transport Systems operating in the 5 GHz frequency band*. 2010.

[5] Draft DSRC message communication minimum performance requirements basic safety message for vehicle safety applications. *SAE Draft Std. J2945.1 Revision 2.2, SAE Int. DSRC Committee*, Apr. 2011.

[6] *ETSI TS 102 687: Decentralized Congestion Control Mechanisms for Intelligent Transport Systems operating in the 5 GHz range; Access layer part*. 2011.

[7] *ns2 Simulator: http://www.isi.edu/nsnam/ns/*. 2011.

[8] K. Bilstrup, E. Uhlemann, E. Strom, and U. Bilstrup. Evaluation of the IEEE 802.11p MAC method for vehicle-to-vehicle communication. In *IEEE 68th Vehicular Technology Conference (VTC)*, Sep. 2008.

[9] K. Bilstrup, E. Uhlemann, E. G. Ström, and U. Bilstrup. On the ability of the 802.11p MAC method and STDMA to support real-time vehicle-to-vehicle communication. *EURASIP J. Wirel. Commun. Netw.*, Jan. 2009.

[10] Q. Chen, D. Jiang, and L. Delgrossi. IEEE 1609.4 DSRC multi-channel operations and its implications on vehicle safety communications. In *IEEE Vehicular Networking Conference (VNC)*, Oct. 2009.

[11] Q. Chen, F. Schmidt-Eisenlohr, D. Jiang, M. Torrent-Moreno, L. Delgrossi, and H. Hartenstein. Overhaul of IEEE 802.11 modeling and simulation in ns-2. In *Proceedings of the 10th ACM Symposium on Modeling, analysis, and simulation of wireless and mobile systems (MSWiM)*, pages 159–168, 2007.

[12] S. Eichler. Performance evaluation of the IEEE 802.11p WAVE communication standard. In *IEEE 66th Vehicular Technology Conference (VTC)*, pages 2199–2203, Oct. 2007.

[13] H. Hartenstein and K. Laberteaux. A tutorial survey on vehicular ad hoc networks. *IEEE Communications Magazine*, 46(6):164 –171, Jun. 2008.

[14] C.-L. Huang, Y. Fallah, R. Sengupta, and H. Krishnan. Intervehicle transmission rate control for cooperative active safety system. *IEEE Trans. onIntelligent Transportation Systems*, 12(3):645 –658, Sep. 2011.

[15] J. B. Kenney. Dedicated short-range communications (DSRC) standards in the United States. *Proceedings of the IEEE*, 99(7):1162 –1182, Jul. 2011.

[16] J. B. Kenney, G. Bansal, and C. E. Rohrs. LIMERIC: a linear message rate control algorithm for vehicular DSRC systems. In *Proceedings of the Eighth ACM international workshop on Vehicular inter-networking (VANET)*, pages 21–30, 2011.

[17] D. Koutsonikolas, T. Salonidis, H. Lundgren, P. LeGuyadec, Y. C. Hu, and I. Sheriff. TDM MAC protocol design and implementation for wireless mesh networks. In *Proceedings of the ACM CoNEXT Conference*, 2008.

[18] H. Lu and C. Poellabauer. Balancing broadcast reliability and transmission range in VANETs. In *IEEE Vehicular Networking Conference (VNC)*, Dec. 2010.

[19] J. Ni, R. Srikant, and X. Wu. Coloring spatial point processes with applications to peer discovery in large wireless networks. In *SIGMETRICS*, 2010.

[20] A. Rao and I. Stoica. An overlay MAC layer for 802.11 networks. In *Proceedings of the 3rd ACM international conference on Mobile systems, applications, and services (MobiSys)*, pages 135–148, 2005.

[21] M. Torrent-Moreno, D. Jiang, and H. Hartenstein. Broadcast reception rates and effects of priority access in 802.11-based vehicular ad-hoc networks. In *Proceedings of the 1st ACM international workshop on Vehicular ad hoc networks (VANET)*, pages 10–18, 2004.

[22] M. Torrent-Moreno, J. Mittag, P. Santi, and H. Hartenstein. Vehicle-to-vehicle communication: Fair transmit power control for safety-critical information. *IEEE Trans. on Vehicular Technology*, 2009.

[23] Y. Wang, A. Ahmed, B. Krishnamachari, and K. Psounis. IEEE 802.11p performance evaluation and protocol enhancement. In *IEEE International Conference on Vehicular Electronics and Safety (ICVES)*, pages 317 –322, Sep. 2008.

[24] Z. Wang and M. Hassan. How much of DSRC is available for non-safety use? In *Proceedings of the fifth ACM international workshop on VehiculAr Inter-NETworking*, VANET '08, 2008.

[25] X. Wu, S. Tavildar, S. Shakkottai, T. Richardson, J. Li, R. Laroia, and A. Jovicic. Flashlinq: A synchronous distributed scheduler for peer-to-peer ad hoc networks. In *48th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 514 –521, Oct. 2010.