

Scalable Cross-Layer Wireless Access Control Using Multi-Carrier Burst Contention

Bogdan Roman, *Student Member, IEEE*, Ian Wassell and Ioannis Chatzigeorgiou, *Member, IEEE*

Abstract—The increasing demand for wireless access in vehicular environments (WAVE) supporting a wide range of applications such as traffic safety, surveying, infotainment etc., makes robust channel access schemes a high priority. The presence of selective fading, variable topologies, high density of nodes and feasibility issues represent important challenges in vehicular networks. We present *Multi-Carrier Burst Contention*, a cross-layer protocol based on a contention scheme that spans both time and frequency domains, employing short and unmodulated energy bursts and a randomized and recursive node-elimination mechanism in order to resolve collisions. It can overcome many of the vehicular environment challenges and provide desirable WAVE features such as scalability, robustness, prioritized access and others. We address physical layer related challenges, present an analytical model, hardware implementation and performance results from theoretical analysis, hardware measurements and simulations, which were run in comparison with the IEEE 802.11p. The results show high scalability and resilience to channel fading and variable topologies and a considerable performance improvement over IEEE 802.11p.

Index Terms—Contention, cross layer design, FFT, IEEE 802.11, MAC, leader election, OFDM, PHY, vehicular networks

I. INTRODUCTION

INTELLIGENT Transportation Systems (ITS) encompass collective efforts to improve traffic safety and advance driving experience by enabling vehicle-to-vehicle, vehicle-to-roadside and vehicle-to-infrastructure communications. Also envisioned are bandwidth-sensitive communications owing to the increasing demand for in-vehicle infotainment and to the potential of vehicles to act as real-time environment surveyors. Such applications can yield networks with very high number and density of nodes, which, coupled with the inherent challenges such as high mobility, highly variable topologies and channel fading, makes efficient medium access control (MAC) a high priority and a challenge.

IEEE 802.11-based solutions initially developed for wireless LANs and mobile ad-hoc networks (MANET) were also investigated for vehicle-to-roadside topologies and vehicular ad-hoc networks (VANET). The IEEE 802.11p amendment [1] to the IEEE 802.11-2007 standard [2], is dedicated to providing wireless access in vehicular environments (WAVE) in the 5.9 GHz band. IEEE 802.11p uses an orthogonal frequency division multiplexing (OFDM) physical layer (PHY)

and employs carrier sense multiple access (CSMA) and binary exponential backoff, very similar in most aspects to IEEE 802.11a. One major difference is that the channel bandwidth is halved to 10 MHz and the OFDM symbol duration is doubled for increased fading resilience.

An early survey of MAC protocols, designed for MANETs but that could be adapted for VANETs, has been presented in [3]. Protocols specifically proposed for vehicular networks include both contention-based and schedule-based approaches. Contention-based schemes [4]–[6] commonly use CSMA variations with different priority levels to increase resilience to high mobility. As with many CSMA-based schemes, these suffer from high collision rate and unpredictable delays under high loads and fading scenarios [7], which are typically found in vehicular networks. Schedule-based schemes generally use time division multiple access (TDMA) techniques. TDMA was proposed in conjunction with space-division multiplexing [8], [9], where time slots are mapped onto road divisions. Although promising, this will likely result in network sub-utilization for sparse traffic and may be hindered by inaccurate location information. Other schemes [10], [11] achieve TDMA transmission opportunity using a CSMA-based random access channel. These however suffer from the same issues as the CSMA-based protocols mentioned above. TDMA was also proposed in conjunction with directional antennas [12] or self-configuring algorithms [13] to provide bounded latency, the obvious drawback being increased sensitivity to topology changes and adverse channel qualities. More recently, cooperative communications have also been incorporated in MAC protocols for vehicular networks [14] to further improve network throughput but at the cost of increased implementation complexity and topology dependency. IEEE 802.11p MAC enhancements have also been suggested [15]–[17], these having in common the decrease of the data traffic load using improved prioritization, e.g. avoiding collisions by using polling [15]. Even though the average delay is decreased, they all suffer from the inherent CSMA/CA drawbacks, the worst case collision scenario still being a major issue.

A successful MAC design for environments such as vehicular networks — subject to variable topologies, adverse channel quality and high density and fluctuations in the number of nodes — should strive to provide real performance scalability and resilience to variable and adverse factors, but it should in equal measure carefully investigate and address the underlying PHY, feasibility, implementation and real-world challenges.

This paper proposes a contention-based cross-layer protocol that distributes contention onto the frequency domain as well as the time domain. The proposed scheme, *Multi-Carrier*

This work was supported in part by Girton College at University of Cambridge, the Cambridge European Trust and Felix Telecom Romania.

B. Roman and I. Wassell are with the Computer Laboratory, University of Cambridge, Cambridge CB3 0FD, UK (e-mail: abr28@cam.ac.uk).

I. Chatzigeorgiou was with the Computer Laboratory, University of Cambridge, UK. He is now with the Norwegian University of Science and Technology, Norway.



Fig. 1. *MCBC* data exchange. T is the sync interval and CW denotes the contention window, which comprises R contention rounds of duration T_r . The *ACK* is transmitted by the receiving node.

Burst Contention (*MCBC*), employs a rapidly converging node elimination algorithm [18] and synchronized rounds. We compare *MCBC* to IEEE 802.11p and show that it is scalable in performance, resilient to environment and topology changes and capable to provide prioritized access. More importantly, we did not restrict ourselves to a theoretical study but also built a hardware testbed to validate the design and performed a thorough feasibility study to address challenges related to the PHY and real-world implementations.

The paper is organized as follows. Section II describes the MAC layer with the contention scheme and prioritized access methods. Section III contains the study and solutions to the most challenging PHY issues pertaining to the novel MAC design, namely transmission and detection of specific non-standard signals, frequency shift, channel fading and time shift. Section IV presents the hardware testbed and time synchronization mechanism, built to validate some of the performance and feasibility claims. In Section V, we show an analysis of the proposed scheme taking into account cumulative fading. Finally, Section VI shows performance results in comparison with IEEE 802.11p and we conclude with Section VII.

We note that throughout the paper, we refer to IEEE 802.11p [1] as being the IEEE 802.11-2007 standard [2] with the latest IEEE 802.11p amendments (draft D9.0, Sep 2009).

II. MEDIUM ACCESS CONTROL LAYER (MAC)

The *MCBC* data exchange is similar to IEEE 802.11. The contention mechanism preceding any data frame is however very different (not based on CSMA or backoff strategies). In brief, contention windows (CW) are synchronized and contain R rounds, during which nodes run the *MCBC* contention algorithm detailed in Section II-A. A high-level example with two nodes, STA_A and STA_B , employing the *MCBC* protocol, is shown in Fig. 1 where STA_A happens to win the contention and thus proceeds to transmit its data frame.

Contrary to IEEE 802.11, *MCBC* does not use a four-way handshake (e.g. RTS/CTS). The feedback mechanism of the contention scheme mitigates the hidden node problem intrinsically. Since this happens during the CW , it has the advantage of a much reduced overhead that also reduces power consumption. *MCBC* exchanges no data during contention: it relies on transmitting and detecting very short (a few μs) and unmodulated energy bursts on random and individual subcarrier frequencies. The frames sent after the CW are compliant with the standard [1]. We also note the absence of distributed inter-frame spaces (DIFS) — prioritized access and Quality of Service (QoS) are achieved differently, as explained in Section II-C. *MCBC* nodes can start contenting at synchronized and fixed intervals of duration T , hence a low frame

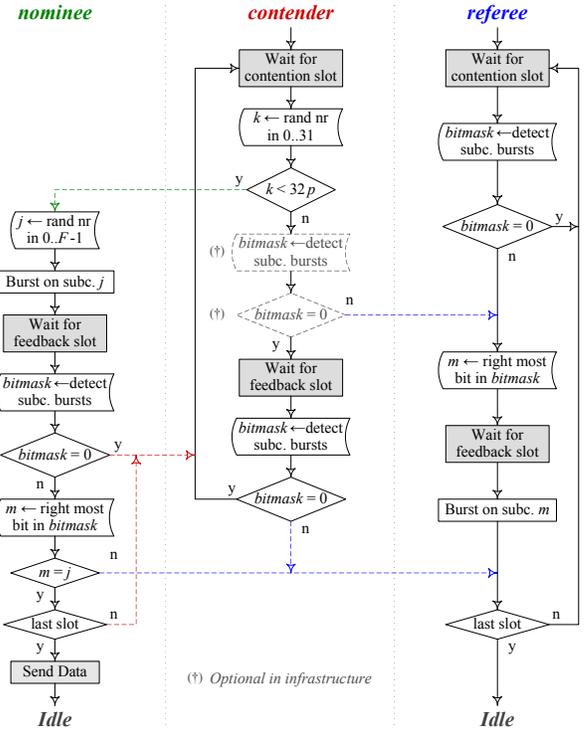


Fig. 2. Contention algorithm running in a node. The node enters contention from Idle as *contender* if it has data to send. In an adhoc network, if it does not have data to send, the node may choose to remain idle or start as *referee*. In an infrastructure network, the AP is the dedicated and only referee, following the same algorithm, and the blue dashed transitions do not exist. Dashed lines represent transitions between *contender*, *nominee* and *referee* states.

collision probability is imperative to maintain a high efficiency. Even though time synchronization is a challenge, we show an effective sync mechanism having a low hardware footprint in Section IV as part of the *MCBC* hardware implementation.

A. Contention Algorithm

The contention scheme is the heart of *MCBC*'s performance. In this section, we assume that each node can transmit and distinguish between F individual subcarrier frequencies. Fig. 3 and its description provide a quick understanding. The process and challenges of transmitting and detecting individual subcarriers is detailed in the next section.

Each CW has a fixed length of R rounds, each comprising a contention slot followed by a feedback slot. Here, a node can be either a *contender*, *nominee* or *referee*. In an infrastructure network, the AP is the dedicated and sole referee. Nodes with data to send start as contenders, while the others will remain idle. In each round, a fraction of contenders are randomly promoted to nominees, meaning they can compete. The nominees who win the round are promoted to contenders for the next round, while those who lose turn idle (if infrastructure) or become referees (if adhoc). The nominees who win the last round initiate their data transmission. Thus, the primary goal of the algorithm is to maximize the probability that there is a single winning nominee in the last round.

Fig. 2 shows the algorithm diagram of a node. In the contention slot, contenders become nominees with probability

	$c=250$ $n=132$	$w=34$	$c=34$ $n=18$	$w=5$	$c=5$ $w=3$	$w=1$
f_4	34		5			
f_3	30		4		1	
f_2	29		6			
f_1	39		3		2	
	ct_1	fb_1	ct_2	fb_2	ct_3	fb_3

$M = 250$
 $R = 3$
 $F = 4$
 $p = [0.5, 0.5, 0.5]$

Fig. 3. Contention session example for an infrastructure network; c and n are the number of contenders and nominees respectively. The shaded cells denote the contention subcarrier activated by the referee in the feedback slots.

p , i.e., they flip a coin. Nominees (those who flipped *yes*) transmit a burst of energy on a subcarrier frequency index j , chosen randomly from $1, \dots, F$. Nodes that flipped *no* listen for bursts on all F subcarriers and, if any are detected, they become referees and record m : the index of the highest indexed subcarrier a burst was heard on — by convention, this is the right-most non-zero bit index in the subcarrier bitmask of F bits provided by the PHY. In the feedback slot, referees transmit an energy burst on subcarrier index m and nominees listen for bursts on all F subcarriers. Nominees which detect their old chosen subcarrier (i.e., nodes for which the right-most m equals j) will change status to contender (won), whilst the others will change status to referee (lost). The contention proceeds into the next round with the new and fewer contenders until the round after which all standing contenders initiate their data transmission.

In an infrastructure network, the AP is the dedicated and unique referee, so nodes that lose a round simply turn idle instead of becoming referees. Here, the blocks marked *optional* in Fig. 2 help in fading channels (see Section VI) to sense whether there were other contention bursts in case the reply from the referee is faded. For good quality channels, these blocks are redundant and can be skipped, i.e., nodes do not have to listen for bursts during the contention slots. In an ad-hoc network these blocks are enabled, and nodes who become referees will remain referees until the end of the last round.

Fig. 3 depicts an example for an infrastructure network of $M = 250$ competing nodes where the contention uses $R = 3$ rounds and $F = 4$ subcarriers. The ad-hoc case was evaluated in [7]. All 250 nodes are initially *contenders* ($c = 250$). During the first contention slot, ct_1 , contenders become *nominees* with probability $p = 0.5$ by flipping a coin. Let's assume that $n = 132$ nodes flip *yes*, thus becoming nominees; subsequently, each of them transmits an energy burst on a random subcarrier f_j , where $j = 1, \dots, F$. The referee listens to the channel to identify the active subcarrier with the highest index, in this case f_4 . During the feedback slot of the first round, fb_1 , the referee transmits an energy burst on f_4 and the 34 nominees that had used it in ct_1 will thus select themselves as round winners. The remaining 118 contenders will detect a different subcarrier than the one they had used in ct_1 and hence will lose the contention. The winners are promoted to contenders ($c = 34$) for the next round and the contention continues similarly until the end of the third round. In fb_3 , the referee selects the subcarrier with the highest index from ct_3 (now f_3) and then the winner initiates its data transmission. If multiple winners survive contention, a packet collision will occur.

The algorithm maximizes the probability that there is a single winner after the last round. This approaches 1 very quickly since approximately $\lceil M(p/F)^R \rceil$ contenders are expected to reach the end of contention, which represents a steep exponential decrease owing to the very low ratio p/F .

If a node does not hear a reply in the feedback slot (due to fading or if there were no referees or if all nodes flipped *no*, so the referees remain silent) then it will select itself as round winner; this also results from the state diagram. Hence, if there is no reply in the feedback slot then all contender nodes win. This ensures protocol correctness: *a*) there is always at least one winner at the end of the contention and *b*) a unique round winner will also be the overall contention winner as well. These theorems are straightforward to prove.

Fairness in *MCBC* is inherent since all contention parameters are fixed and equal among all nodes and there is no dependency on previous attempts. Hence, statistically, there are no nodes having priority over the rest. Fixed parameters also helps making *MCBC* resilient to topology changes.

B. Performance Extensions

In its very essence, the algorithm reduces the number of contenders. We can achieve a steeper descent as the algorithm progresses by applying two performance improvements — at no hardware complexity or delay costs: a different flipping probability p in each round and a non-uniform distribution to choose from the F subcarriers. Since we should avoid the case where all nodes flip *yes* or all nodes flip *no* (in both cases, all nodes win and the round is wasted) and since the number of nodes is being reduced on average after each round, we can enforce a more rapid decrease in the first rounds (when there are more nodes) and a slower decrease in the last rounds (when there are fewer nodes). Thus, p should be an increasing function of the round number r , noting it with p_r . This is taken into account in Section V and used in Section VI.

We can further obtain fewer winners on average by giving a lower probability to the higher indexed subcarriers. Similarly, we should avoid the case where all nodes pick the same subcarrier index so we should employ a probability distribution that is selective in the first rounds and tends to be uniform in the last rounds. We can employ a modified geometric distribution for this purpose having the probability mass function:

$$q_{r,f} = \frac{(1 - \alpha_r) \alpha_r^{f-1}}{1 - \alpha_r^F}, \quad (1)$$

where $r = 1, \dots, R$ is the round number, $f = 1, \dots, F$ is the subcarrier index and $\alpha_r \in [0, 1]$ is the distribution parameter. It is easy to show that $\lim_{\alpha_r \rightarrow 1} q_{r,f} = 1/F$. So, by making α_r an increasing function of r , we obtain the desired effect mentioned above where $q_{r,f}|_{r=R}$ is close or equal to $1/F$.

In hardware terms, different p_r values are achieved using different (rather than equal) comparison thresholds for the random number generator (RNG) and geometric distributed random numbers are achieved using geometrically (rather than uniformly) spaced thresholds. All thresholds are tabulated in lookup tables so complexity and delay are unaffected.

C. Prioritized Access and QoS Extensions

It is important for a MAC protocol to offer prioritized access, especially in vehicular networks, e.g. for safety messages or infrastructure communication. Considering the current trend of multimedia content streaming over the wireless link, QoS capability is also a desired feature.

Here, the first contention round is used for priority signalling: the lower subcarrier indexes are mapped to QoS classes and the higher subcarrier indexes are reserved for emergency signalling. The remaining contention rounds are then used for normal contention. Contenders would burst on the subcarriers corresponding to their message priority, the feedback slot ensuring that only the highest priority nodes would continue in the subsequent contention rounds. A round-robin scheme is inherently being created among same priority nodes. This happens if there is sufficient bandwidth to accommodate all traffic. Otherwise, nodes increase their priority based on their waiting time, to avoid starvation.

The 2004 FCC ruling [19] suggests that safety messages in the Dedicated Short Range Communications (DSRC) context are to be sent on the control channel. Since maximum reliability is required, these need to take precedence over any lower priority messages. In this case, an emergency node will burst in both contention and feedback slots on the reserved subcarriers to silence all other nodes. To avoid collision between safety messages of equal priority, the subsequent rounds are used for contention, the protocol guaranteeing that at least one wins. Since a contention round is very short, around $20 \mu\text{s}$, *MCBC* can afford to employ several contention rounds to achieve maximum reliability: Section VI shows a success probability at the network level of virtually 1.0 when 4 rounds are used. This means that a safety message incurs virtually no additional delay other than potentially the delay caused by other equal or higher priority safety messages who won the channel.

The second method is based on the concepts in Section II-B. A flexible QoS access scheme is achieved if a node selects multiple subcarriers instead of only one, thus increasing its chance of winning, according to its data traffic class. This can be accompanied by an increase of p_r or decrease of α_r . This way, starvation is avoided when there is not enough bandwidth to accommodate all traffic. Otherwise, it becomes less reliable than the above method, i.e., even though fairness among same priority nodes is ensured statistically, there are short-term random delays.

We note that further analysis and evaluation of these features is out of the scope of the paper.

III. PHYSICAL LAYER (PHY)

Using unmodulated individual subcarriers as the basis for the contention algorithm introduces several non-trivial challenges. In this section, we describe the enhancements to the PHY that allow the implementation of *MCBC*, explore various real-world issues that can severely hinder protocol feasibility and we present our solutions to address them.

MCBC uses an augmented version of the OFDM PHY of IEEE 802.11p [1]. Data transmission is unchanged. During contention it transmits and senses custom shaped energy bursts

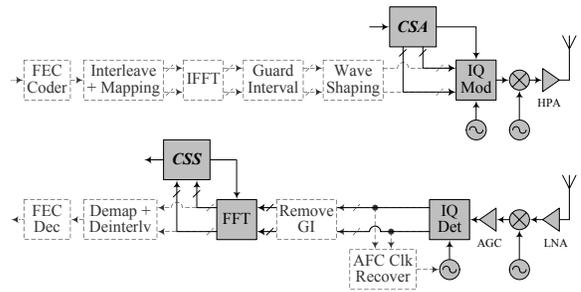


Fig. 4. PHY transceiver. During the CW, only the shaded blocks are active, the dashed blocks are skipped and disabled. During data exchange, the CSA and CSS blocks are skipped and disabled while the other blocks are active.

on individual subcarrier frequencies. We define:

- a *contention subcarrier* or *MCBC subcarrier* as the FFT index of a subcarrier that is used during contention
- a *contention burst* or *MCBC burst* as an unmodulated signal of short duration, carrying no data (a blind unmodulated energy burst), on the frequency of an *MCBC* subcarrier.

The PHY enhancements are two blocks shown in Fig. 4: the Contention Subcarrier Activation (CSA) in the transmitter and Contention Subcarrier Sensing (CSS) in the receiver.

It is possible to transmit and detect energy bursts on a specific subcarrier using the properties of the Fast Fourier Transform (FFT): a signal constellation with a non-zero value at subcarrier index j and zero at all subcarrier indexes $i \neq j$ given as input to the Inverse FFT (IFFT) block will yield a signal with the energy concentrated on subcarrier j . When transmitting (activating) contention bursts, none of the conventional PHY transmitter blocks prior to the IQ modulator are used. This is because the burst signal samples corresponding to the F contention subcarriers are tabulated in look-up tables (or ROMs) inside the CSA block which controls and feeds directly the IQ modulator. When receiving (sensing) contention bursts, none of the receiver blocks after the IQ modulator are used, except for the FFT. This is because we are interested only in extracting the energy levels of the F specific subcarriers at the FFT outputs. The CSS block controls (triggers) the FFT and transforms its output into a simple F -bit bitmask signal, which is passed to the MAC. The bitmask is 1 at all subcarrier indexes where the energy is above a determined threshold and 0 at the others. More details regarding the hardware blocks can be found in Section IV.

With this approach, any signal shape and size can be tabulated which results in a considerably reduced delay during contention, considering that only a small subset of the PHY is used. However, since the bursts are not frequency synchronized and since multiple nodes transmit simultaneous bursts on random subcarriers, the shape and spectrum features of the contention bursts as well as the detection efficiency and reliability are crucial aspects concerning protocol feasibility.

A. Contention Bursts and Frequency Shift

MCBC performs no frequency synchronization for contention (it does for data); no OFDM preamble or pilot subcarriers are used for the contention bursts in order to minimize contention duration and to maximize spectrum efficiency.

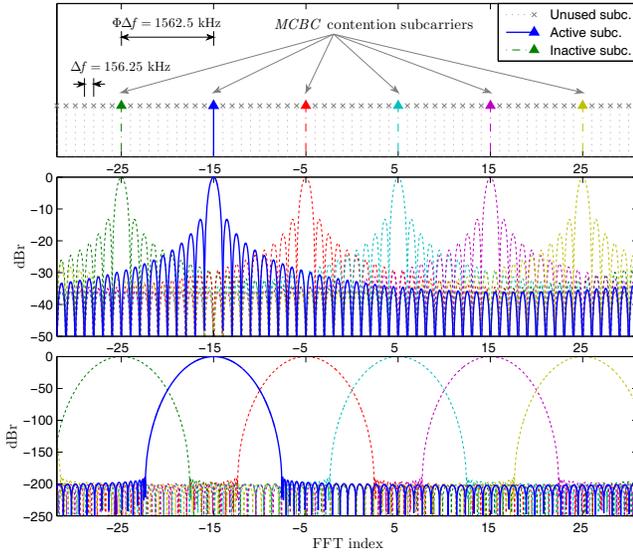


Fig. 5. $F = 6$ contention subcarriers spaced with $\Phi = 10$ times the normal OFDM subcarrier spacing, $\Delta f = 156.25$ kHz [1]. Symbolic spectrum (top), normalized spectrum using a rectangular window (middle) and Ultraspherical window (bottom).

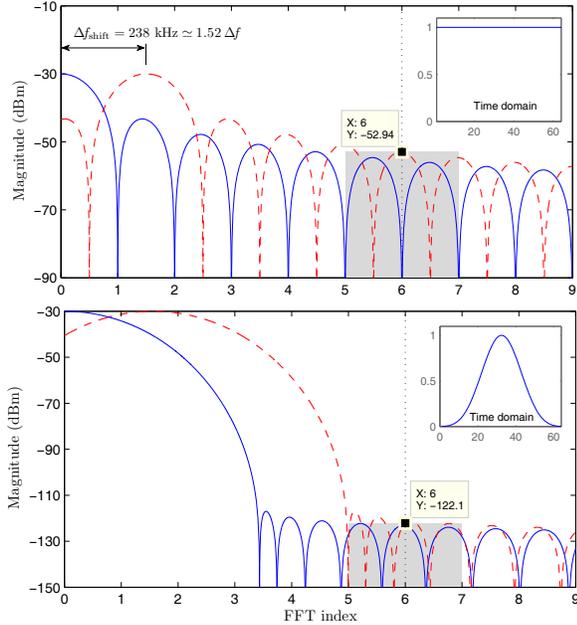


Fig. 6. Frequency response of the rectangular window (left) and Ultraspherical window (right). Full lines represent the ideal response for synchronized center frequencies between the transmitters and the receiver while dashed lines represent the actual response due to frequency shift. The shaded boxes represent the effective FFT filter bandwidth per subcarrier, $B = 2\Delta f$.

However, without frequency synchronization, frequency shift is a major issue that needs to be addressed. The subcarrier spacing in the 10 MHz OFDM channel bandwidth specified in the standard [1] is $\Delta f = 156.25$ kHz given by the 64-point FFT. The standard specifies $C = 52$ subcarriers and a center frequency tolerance of $\max \pm 20$ ppm which gives ± 118 kHz maximum shift at 5.9 GHz. The maximum relative shift between a transmitting and a receiving node is thus $\Delta f_{\text{shift}}^{\max} = 236$ kHz. This is very large, equal to $1.5 \Delta f$, and rules out the possibility of using adjacent subcarriers

(without frequency synchronization). Under these conditions, other sources for frequency shift, such as Doppler, are a negligible problem, e.g. at 100 mph (161 km/h) and 5.9 GHz center frequency, the Doppler shift is only 0.88 kHz.

MCBC addresses the frequency shift problem in two stages, without increasing hardware complexity. First, it uses a subset of $F < C$ subcarriers and spaces them so that any non-zero contribution from adjacent subcarriers is reduced, as shown in Fig. 5 where $F = 6$. For a desired number of contention subcarriers F , the maximum frequency spacing Δf_{cont} normalized to Δf (i.e., the number of FFT bins between two adjacent contention subcarriers) is given by:

$$\Phi = \frac{\Delta f_{\text{cont}}}{\Delta f} = \left\lfloor \frac{C - 1}{F - 1} \right\rfloor. \quad (2)$$

Fig. 5 (middle) shows the spectrum of $F = 6$ contention subcarriers spaced with $\Phi = 10$. This is the FFT response of a 64-point rectangular time window as specified in the standard [1]. Its side-lobe attenuation and fall-off ratio are low; considering the large frequency shift and that multiple nodes can simultaneously burst on the same subcarrier, increasing the received power at the receiver, this can lead to false positives at neighboring contention subcarriers.

Let us assume a worst case scenario, depicted in Fig. 6. Even though this scenario is virtually impossible in practice, it allows us to determine the bounds and safe values for some of the *MCBC* parameters. In this scenario, 20 nodes happen to pick and transmit on the same subcarrier and they travel at very high velocity relative to the receiver, 230 mph (370 km/h), giving a Doppler shift of 2 kHz; the receiver is shifted in frequency with the maximum value $+118$ kHz and every transmitter is shifted in frequency with the opposite maximum value -118 kHz, giving the maximum relative frequency shift between each transmitter and the receiver, i.e., $\Delta f_{\text{shift}} = 238$ kHz $\approx 1.52 \Delta f$, which includes the Doppler shift. Let us also consider a relatively high received power from each node of -43 dBm which gives a total received power of -30 dBm at the receiver. Assuming that the contention subcarrier spacing is $\Phi = 6$, we notice that the minimum attenuation at the adjacent contention subcarrier frequency (i.e., at $f = 6\Delta f$) is only 22.9 dB in the case of the rectangular window (Fig. 6 top). For a detection threshold of -65 dBm [1], the energy sensed in the equivalent FFT filter bandwidth $B = 2\Delta f$ would be higher than the threshold and would yield a false positive at that subcarrier index for the *MCBC* contention algorithm.

The larger spacing Φ allows *MCBC* to use a custom time window. Recently, windows based on Ultraspherical polynomials have been proposed for signal processing [20]. We employ an Ultraspherical window (Figs. 5 and 6 bottom) because of its customization abilities. For our purposes, the advantages are threefold: it has a higher side-lobe attenuation (reduces contribution of neighbor subcarriers) which, consequently, allows to span a slightly larger bandwidth (i.e., use $C > 52$ FFT bins) and has a wider main lobe (improves detection of bursts in case of frequency shift). We note that the resulting spectrum still needs to meet the specified spectrum mask requirements [1]. There is however a compromise: a

TABLE I
MCBC CONTENTION SUBCARRIERS SPECTRUM CHARACTERISTICS

Φ	F_{\max}	FFT bins	Neighbor sub min attn (dB) for $\Delta f_{\text{shift}}^{\max} = 238$ kHz	Max attn (dB) at $f \leq$ 238 kHz
3	21	-30, -27, ..., 30	33.6	34.2 ²
4	15	-28, -26, ..., 28	60.7	9.1
5	12	-28, -23, ..., 27	86.1	7.6
6	10	-27, -21, ..., 27	113	6.3
7	9	-28, -21, ..., 28	140	5.1
8	7	-24, -16, ..., 24	167	4.3
9	7	-27, -18, ..., 27	195	3.6
10	6	-25, -15, ..., 25	224	3.2

wider main lobe (or higher side-lobe attenuation) increases the minimum contention subcarrier spacing Φ_{\min} which in turn decreases the maximum number of contention subcarriers F_{\max} . Noting with W the main lobe width in Hz at the desired side-lobe attenuation level [21], we have:

$$\Phi_{\min} = \left\lceil \frac{W + B + 2\Delta f_{\text{shift}}^{\max}}{2\Delta f} \right\rceil. \quad (3)$$

Plugging (3) into (2), and considering $C = 63$ and the FFT periodicity¹, we obtain:

$$F_{\max} = \left\lfloor \frac{63}{\Phi_{\min}} \right\rfloor. \quad (4)$$

Since it is out of the scope of this paper to investigate window design, we list in Table I the parameters and spectral properties of contention bursts based on Ultraspherical windows which give good performance, for various Φ , F and worst case frequency shift, $\Delta f = 238$ kHz. Performance was evaluated in terms of neighboring MCBC subcarrier attenuation, transmit mask compliance [1] and detection error, as we will show in the next subsections.

Spacing subcarriers and using a different time window does not increase hardware complexity or processing delay. All signal samples, i.e., the time window samples of the F contention subcarriers, are hardwired in look-up tables in the CSA block that feeds the IQ modulator, offering full flexibility with no complexity penalty. The 64 values for subcarrier j cached in the CSA block are:

$$s_{j,k} = w[k] e^{-i2\pi jk\Phi \Delta f} \quad (5)$$

where $j = 1, \dots, F$, $k = 0, \dots, 63$, $w[k]$ is the k^{th} window sample, i is the imaginary number, Φ is the MCBC spacing from Table I, and $\Delta f = 156.25$ kHz is the OFDM subcarrier spacing [1].

B. Time Shift

High mobility, a characteristic of vehicular networks, makes the distance between nodes variable, so simultaneous contention bursts will arrive at the receiver at different time

¹The FFT response is periodic so a wide main lobe at one end (e.g. FFT bin 0) gives a high contribution at the other (FFT bin 63), hence a false positive for the MAC algorithm if all 64 bins are used. Hence, F_{\max} includes all possible contention subcarriers less 1 and the F subcarriers are centered within the 64-point spectrum to ensure spectrum roll-off for adjacent channels.

² $\Phi = 3$ will not be used as it yields an out of mask transmit spectrum and the main lobe is too narrow, affecting burst detection.

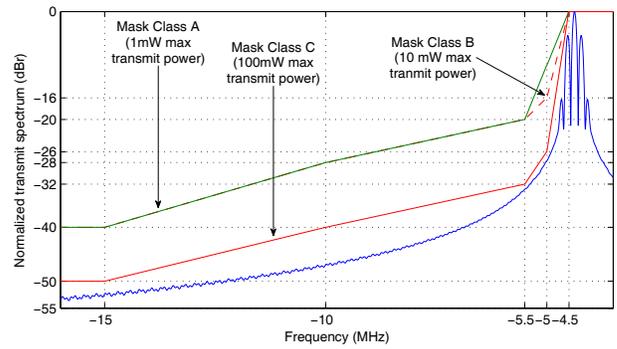


Fig. 7. Worst case transmit spectrum: $\Phi = 4$, $N_{\text{CP}} = 32$ and active MCBC subcarrier index -28 . Mask Class D is not shown as it is meant only for very high transmit powers (max 760 mW).

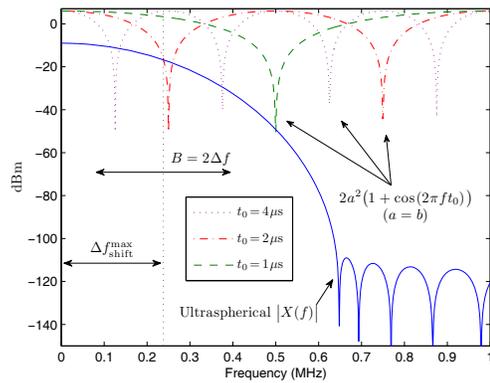


Fig. 8. FFT response of Ultraspherical burst and worst case combined time shift component from two simultaneous transmissions (t_0 is the time shift).

instants due to propagation delays and delay spreads. Independently, even though nodes are time synchronized (time synchronization is detailed in Section IV), there is unavoidably also a time synchronization error. We note the combined maximum time shift of all these effects with t_{shift}^{\max} .

If many of the 64 samples of the Ultraspherical window are missed due to time shift then the spectrum features described in Section III-A will be invalidated. To address this, we use a cyclic prefix (CP) of N_{CP} samples, i.e., a longer time window of $N = 64 + N_{\text{CP}}$ samples where the samples following the 64th sample are circular copies of the first samples, i.e., $w[64m + k] = w[k]$ where $k, m \geq 0$ are integers. With a sufficiently large N_{CP} , the receiver is then able to take 64 samples of the original window $w[k]$. For a given t_{shift}^{\max} , the window size including the CP becomes $N_{\min} = 64 + \lceil f_s t_{\text{shift}}^{\max} \rceil$ where f_s is the sampling frequency, $f_s = 10$ MHz in IEEE 802.11p [1].

Since the FFT sees the time signal as periodic, a CP can cause a discontinuity between the first and last samples (unless $N_{\text{CP}} = 64m$) and the side lobe attenuation of the transmitted burst decreases. This may cause the transmit spectrum to no longer meet the spectrum mask requirements. The Ultraspherical windows parameters we computed in Table I however yield signals that meet the transmit mask requirements for all scenarios, including the worst case shown in Fig. 7.

We need to extract only the received energy level so only the spectral magnitude is of interest. A signal $x(t - t_0)$, where t_0 is the time shift, has the FFT response $X(f) e^{-2\pi i f t_0}$, where

$X(f)$ is the FFT response of $x(t)$. Knowing that $|e^{-2\pi if t_0}| = 1$, the magnitude of the time shifted signal is $|X(f)|$. In *MCBC*, multiple nodes transmit bursts simultaneously. Let us consider two simultaneous bursts (the general case follows similarly), time shifted with t_0 , i.e., $a x(t) + b x(t - t_0)$. Assuming N is large enough to accommodate t_0 , the combined FFT response magnitude is $|X(f)| \cdot |a + b e^{-2\pi if t_0}| = |X(f)| \cdot (a^2 + b^2 + 2ab \cos(2\pi f t_0))$. As shown in Fig. 8, the worst case scenario is when $a = b$ which gives a null at $f = \frac{1}{2t_0}$. This is highly unlikely to happen in practice, yet worth investigating. Further still, if the frequency shift Δf_{shift} equals $\frac{1}{2t_0}$ then the energy level sensed in the effective FFT bandwidth per subcarrier, $B = 2\Delta f = 0.3125$ MHz, would be the lowest. However, we see that for $t_0 < 2 \mu\text{s}$, the window response is little affected (on average it is actually amplified, as expected) for all frequency shifts $\Delta f_{\text{shift}} \leq \Delta f_{\text{shift}}^{\text{max}} = 238$ kHz. Thus, ideally, the maximum time and frequency shifts would satisfy:

$$t_{\text{shift}}^{\text{max}} \leq \frac{1}{2\Delta f_{\text{shift}}^{\text{max}}} \quad (6)$$

Hence, by decreasing either $t_{\text{shift}}^{\text{max}}$ or $\Delta f_{\text{shift}}^{\text{max}}$ the scheme's robustness is improved.

In order to *catch* 64 valid samples of the CP-protected Ultraspherical burst, choosing the right time instant to trigger the FFT is vital. Assuming a maximum propagation delay δ_{max} and synchronization error t_{sync} , the CP length needs to be $N_{\text{CP}} = \lceil f_s (\delta_{\text{max}} + t_{\text{sync}}) \rceil$ and the FFT should be triggered after $N_{\text{CP}}/f_s + T_{\text{PHY}}$ seconds, where T_{PHY} is the PHY delay.

C. Burst Detection

We ran exhaustive MATLAB channel simulations for various parameter values: M nodes transmit simultaneously, each picking a random contention subcarrier from the F subcarriers in Table I with the respective *MCBC* spacing Φ . Each node was subjected to -174 dBm/Hz typical receiver noise, random frequency shift in the range $[-238, 238]$ kHz and random time shift in the range $[0, 4]$ μs , which includes values of $t_{\text{shift}} > \frac{1}{2\Delta f_{\text{shift}}^{\text{max}}}$ and show that they are not problematic. Channel fading and the false-negative detection error are discussed separately in Section III-D and analyzed in Section V.

An important adverse factor for any carrier sense mechanism is external interference, which can be from nodes using the same protocol, or unknown interference. We note that in the DSRC/WAVE context, because the spectrum is licensed the amount of unknown interference is much reduced. To validate the *MCBC* burst detection mechanism, we forced every node to avoid picking the 6th indexed *MCBC* subcarrier, i.e., FFT bins $-8, -3, 7$ and 18 for Φ equals 4, 5, 7 and 9 respectively. This allowed us to evaluate the interference caused by all other subcarriers activated by other nodes, at the frequency of this (inactive) subcarrier and the false-positive detection error. Fig. 9 shows typical FFT responses and we notice that the power level of the inactive (6th) contention subcarrier is well below all other FFT bins from the horizontal axis, even for many simultaneously transmitting nodes ($M > 250$) or narrow frequency spacing ($\Phi = 4$), cases representing worst case scenarios. To implement a detection threshold,

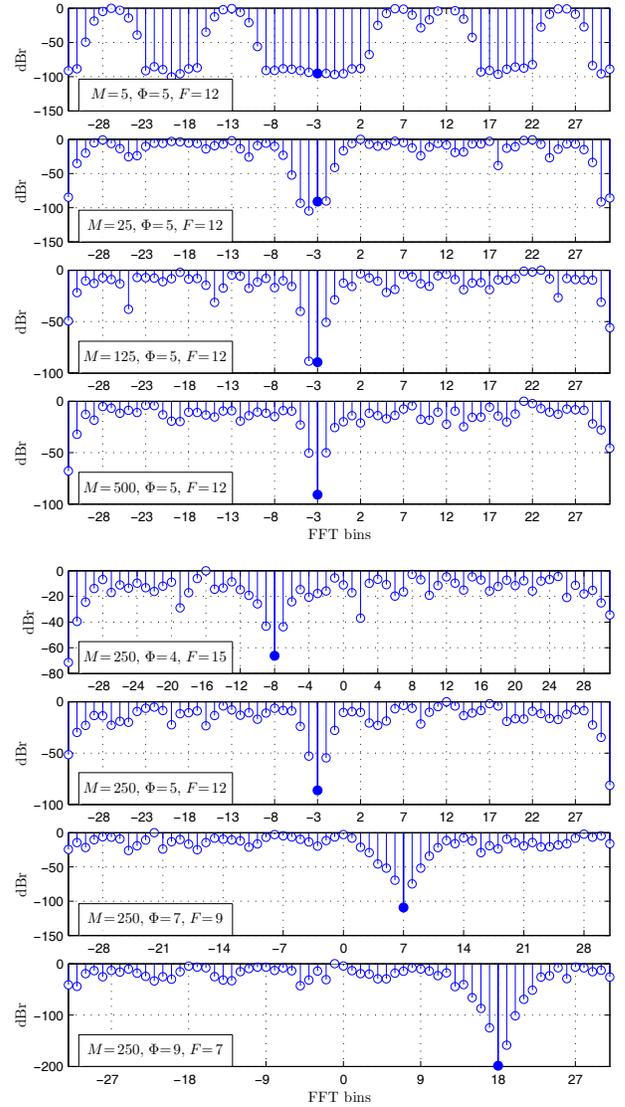


Fig. 9. Normalized response of *MCBC* bursts spaced in frequency with $\Phi \Delta f$ for M simultaneously transmitting nodes, random time shift in the range $[0, 4]$ μs , random frequency shift in the range $[-238, 238]$ kHz and -174 dBm/Hz receiver noise. The 6th *MCBC* subcarrier index (full marker) was intentionally avoided by all nodes to validate the burst detection mechanism. Only the numbered FFT bins are read by *MCBC* during contention.

we need to evaluate the difference between the inactive (6th) subcarrier power and the lowest received power among all active subcarriers (worst case). Fig. 10 shows the Cumulative Distribution Function (CDF) of the two received normalized power values for $M = 250$ simultaneous random bursts. Increasing Φ will increase the distance between the two curves, an expected result since a larger Φ allows using windows with higher side-lobe attenuation. Thus, it is sufficient to investigate only the worst case, $\Phi = 4$. Fig. 11 shows the detection error for various thresholds, revealing a large range of satisfying thresholds for both high and low number of simultaneous bursts ($M = 250$ and $M = 5$). A threshold of -50.4 dB yields a detection error of $4 \cdot 10^{-4}$ and any threshold in the range $[-52, -42]$ gives less than $4 \cdot 10^{-3}$ detection error.

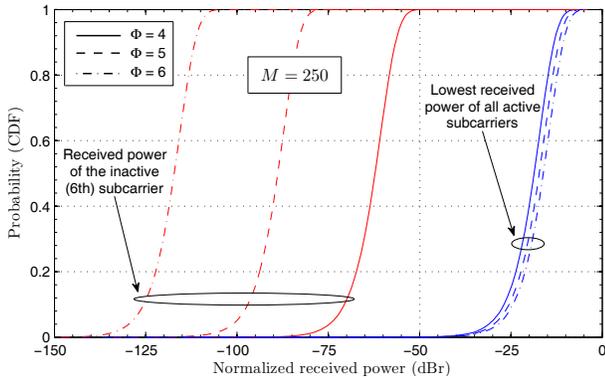


Fig. 10. Received power CDF of the inactive (6th) MCBC subcarrier and of the highest attenuated active MCBC subcarrier.

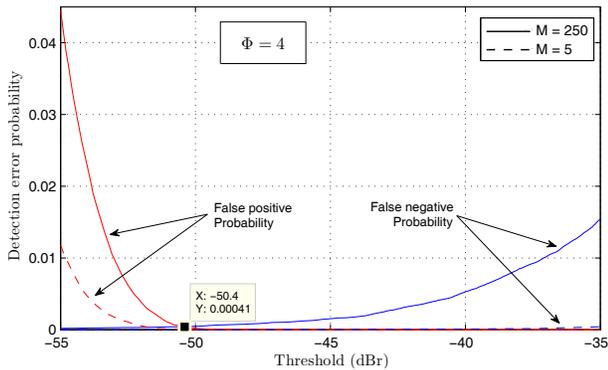


Fig. 11. Detection error probability versus detection threshold.

D. Channel Fading

In realistic environments, particularly in vehicular networks, channel fading represents a serious challenge. For vehicular networks, because of the very high mobility, the channel state can vary greatly both in time and frequency and thus we should consider fast fading but also frequency selective fading, i.e., different subcarriers have different fading characteristics. Selective fading can reduce the received power of an active contention subcarrier at the receiver below the detection threshold thus causing a false negative for the MAC algorithm. Evidently, this adversely affects the MAC contention algorithm, increasing the frame collision probability.

1) **Ricean Fading:** In vehicular networks, we can assume that there may exist both a direct line of sight (LOS) and multipath propagation between transceivers. The Rice distribution [22] fits those characteristics and is discussed here as an example; the analytic model presented in Section V is agnostic to the fading model so any model that is deemed appropriate (e.g. [23]) can be used, following the description below.

Let us consider the general case of a node-pair i, j , where node i transmits and node j senses the channel to determine if it is idle or busy. We note with $Y_{i,j}$ the random variable representing the instantaneous received power at node j when node i is transmitting. For simplicity, we consider links between nodes to be statistically similar, described by an average received power, $\bar{y} = \mathbb{E}[Y_{i,j}]$ for all i, j , where $\mathbb{E}[\cdot]$ denotes the expectation operation. The received signal power probability

density function (PDF) of the Rice distribution is [22]:

$$f_Y(y_{i,j}) = \frac{1}{2\sigma^2} \exp\left(-\frac{y_{i,j} + s^2}{\sigma^2}\right) I_0\left(\frac{\sqrt{y_{i,j}} s}{\sigma^2}\right) \quad (7)$$

where I_0 is the modified Bessel function of the first kind with order zero [24] and s^2 and $2\sigma^2$ are the average received power of the LOS component and of the multipath component respectively. The overall average received power is:

$$\bar{y} = \mathbb{E}[Y_{i,j}] = \int_0^\infty y_{i,j} f_Y(y_{i,j}) dy_{i,j} = s^2 + 2\sigma^2. \quad (8)$$

Node j can correctly identify a busy channel if $Y_{i,j}$ is at least equal to a predefined detection threshold y_0 . However, in the event that $Y_{i,j} < y_0$ then node j incorrectly detects the channel as idle and yields a *false negative*. The probability ξ that node i will not be overheard by node j is called *outage probability* and is obtained from the CDF:

$$\xi = \Pr\{Y_{i,j} < y_0\} = \int_0^{y_0} f_Y(y_{i,j}) dy_{i,j} = 1 - Q_0\left(\frac{s}{\sigma}, \frac{\sqrt{y_0}}{\sigma}\right), \quad (9)$$

where Q_0 is the Marcum Q function [24].

2) **Cumulative Fading:** In much of the MAC literature, a binary hidden/visible model is employed regardless of the number of simultaneous transmissions. This is unrealistic since for $n > 1$ simultaneously transmitting nodes, the transmitted signals cumulate at the receiver, and so, individual hidden nodes can still yield a received signal power above the detection threshold if more of them transmit simultaneously. It is thus appropriate to assume a proper distribution to model this case. In our Ricean example, for n simultaneous transmissions the received signal power follows a non-central chi-square distribution with $2n$ degrees of freedom and non-centrality parameter s^2 [22]. The outage probability ξ_n with respect to the same threshold y_0 is given by its CDF:

$$\xi_n = 1 - Q_n\left(\frac{s}{\sigma}, \frac{\sqrt{y_0}}{\sigma}\right), \quad (10)$$

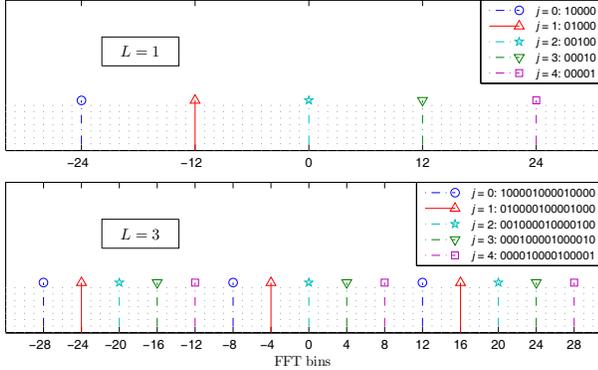
where Q_n is the generalized Marcum Q function [24]. We can also describe Ricean fading using the *fading factor* $K = s^2/2\sigma^2$ (the ratio between the LOS and multipath average received power), which is a measure of the fading severity: when $K = 0$ it reverts to Rayleigh fading (no LOS component) and when $K \rightarrow \infty$ there is no fading (no multipath component). Hence, we have:

$$\xi_n = 1 - Q_n\left(\sqrt{2K}, \sqrt{2(K+1)\frac{y_0}{y}}\right). \quad (11)$$

3) **Fading Resilience:** MCBC has an inherent degree of resilience to selective fading thanks to the random selection from the F subcarriers and the multiple rounds, similar to frequency hopping schemes. However, it can benefit from additional redundancy similar to error-correction schemes. A straightforward yet effective method for improving resilience to selective fading is for a node to burst on 2 or more subcarriers simultaneously, instead of only 1, by repeating (copying) its chosen subcarrier across the spectrum with distance $k\Phi F$, where k is an integer. Thus, for L copies, a node transmits simultaneous bursts on subcarrier indexes $j, j + \Phi F, j + 2\Phi F,$

TABLE II
POSSIBLE ACTIVE SUBCARRIERS FOR $F = 5$

j	No repetition ($L = 1$)	Repeated twice ($L = 3$)
0	1 0 0 0 0	1 0 0 0 0 1 0 0 0 0 1 0 0 0 0
1	0 1 0 0 0	0 1 0 0 0 0 1 0 0 0 0 1 0 0 0 0
2	0 0 1 0 0	0 0 1 0 0 0 0 1 0 0 0 0 1 0 0 0
3	0 0 0 1 0	0 0 0 1 0 0 0 0 1 0 0 0 0 1 0 0 0
4	0 0 0 0 1	0 0 0 0 1 0 0 0 0 1 0 0 0 0 1



$\dots, j + (L-1)\Phi F$, with j randomly selected from $0, \dots, F-1$. A binary representation is possible and is depicted in Table II where a randomly selected row represents the binary bitmask of the activated subcarrier(s). The corresponding symbolic spectrum is shown in the figure underneath. We see that for any j , the corresponding L subcarriers span a large spectrum bandwidth. Also, different j values do not have overlapping active subcarriers (bits of 1) which allows to uniquely identify a selected row for up to $L-1$ faded subcarriers (errors). These two frequency diversity features offer very good selective fading resilience as shown in Section VI.

The MAC can perform bitwise *AND* operations against the received bitmask from the PHY to find the winning j index. By caching all $2^{F \cdot L}$ bitmasks into look-up tables, bitwise operations are not needed and hardware complexity and delay remain unaffected by this fading resilience measure. Also, no hardware changes to the CSA block are needed. The CSA block will simply hold different 64-point window samples; for a choice $j = 0, \dots, F-1$ these are (see (5)):

$$s_{j,k} = w[k] \sum_{l=1}^L e^{i2\pi k(j+l\Phi F)\Delta f} \quad (12)$$

It is important to note that the total number of contention subcarriers with repetition must still be lower than the maximum available contention subcarriers from (4), i.e., $F \cdot L \leq F_{\max}$ must hold, e.g. a value $F_{\max} = 15$ allows $F = 5$ different subcarriers (rows) copied $L = 3$ times, or any combination of F and L which satisfies the above inequality. Obviously, this is a compromise between fading resilience (higher L) and frequency diversity (higher F).

Usually, fading levels for subcarriers spaced with Δf are correlated. In vehicular networks however, due to high frequency selectivity, we can assume that correlation between subcarriers spaced with more than $\Phi \Delta f$ is very low. Hence, when n nodes transmit simultaneously on the same subcarrier and a repetition length L is used, the outage probability of the cumulative contention burst is reduced to ξ_n^L .



Fig. 12. *MCBC* testbed: IQ modulator with Tx antenna (left), oscilloscope showing the Ultraspherical signal (top center), spectrum analyzer showing its spectrum (bottom center), logic analyzer showing synchronized strobes in its upper part and the random contention signals in its lower part (top right), FPGA platform and independent quartz oscillators board (bottom right).

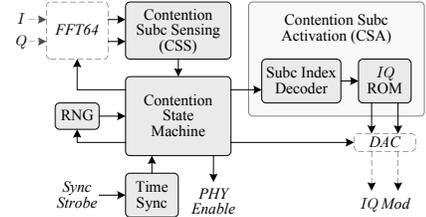


Fig. 13. Block diagram of one *MCBC* core. The dashed blocks are conventional blocks from the OFDM PHY.

IV. HARDWARE TESTBED

To validate the algorithm and its timings, the time synchronization and all concepts described in Sections II and III, we built a hardware testbed, shown in Fig. 12, based on an Altera EP2C35F FPGA platform board. We note that the goal was not to build fully working prototypes since a real-world evaluation of several tens or hundreds of simultaneous transmitting nodes requires as many prototypes. A testbed emulating a real-world network of nodes allows validating many of the performance claims that would otherwise require numerous prototypes. The testbed consists partly of multiple *MCBC* cores implemented on the same chip and a separate analogue board with independent free running quartz oscillators from different manufacturers so that each core is independently clocked to emulate real nodes. Having cores independently clocked also allows to investigate frequency shift (see Section III-A) and test the time sync mechanism detailed below.

The *MCBC* block diagram is shown in Fig. 13. When receiving a burst, the CSS block reads the FFT outputs, extracts the signal strength of the F contention subcarriers (phase information along with the signal strength of the other $64 - F$ subcarriers is discarded), compares the F values with the predefined detection threshold (see Section III-C) and ultimately passes an F -bit bitmask to the main state machine. The bitmask is 1 for subcarriers where the signal strength was above the threshold and 0 for the ones below the threshold. The state machine runs the contention algorithm, triggers the FFT and the other conventional PHY blocks (see Fig. 4), controls the RNG and synchronizes all signals on the pulse received from the time synchronizing block. The RNG was implemented using a combination of linear feedback and cellular automata shift registers (LFSR+CASR) [25]; this provides better bit independence, permitting to generate long sequences once per contention session and to use groups of

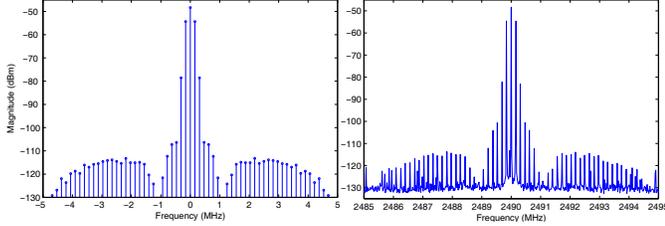


Fig. 14. Spectrum for $\Phi = 4$ after removal of CP, spanning 10 MHz and 64 FFT bins. Theoretical (left) vs Measurement (right).

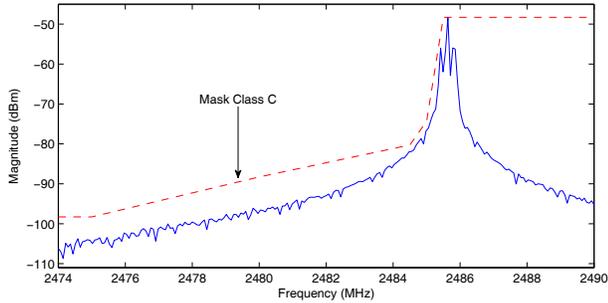


Fig. 15. Spectrum of worst case *MCBC* burst with $\Phi = 4$ with cyclic prefix, at 2.49 GHz center frequency. A good match is observed with the equivalent theoretical spectrum shown in Fig. 7.

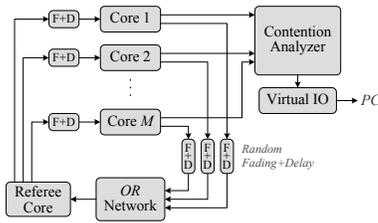


Fig. 16. Block diagram of the multi-core testbed emulating a real-world network of M nodes. Cores are clocked with M independent external quartz oscillators from different manufacturers.

bits to minimize delay. When transmitting a burst, the state machine generates a random F -bit contention subcarrier bitmask which is mapped by a decoder to appropriate enable signals, driving the I and Q ROMs in the CSA block; these hold the Ultraspherical samples $s_{j,k}$ from (12) (see Section III-D3) and feed the PHY's IQ modulator.

First, to validate the *MCBC* bursts spectral features and the detection mechanism from Sections III-C and III-D, an *MCBC* core was connected to the IQ modulator, tuned at 2.49 GHz (hardware limitation), using 12-bit D/A converters available on the FPGA board. The transmit spectrum of the bursts was observed using a spectrum analyzer in both close (same room) and far+occluded (multiple walls) locations from the transmitter. Fig. 14 shows the received spectrum of a 64 samples *MCBC* burst (i.e., after removal of CP) and a good match between the theoretical and measured spectrum can be observed. Fig. 15 shows the measured transmit spectrum of the worst case *MCBC* burst with CP as detailed in Section III-B, which matches the theoretical spectrum from Fig. 7 and meets the transmit mask requirements.

Second, to measure timings, validate the time sync mecha-

nism and take contention measurements, up to 16 cores were implemented (the EP2C35F has only 16 clock buffers) to emulate a real network, shown in Fig. 16. A block emulating a real-world *OR* channel was implemented to add subcarrier energy levels and to feed a referee core. The channel inputs and outputs of each core pass through Fading+Delay blocks that apply random Ricean fading and random time delay. Each core's state machine feeds a contention analyzer block and a Virtual IO that taps directly into the internal FPGA signals interfaces with the PC.

Time Synchronization

IEEE 802.11p [1] specifies Timing Advertisement frames which are beacon-like frames advertising the estimate of the difference to an external master clock and an estimate of the standard deviation. The implementer is free to choose the external time source and the synchronization mechanism. As suggested in [1], in vehicular networks it is sensible to assume the existence of a GPS receiver that also provides a timing strobe, typically 1 pulse per second (pps). We note the existence of inexpensive off-the-shelf GPS timing products with good accuracy (below 50 ns [26], [27]) and good holdover ability (below 200 ns/h drift after loss of GPS steering [28]). For indoors networks, nodes can synchronize on the AP beacons using the same mechanism described below.

In our implementation, we did not use an accurate master clock, time sync between cores being implemented using a low quality ± 100 ppm oscillator, down-converted to output 1, 2, 5 and 10 pps strobes. The oscillators clocking the $M + 1$ cores were ± 20 ppm [1]. Without synchronization, we measured a maximum $34 \mu\text{s/s}$ core-to-core relative drift and $93 \mu\text{s/s}$ core-to-master clock relative drift, which is far too large. As the master clock is also drifting, this calls for an adaptive sync scheme. We implemented a linear prediction design having a low hardware footprint, based on parallel counters and look-up tables, to estimate the number of overrunning or underrunning cycles and to maintain a deviation estimate. Noting with f_c the local clock frequency and Δf_c the measured drift from the master clock, then the number of cycles to trigger an adjustment of the synchronized pulse is:

$$l_{\text{adj}} = \left\lfloor \frac{f_c + \Delta f_c}{\Delta f_c} \right\rfloor. \quad (13)$$

The remainders are optionally used for deviation estimation and fine adjustment. Thus, the scheme can follow the master clock and the maximum drift from the master clock is:

$$t_{\text{drift}}^{\text{max}} = l_{\text{adj}} \left| \frac{1}{f_c} - \frac{1}{f_c + \Delta f_c} \right| \leq \frac{1}{f_c}. \quad (14)$$

Fig. 17 shows measurements for a 1 pps master strobe and 3 synchronized *MCBC* cores where the oscilloscope persistence reveals maximum 100 ns drift from the master strobe. For 30 s training period, the holdover drift (i.e., after loss of master clock signal) was measured below $1.2 \mu\text{s/h}$. Considering that the PHY and MAC processing delays are fixed, this sync method would also work when using the Time Advertisement frame as master clock, even if it is not periodically received (due to collisions or fading) as the frame contains the time

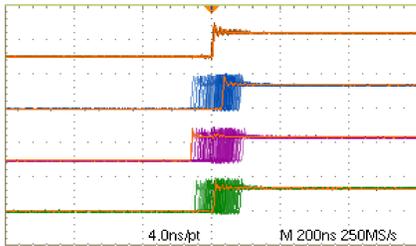


Fig. 17. Persistent oscilloscope showing synchronized slot strobes (100 ns sync accuracy) of 100 kHz derived from 100 ppm independent oscillators. The sync strobe (top) is 1 pps derived from a 10 MHz 100 ppm oscillator.

info; the disadvantage over using a GPS time strobe is that nodes would receive the frame at different time instants due to propagation delay so the overall sync time-guard increases.

The FPGA footprint of an *MCBC* core is 520 logic cells (excluding IQ ROMs), a big part being the divider from (13). The maximum operating frequency was 368 MHz. The MAC processing delay, i.e., the time elapsed since the contention round strobe until the full reply is ready for the PHY, was less than 50 ns when using a 20 MHz local clock.

V. ANALYTIC MODEL

As packet collisions can drastically hinder performance, it is important to analyze the contention scheme's performance, i.e., the probability that there is a unique winning node after R rounds, taking into account cumulative channel fading. The contention mechanism does not depend on past attempts or packet error (which is out of the scope of this paper), so the equations presented here hold for all cases. Moreover, since contention sessions are independent, various packet arrival distributions, e.g. exponential, only affect the number of initial contender nodes M . Thus, we analyze saturation arrival without loss of generality since M is variable. We consider the more challenging case of non-reciprocal channels, i.e., fading characteristics on the forward channel (nodes to referee) differ from those on the feedback channel (referee to nodes), and note the feedback channel outage probability with ξ_{ref} , given by (9). For reciprocal channels $\xi_{\text{ref}} = \xi_1$. Due to space limitations, in the analytical model nodes rely solely on the referee reply, hence they do not listen for bursts during contention slots (the optional blocks in Fig. 2 are skipped) and we assume that a referee is present.

It is worth mentioning that the model presented here applies to a large extent to other systems that use individual OFDM subcarriers for higher level applications, e.g. SMACK [29] or Spectrum Pooling [30], and for general leader election schemes that use parallel random selections.

A. Success Probability

The key probabilities used in the analysis are listed in Table III. Following the protocol's recursive nature and correctness (Section II-A), i.e., *a*) there is always at least one winner in a round and *b*) a single round winner will also be the single overall contention winner, then $P_{c,r}$ is:

$$P_{c,r} = \mu_{c,1,r} + \sum_{w=2}^c \mu_{c,w,r} P_{w,r-1}, \quad (15)$$

TABLE III
KEY PROBABILITIES USED IN THE ANALYSIS.

p_r	the <i>flipping probability</i> in round r (see Section II-A)
$q_{r,f}$	the <i>choice probability</i> in round r or the probability to pick sub-carrier index f from the predefined F contention subcarriers; when uniform then $q_{r,f} = 1/F$ (Section II-B)
$\xi_n^L, \xi_{\text{ref}}^L$	the outage probability due to channel fading for n cumulative simultaneous bursts and for the referee burst respectively, for a repetition length L (see Sections III-D2 and III-D3)
$P_{c,r}$	the <i>unique-winner probability</i> , i.e., the probability that there is a single winner after r rounds, given there were c contender nodes at the beginning of the first round. For simplicity, rounds are assumed to decrease from r to 1, where $r = 1, \dots, R$
$\mu_{c,w,r}$	the <i>condenters-winners probability</i> , i.e., the probability that round r finishes with w winning nodes, given that there were c contender nodes at the beginning of the round
$\nu_{n,w,f}$	the <i>nominees-winners probability</i> , i.e., the probability that, given there were n nominee nodes in a round (n nodes flipped yes), there are w winning nodes and the winning subcarrier index is f , where $f = 1, \dots, F$

which is a recursive form showing there is a single winner when either the current round r yields a single winner out of c contender nodes or when the next round $r-1$ yields a single winner out of the w winners of the current round. The obvious recursion stop condition is $P_{c,1} = \mu_{c,1,1}$.

Naturally, the *success probability* P_s after R rounds in a network of M active nodes is:

$$P_s = P_{c,r}. \quad (16)$$

1) **Contenders-Winners Probability** $\mu_{c,w,r}$: We know that if a node does not hear the referee reply, it will select itself as winner. In ideal channels, this can only happen if all contender nodes flip *no*, so the referee will not reply. In fading channels however, a node who flips *no* may simply be unable to hear the referee due to fading and thus select itself as winner, which is undesirable. So, depending on the number of nodes n who flip *yes*, the w winning nodes are split into w' nodes who flipped *yes* and $w-w'$ who flipped *no*. This is a hypergeometric distribution [31] with support $w' \in \mathcal{H} = \{\max(1, n - (c - w)), \dots, \min(n, w)\}$. Thus, the probability $\mu_{c,w,r}$ that round r finishes with $w < c$ winners, given there were c contender nodes at the beginning of the round, is:

$$\mu_{c,w,r} = \sum_{f=1}^F \sum_{n=1}^c \binom{c}{n} p_r^n (1-p_r)^{c-n} \sum_{w' \in \mathcal{H}} \nu_{n,w',f} \cdot \binom{c-n}{w-w'} (\xi_{\text{ref}}^L)^{w-w'} (1-\xi_{\text{ref}}^L)^{c-n-(w-w')}, \quad (17)$$

which holds only for $w < c$ as it omits the case when all nodes flip *yes*. As *MCBC* ensures at least one winner in any round, the normalization condition $\sum_{w=1}^c \mu_{c,w,r} = 1$ holds, so:

$$\mu_{c,c,r} = 1 - \sum_{w=1}^{c-1} \mu_{c,w,r}. \quad (18)$$

2) **Nominees-Winners Probability** $\nu_{n,w,f}$: We employ an *augmented* multinomial distribution to compute $\nu_{n,w,f}$. Let X_j be the random variable representing the number of nodes that picked subcarrier j in round r , where $j = 1, \dots, F$. The

standard multinomial distribution [31] gives:

$$\begin{aligned} \Pr\{X_1 = n_1, X_2 = n_2, \dots, X_F = n_F\} &= \\ &= \binom{n}{n_1, n_2, \dots, n_F} q_{r,1}^{n_1} q_{r,2}^{n_2} \dots q_{r,F}^{n_F}, \end{aligned} \quad (19)$$

where $\sum_{j=1}^F n_j = n$ and $\binom{n}{n_1, n_2, \dots, n_F}$ is the multinomial coefficient defined as:

$$\binom{n}{n_1, n_2, \dots, n_F} \triangleq \binom{n}{n_1} \binom{n-n_1}{n_2} \binom{n-n_1-n_2}{n_3} \dots \binom{n_F}{n_F} = \frac{n!}{n_1! n_2! \dots n_F!}$$

which is the number of ways to split n nodes into F constrained groups. We can rewrite (19) as:

$$\Pr\{X_1 = n_1, X_2 = n_2, \dots, X_F = n_F\} = \prod_{j=1}^F Q_j, \quad (20)$$

where Q_j is defined as:

$$Q_j = \left(\sum_{k=j}^F n_k \right) q_{r,j}^{n_j} = \binom{n - \sum_{k=1}^{j-1} n_k}{n_j} q_{r,j}^{n_j} \quad (21)$$

With fading, a node may also win if it does not hear the referee reply. We can thus define the random variable W_j representing the number of winners that picked subcarrier j out of n_j nominees. This follows a binomial distribution with respect to n_j and ξ_{ref}^L :

$$B_j = \Pr\{W_j = w_j\} = \binom{n_j}{w_j} (\xi_{\text{ref}}^L)^{w_j} (1 - \xi_{\text{ref}}^L)^{n_j - w_j}, \quad (22)$$

where $\sum_{j=1}^F w_j = w$ and $w_j \leq n_j$. Given n nominees and w winners, if the highest indexed subcarrier that is heard by the referee is f then all the nodes who picked f will win regardless of whether they can hear the referee reply or not. Hence, in this case we have $n_f = w_f = w' \geq 1$ and (X_j, W_j) follows an *augmented* multinomial distribution given by:

$$\begin{aligned} \Pr\{X_1 = n_1, W_1 = w_1, \dots, X_f = W_f = w', \dots \\ \dots, X_F = n_F, W_F = w_F\} = \prod_{j=1}^F A_{f,j} B_j, \end{aligned} \quad (23)$$

where $\sum_{j=1}^F w_j = w$, $\sum_{j=1}^F n_j = n$, $w_j \leq n_j$ and $A_{f,j}$ is the augmented version of Q_j :

$$A_{f,j} = \begin{cases} \binom{n-w'-\sum_{k=1}^{j-1} n_k}{n_j} q_{r,j}^{n_j}, & j < f \\ \binom{n}{w'} q_{r,j}^{w'} (1 - \xi_{\text{ref}}^L), & j = f \\ \binom{n-w'-\sum_{k=1, k \neq f}^{j-1} n_k}{n_j} q_{r,j}^{n_j} \xi_{\text{ref}}^L, & j > f. \end{cases} \quad (24)$$

One can notice that, in order for f to be the winning subcarrier, the cumulated burst from the w' nodes must be heard by the referee (with probability $1 - \xi_{\text{ref}}^L$) and the cumulated bursts from the nodes who picked any subcarrier $j > f$ must be faded at the referee (with probability ξ_{ref}^L).

Now, $\nu_{n,w,f}$ is computed by summing over all values of w' , n_1, \dots, n_{F-1} and w_1, \dots, w_{F-1} in the augmented multinomial context, i.e., $n_f = w_f = w'$, $\sum_{j=1}^F w_j = w$, $\sum_{j=1}^F n_j = n$ and $w_j \leq n_j$:

$$\nu_{n,w,f} = \sum_{w'=1}^w \sum_{\substack{n_1, n_2, \dots, n_{f-1}, \\ n_{f+1}, \dots, n_{F-1}}} \sum_{\substack{w_1, w_2, \dots, w_{f-1}, \\ w_{f+1}, \dots, w_{F-1}}} \prod_{j=1}^F A_{f,j} B_j. \quad (25)$$

Using the following identities and notations:

- $\sum_{w_j=0}^w \binom{n_j}{w_j} \binom{n-n_j}{w-w_j} = \binom{n}{w}$, $\forall w_j \leq w \leq n$ (Vandermonde)
- $\binom{n}{w'} \binom{n-w'}{w-w'} = \binom{n}{w} \binom{w}{w'}$, $\forall w' \leq w \leq n$
- $\Omega_{f,j} = \sum_{k=1, k \neq f}^j n_k$ (where $\Omega_{f,F} = n - w'$ and $\Omega_{f,0} = 0$)
- $\Upsilon_{f,j} = \sum_{k=1, k \neq f}^j w_k$ (where $\Upsilon_{f,F} = w - w'$ and $\Upsilon_{f,0} = 0$)

we can simplify (25) such that:

$$\begin{aligned} \nu_{n,w,f} &= \sum_{w'=1}^w \binom{n}{w'} q_{r,f}^{w'} (1 - \xi_{\text{ref}}^L) \sum_{\substack{n_1, n_2, \dots, n_{f-1}, \\ n_{f+1}, \dots, n_{F-1}}} \prod_{\substack{j=1 \\ j \neq f}}^F A_{f,j} \cdot \\ &\cdot \sum_{w_1=0}^{w-w'} B_1 \sum_{w_2=0}^{w-w'-\Upsilon_{f,1}} B_2 \dots \sum_{w_{f-1}=0}^{w-w'-\Upsilon_{f,f-2}} B_{f-1} \cdot \\ &\cdot \sum_{w_{f+1}=0}^{w-w'-\Upsilon_{f,f}} B_{f+1} \sum_{w_{f+2}=0}^{w-w'-\Upsilon_{f,f+1}} B_{f+2} \dots \underbrace{\sum_{w_{F-1}=0}^{w-w'-\Upsilon_{f,F-2}} B_{F-1} \cdot B_F}_{\substack{(n-w'-\Omega_{f,F-2}) \cdot (\xi_{\text{ref}}^L)^{w-w'-\Upsilon_{f,F-2}} \\ \cdot (1-\xi_{\text{ref}}^L)^{n-\Omega_{f,F-2}-(w-\Upsilon_{f,F-2})}}}, \\ &\cdot \underbrace{\sum_{w_{f+1}=0}^{w-w'-\Upsilon_{f,f+1}} B_{f+1} \cdot B_{f+2} \dots B_{F-1} \cdot B_F}_{\substack{(n-w'-\Omega_{f,f+1}) \cdot (\xi_{\text{ref}}^L)^{w-w'-\Upsilon_{f,f+1}} \\ \cdot (1-\xi_{\text{ref}}^L)^{n-\Omega_{f,f+1}-(w-\Upsilon_{f,f+1})}}}. \end{aligned} \quad (26)$$

which, after more algebraic manipulations, becomes:

$$\begin{aligned} \nu_{n,w,f} &= \binom{n}{w} (1 - \xi_{\text{ref}}^L)^{n-w} \sum_{w'=1}^w \binom{w}{w'} q_{r,f}^{w'} (1 - \xi_{\text{ref}}^L) \cdot \\ &\cdot (\xi_{\text{ref}}^L)^{w-w'} \cdot \sum_{\substack{n_1, n_2, \dots, n_{f-1}, \\ n_{f+1}, \dots, n_{F-1}}} \prod_{\substack{j=1 \\ j \neq f}}^F A_{f,j}. \end{aligned} \quad (27)$$

The number of sums in (27) is variable, depending on F , due to the multinomial distribution. The recursive function:

$$\begin{aligned} \psi_{n,f,j} &= \\ &= \begin{cases} \sum_{k=0}^n \binom{n}{k} q_{r,j}^k \psi_{n-k,f,j+1}, & f < F, j = 1, \dots, f-1 \\ \sum_{k=0}^n \binom{n}{k} q_{r,j+1}^k \xi_{\text{ref}}^L \psi_{n-k,f,j+1}, & f < F, j = f, \dots, F-2 \\ q_{r,j+1}^n \xi_{\text{ref}}^L, & f < F, j = F-1 \\ q_{r,j}^n, & f = F, j = F-1. \end{cases} \end{aligned} \quad (28)$$

implements the variable multinomial sums in (27) and an expression suitable for numerical implementation is:

$$\begin{aligned} \nu_{n,w,f} &= \binom{n}{w} (1 - \xi_{\text{ref}}^L)^{n-w} \sum_{w'=1}^w \binom{w}{w'} q_{r,f}^{w'} (1 - \xi_{\text{ref}}^L) \cdot \\ &\cdot (\xi_{\text{ref}}^L)^{w-w'} \psi_{n-w',f,1} \end{aligned} \quad (29)$$

which, finally, plugs into (17) and in turn into (15) for the complete solution.

3) **Ideal Case:** In ideal channels, $\xi_{\text{ref}} = \xi_n = 0, \forall n$. For a uniform subcarrier choice³, i.e., $q_{r,f} = 1/F$, then $\psi_{n,f,1} = (f-1)^n/F^n$ which gives $\nu_{n,w,f} = \binom{n}{w} (f-1)^{n-w}/F^n$ so:

$$\mu_{c,w,r} = \begin{cases} \frac{1}{F^c} \binom{c}{w} p_r^w \sum_{f=1}^F (F - p_r f)^{c-w}, & w < c \\ (1 - p_r)^c + F \left(\frac{p_r}{F}\right)^c, & w = c. \end{cases} \quad (30)$$

B. Throughput and Delay

Let S be the normalized network throughput defined as the fraction of time the channel is used to successfully transmit payload bits. Contention sessions are independent so S is directly proportional to P_s and since successful transmissions and collisions are equal in duration, we have:

$$S = \frac{P_s N_d/B}{T}, \quad (31)$$

where P_s is given by (16), N_d is the data frame size in bits, B is the PHY bit rate in bits/s and T is the duration of one full contention+transmission (see Fig. 1) given by:

$$T = \underbrace{2T_{\text{slot}} \cdot R}_{T_r} + T_d + 2 \cdot \text{SIFS} + T_{\text{ACK}}. \quad (32)$$

T_{slot} is the duration of a contention or feedback slot, T_d (T_{ACK}) is the time the PHY uses to transmit a data (acknowledgement) frame and SIFS is the Short Inter-Frame Space interval. The propagation delay, δ , is included in SIFS and T_{slot} . SIFS is the same as in the standard [1] and T_{slot} is:

$$T_{\text{slot}} = \frac{64}{f_s} + 2(\delta + t_{\text{sync}}) + T_{\text{MAC}}^{\text{cont}} + T_{\text{RXTx}}^{\text{cont}}, \quad (33)$$

$T_{\text{MAC}}^{\text{cont}}$ and $T_{\text{RXTx}}^{\text{cont}}$ are the MAC processing delay and RxTx turnaround time during contention (see also Section III-B). T_d (and, similarly, T_{ACK}) is defined [1] as:

$$T_d = T_{\text{pre}} + \frac{N_{\text{sgn}}}{B_{\text{sgn}}} + T_{\text{sym}} \left[\frac{N_{\text{srv}} + N_d + N_{\text{MAC}} + N_{\text{tail}}}{T_{\text{sym}} \cdot B} \right], \quad (34)$$

where we identify the duration of the PHY preamble (T_{pre}), size and bit rate of the signal field ($N_{\text{sgn}}, B_{\text{sgn}}$), duration of one OFDM symbol (T_{sym}), size of the service field (N_{srv}), MAC header size (N_{MAC}) and number of tail bits (N_{tail}).

The network delay is defined as the time elapsed since the scheduling of the data frame for transmission until (and including) its successful transmission:

$$D = T(\mathbb{E}[N_c] + 1) = T/P_s, \quad (35)$$

where $\mathbb{E}[N_c] = 1/P_s - 1$ is the average number of collisions between two successful transmissions.

As shown by (31) and (35), maximizing P_s is crucial to obtain a high throughput and low delay.

³For a geometric choice (see Section II-B) then $\psi_{n,f,1} = \left(\frac{1-\alpha_r^{f-1}}{1-\alpha_r^F}\right)^n$ and $\nu_{n,w,f} = \binom{n}{w} \left(\frac{1-\alpha_r^{f-1}}{1-\alpha_r^F}\right) \left(\frac{(1-\alpha_r)\alpha_r^{f-1}}{1-\alpha_r^F}\right)^{n-w}$ and $\mu_{c,w,r}$ is computed using (17) and (18).

TABLE IV
SIMULATION PARAMETERS.

Parameter	802.11p	MCBC
MAC payload (bits)	8184	8184
PHY bit-rate B (Mbits/s)	12	12
δ {min,max} ⁴ (μ s)	{0.5, 3}	{0.5, 3}
Delay spread t_{spread} (μ s)	< 1.6	< 1.6
T_{RXTx} (μ s)	< 2	< 1.2
T_{MAC} (μ s)	< 2	< 0.4
Sampling freq f_s (MHz)	10	10
$\Delta f_{\text{shift}}^{\text{max}}$ (MHz)	0.238	0.238
Sgn-Thr Ratio $\gamma_0/\bar{\gamma}$ (dB)	3	3
Ricean K {min,max}	{0.1, 50}	{0.1, 50}
CW {min,max}	{15, 1023}	-
Retry limit {short,long}	{4, 7}	-
Sync error t_{sync} (μ s)	-	< 0.4
Φ	-	4
F	-	15
L	-	1
R	-	3
$[p_1, p_2, p_3]$	-	[0.12, 0.77, 0.86]
$[\alpha_1, \alpha_2, \alpha_3]$	-	[0.60, 0.90, 0.98]
T_{slot} {min,max} ⁴ (μ s)	{13, 18}	{11, 15}

VI. RESULTS AND DISCUSSION

To evaluate *MCBC* we obtained analytic results and performed hardware measurements and simulations in comparison with IEEE 802.11p, considering all PHY challenges from Section III. Unless otherwise stated, the parameters used are those in Table IV, all others being the same for both protocols as the ones in [1] for the 10 MHz OFDM channel.

The received signal power for each link (IEEE 802.11p) and each link and each subcarrier (*MCBC*) was independently subjected to fast Ricean fading considering the cumulative effect, as described in Section III-D2. For a fair comparison, we used equivalent carrier sense thresholds for the two protocols⁵ using an average power of 3 dB above the detection threshold and considering a range of the Ricean fading factor $0.1 \leq K \leq 50$. This allowed us to simulate scenarios varying from deep fading (high multipath \leftrightarrow dense buildings) to moderate fading (comparable multipath and LOS \leftrightarrow sparse buildings) to almost ideal channels (mostly LOS \leftrightarrow highways, open roads), giving a wide range of outage probability.

By varying the fading levels in time, i.e., during the same measurement and simulation run, we effectively obtained variable topologies and random out-of-range conditions (links between nodes break or hold depending on the resulted received signal strength) to simulate networks where nodes can go in and out of range both randomly and rapidly – an expected scenario in vehicular networks.

Simulations and measurements matched the analytical expressions very well as shown in Fig. 18. The wave shape of the top curves in Fig. 18 (left) is due to the usage of multiple contention rounds ($R > 1$) and so the p values can shift the position of the peaks. However, using $R \geq 3$ rounds and an

⁴The IEEE 802.11-2007 standard [2] considers a 0.5 μ s propagation delay. However, WAVE aims for up to 1000 meters range, i.e., $\delta \simeq 3.33 \mu$ s, in which case T_{slot} is increased accordingly. For *MCBC*, T_{slot} is larger than given by (33) to include delay spread, but that is not a concern during contention as *MCBC* transmits only single bursts, not successive OFDM symbols.

⁵Here we refer to the fact that IEEE 802.11 considers the entire channel bandwidth while *MCBC* considers individual subcarriers.

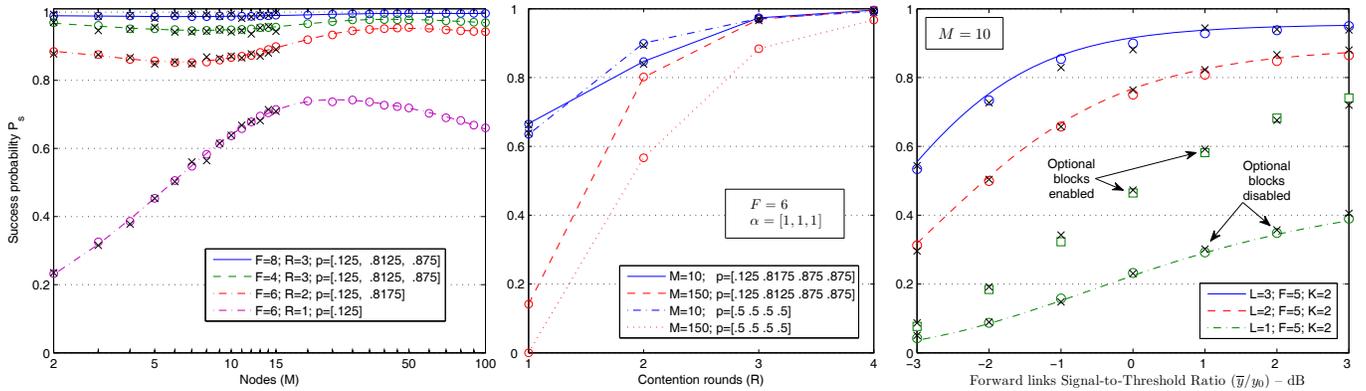


Fig. 18. *MCBC* success probability P_s ; Analysis vs. measurements vs. simulations: variable number of nodes in ideal channel (left), variable number of rounds in ideal channel (center) and variable fading (right). Lines represent theoretical analysis, circle markers represent software simulations and crosshair markers show hardware measurements. The square markers (right) represent results when the optional blocks of the algorithm (see Fig. 2) are enabled.

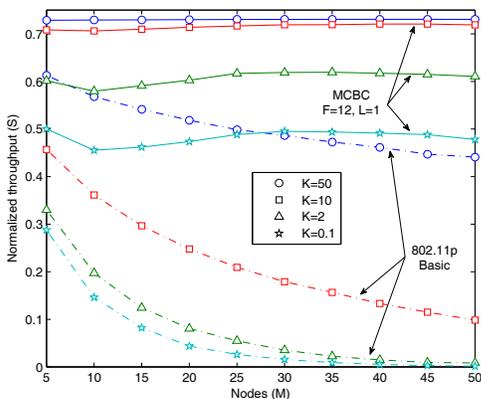


Fig. 19. *MCBC* without fading protection ($L = 1$) and IEEE 802.11p in basic access mode.

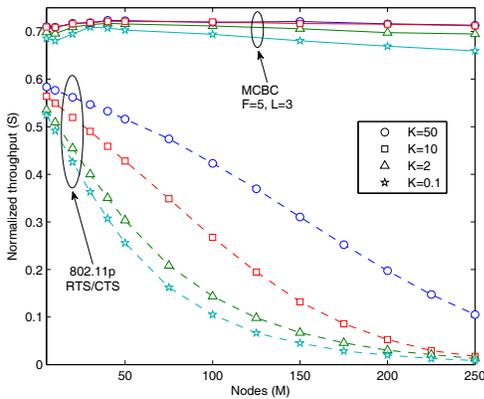


Fig. 20. *MCBC* with fading protection ($L = 3$) and IEEE 802.11p with RTS/CTS enabled.

appropriate F , we notice a high and almost flat performance. Fig. 18 (center) shows that P_s converges very quickly and that using $R > 3$ rounds brings little benefit when an appropriate p vector is used. As described in Section III-D3 and reflected in Fig. 18 (right), using $L > 1$ visibly improves P_s under fading. The small difference between simulations and analytic curves is because the analysis does not consider background noise. Activating the optional processing blocks (see Fig. 2), further

improves performance under fading, as expected, since many nodes that flipped *no* can sense the bursts of those who flipped *yes* and are thus silenced before the referee reply.

Most IEEE 802.11 devices have RTS/CTS disabled by default [7], operating in basic access mode, mainly to save power and minimize overhead and delay in light to medium loaded networks. In fading channels however, Fig. 19, the basic mode throughput drops considerably since collisions can occur anywhere in the data frame as nodes can falsely detect a busy medium as idle. *MCBC* maintains a flat and higher throughput, even without fading protection ($L = 1$). In deep fading (small K), a node has a low chance of overhearing its neighbors in the contention slot and hence relies mostly on the referee's reply in the feedback slot. This in turn is more likely to be faded when activating only one subcarrier and results in a higher collision rate and hence a lower throughput. Still, using $F > 1$ random subcarriers helps combat fading overall: the contention algorithm is oblivious to the identity of the winning node, so as long as at least one subcarrier can be sensed then the number of round winners is exponentially reduced, decreasing collisions.

IEEE 802.11p has a much lower dependence on fading when it uses RTS/CTS as can be seen in Fig. 20. The short RTS frames have a high chance of being accommodated in the backoff windows of hidden nodes so longer data frames can be used to increase efficiency as shown in Fig. 21. When RTS/CTS is disabled (basic mode), increasing the data frame size has an adverse effect as explained above, yet there exists an optimum frame size, an effect shown in Fig. 21. However, the energy cost of an RTS/CTS handshake is high, which is a major reason for disabling RTS/CTS in portable devices. The other downside is that it adds overhead.

MCBC combats fading by bursting on $L > 1$ simultaneous subcarriers as detailed in Section III-D3. Performance improves visibly in fading channels, maintaining both a high and flat performance profile even for many nodes and deep fading as seen in Figs. 20, 22 and 23. We notice from Figs. 21 and 22 that performance in good quality channels (high K) for $L > 1$ is negligibly worse than for $L = 1$. Considering that activating a few more subcarriers for $< 20 \mu\text{s}$ introduces zero

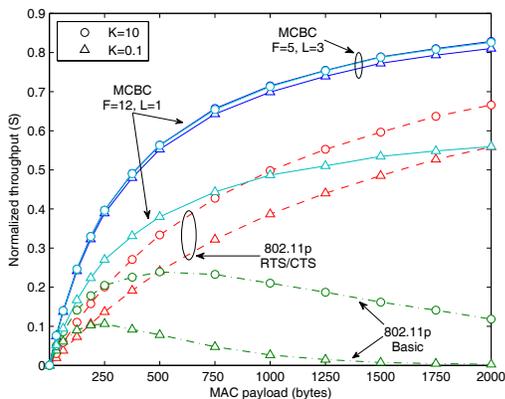


Fig. 21. Normalized throughput vs. MAC payload for deep ($K=0.1$) and low fading ($K=10$) and $M=25$ active nodes.

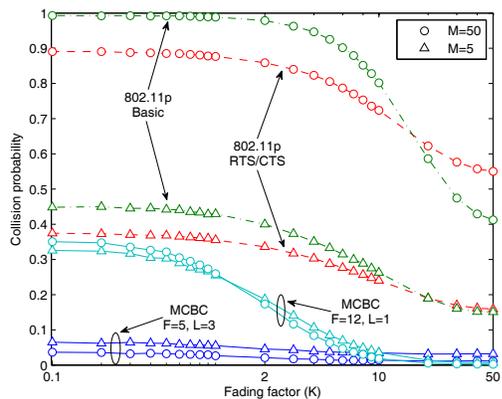


Fig. 22. Collision probability vs. fading severity, from deep (low K) to no fading (high K), for 5 and 50 active nodes.

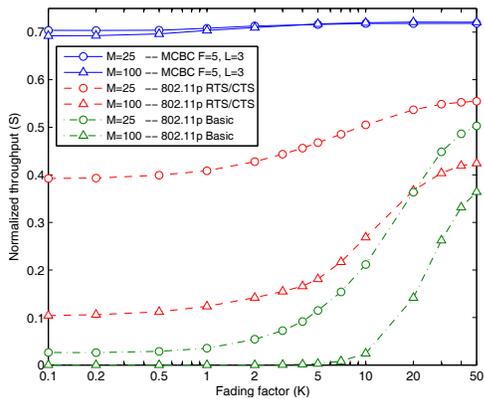


Fig. 23. Normalized throughput vs. fading severity, from deep (low K) to no fading (high K), for 25 and 100 active nodes.

overhead and negligibly more energy consumption, *MCBC* can use $L > 1$ at all times and in all types of scenarios and topologies. This important result highlights the scalability and robustness of the protocol, i.e., it provides a stable and high performance profile under variable or harsh conditions, without delay or energy consumption compromises.

Fig. 24 shows the actual throughput for all PHY bit-rates, in both high and low fading conditions. We notice that the throughput gain of *MCBC* over IEEE 802.11p increases with the bit-rate, thanks mainly to its low collision probability and

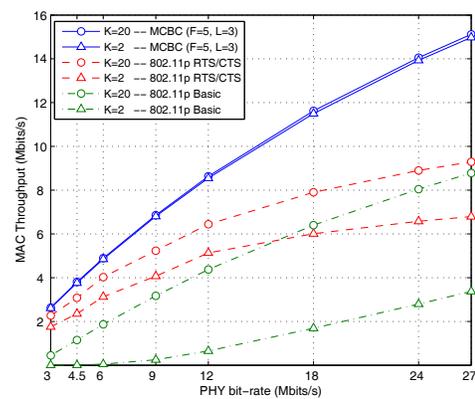


Fig. 24. Throughput for all PHY bit-rates in low ($K=20$) and moderate fading ($K=2$) channels for $M=25$ nodes.

low overhead – there are no random delays or distributed inter-frame spaces (DIFS) in *MCBC*. Overall, it offers above 57% more throughput on average- and above 81% more throughput at high bit-rates (> 18 Mbits/s) than the more robust IEEE 802.11p with RTS/CTS. In deep fading, depending on the bit-rate, *MCBC* can achieve more than 100% gain as the fading protection scheme adds no overhead.

VII. CONCLUSION AND FUTURE WORK

The *MCBC* cross-layer protocol was presented in depth with a keen eye towards realistic conditions. We discussed various real-world challenges pertaining to vehicular networks, e.g. high density of nodes, selective fading, time and frequency shift, synchronization and implementation issues, and presented a set of solutions as well as an analytical model to validate the performance. To further verify the theoretical and feasibility claims, we built a multi-core hardware testbed and ran simulations, at both PHY and MAC level. The simulation results in comparison with IEEE 802.11p show a considerable performance gain and markedly better resilience to channel fading and topology changes.

There are a number of key features of *MCBC* that are worth highlighting. The contention scheme spans both time and frequency domains and uses random individual subcarriers, reducing contention overhead and offering a very high success probability and hence improved reliability. The protocol maintains a high and very stable performance profile across a wide range of network loads and adverse channel conditions, thus offering high scalability and robustness – issues of particular relevance to vehicular networks. Considering also that contention parameters are fixed and contention sessions are very short ($< 70 \mu\text{s}$), it ensures resilience to variable topologies – another vehicular networks challenge – and adds to the protocol robustness. Finally, the fading protection mechanism adds no overhead and is insignificant in terms of energy cost so it can be enabled at all times, contrary to popular solutions like RTS/CTS.

Reliable QoS and prioritized access, which were design features of *MCBC*, were introduced in Section II-C and are the subject of future publications. We are currently in the process

of building fully featured *MCBC* prototypes and developing IEEE 802.11 compatibility mechanisms.

REFERENCES

- [1] 802.11p/D9.0, *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications - Wireless Access in Vehicular Environments (WAVE)*, IEEE Draft Std., Sep. 2009.
- [2] 802.11-2007, *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, IEEE Std., Jun. 2007.
- [3] H. Menouar, F. Filali, and M. Lenardi, "A Survey and Qualitative Analysis of MAC Protocols for Vehicular Ad Hoc Networks," *IEEE Wireless Commun. Mag.*, vol. 13, no. 5, pp. 30–35, Oct. 2006.
- [4] A. Pal, A. Dogan, F. Özgüner, and Ü. Özgüner, "A MAC Layer Protocol for Real-time Inter-vehicle Communication," in *IEEE ITSC*, 2002.
- [5] R. M. Yadumurthy, A. Chimalakonda, M. Sandashivaiah, and R. Makanaboyina, "Reliable MAC Broadcast in Directional and Omni-Directional Transmissions for Vehicular Ad Hoc Networks," in *2nd ACM Int. Workshop on Vehicular Ad Hoc Networks*, 2005.
- [6] S. Yang, H. H. Refai, and X. Ma, "CSMA based inter-vehicle communication using distributed and polling coordination," in *IEEE ITSC*, 2005.
- [7] B. Roman, I. Chatzigeorgiou, I. Wassell, and F. Stajano, "Performance Evaluation of Multi-Carrier Burst Contention and IEEE 802.11 with Fading During Channel Sensing," in *IEEE PIMRC*, 2009.
- [8] S. Bana and P. Varaiya, "Space Division Multiple Access (SDMA) for Robust Ad Hoc Vehicle Communication Networks," in *IEEE ITSC*, 2001.
- [9] J. J. Blum and A. Eskandarian, "A Reliable Link-Layer Protocol for Robust and Scalable Intervehicle Communications," *IEEE Trans. Intell. Transp. Syst.*, vol. 8, no. 1, pp. 4–13, Mar. 2007.
- [10] M. Lott, R. Halfmann, E. Schultz, and M. Radimirsch, "Medium Access and Radio Resource Management for Ad Hoc Networks Based on UTRA TDD," in *ACM MobiHoc*, 2001.
- [11] F. Borgonovo, L. Campelli, M. Cesana, and L. Coletti, "MAC for Ad-Hoc Inter-Vehicle Network: Services and Performance," in *IEEE VTC Fall*, 2003.
- [12] G. Lu, B. Krishnamachari, and C. S. Raghavendra, "An Adaptive Energy Efficient and Low-latency MAC for Data Gathering in Wireless Sensor Networks," in *IEEE IPDPS*, 2004.
- [13] F. Yu and S. Biswas, "Self-Configuring TDMA Protocols for Enhancing Vehicle Safety with DSRC Based Vehicle-to-Vehicle Communications," *IEEE J. Sel. Areas Commun.*, vol. 25, no. 8, pp. 1526–1537, Oct. 2007.
- [14] J. Zhang, Q. Zhang, and W. Jia, "VC-MAC: A Cooperative MAC Protocol in Vehicular Networks," *IEEE Trans. Veh. Technol.*, vol. 58, no. 3, pp. 1561–1571, Mar. 2009.
- [15] N. Choi *et al.*, "A Solicitation Based IEEE 802.11p MAC Protocol for Roadside to Vehicular Networks," in *IEEE INFOCOM*, 2008.
- [16] C. Suthaputthakun and A. Ganz, "Priority Based Inter-Vehicle Communication in Vehicular Ad-Hoc Networks Using IEEE 802.11e," in *IEEE VTC Spring*, 2007.
- [17] S. Shankar and A. Yedla, "MAC Layer Extensions for Improved QoS in 802.11 Based Vehicular Ad Hoc Networks," in *IEEE ICVES*, 2007.
- [18] B. Roman, F. Stajano, I. Wassell, and D. Cottingham, "Multi-Carrier Burst Contention (MCBC): Scalable Medium Access Control for Wireless Networks," in *IEEE WCNC*, 2008.
- [19] F. 03-024, *FCC Report and Order*, Federal Communications Commission Std., Feb. 2004.
- [20] S. W. Bergen and A. Antoniou, "Design of Ultraspherical Window Functions with Prescribed Spectral Characteristics," *EURASIP Journal on Applied Signal Processing*, vol. 13, pp. 2053–2065, Jan. 2004.
- [21] R. Lyons, *Understanding Digital Signal Processing*. Upper Saddle River, NJ: Prentice Hall PTR, 2001.
- [22] J. Proakis, *Digital Communications*, 4th ed. McGraw-Hill, 2000.
- [23] G. Karagiannidis, N. Sagias, and P. Mathiopoulos, "N* Nakagami: A Novel Stochastic Model for Cascaded Fading Channels," *IEEE Trans. Commun.*, vol. 55, no. 8, pp. 1453–1458, Aug. 2007.
- [24] M. Abramowitz and I. Stegun, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, 9th ed. Dover, 1972.
- [25] T. E. Tkacik, "A Hardware Random Number Generator," *CHES 2002, Springer LNCS*, vol. 2523, pp. 450–453, 2003.
- [26] Trimble Navigation Ltd., "Resolution T," 2004. [Online]. Available: <http://www.trimble.com/resolutiont.shtml>
- [27] NavSync Ltd., "CW25-TIM," 2005. [Online]. Available: http://www.navsync.com/docs/CW25-TIM_DS.pdf
- [28] —, "CW25-TIM Holdover Test," 2005. [Online]. Available: http://www.navsync.com/docs/AN03_GPS_Timing.pdf
- [29] A. Dutta, D. Saha, D. Grunwald, and D. Sicker, "SMACK: a Smart ACKnowledgment scheme for broadcast messages in wireless networks," *SIGCOMM Comput. Commun. Rev.*, vol. 39, no. 4, pp. 15–26, 2009.
- [30] T. Weiss and F. Jondral, "Spectrum Pooling: An Innovative Strategy for the Enhancement of Spectrum Efficiency," *IEEE Commun. Mag.*, vol. 42, pp. 8–14, Mar. 2004.
- [31] M. Evans, N. Hastings, and B. Peacock, *Statistical Distributions*. New York: Wiley, 2000.



Bogdan Roman received the BEng. degree in Electrical and Communications Engineering from the Politehnica University of Bucharest, Romania, in 2002 and the MSc. degree in Mobile Communications from University of Versailles, France, in 2002. Until 2006 he held positions in the fields of mobile networks and supercomputing at Vodafone and Esther Graphics, France. He is now working towards his PhD degree at the Computer Laboratory, University of Cambridge, UK. His research is mainly on cross-layer and high efficiency channel access schemes. In

2010 his project was awarded a Microsoft Research Software Radio Academic Program Award to further research into Multi-Carrier Bursting techniques. Bogdan's other interests include the design, modelling and implementation of randomized and adaptive techniques.



Ian Wassell is a Senior Lecturer at the University of Cambridge Computer Laboratory. He received the PhD degree from the University of Southampton in 1990 and the BSc., BEng. (Honours) Degrees (First Class) from the University of Loughborough in 1983. He has in excess of 15 years experience in the simulation and design of radio communication systems gained via a number of positions in industry and higher education. He has published in excess of 170 papers concerning wireless communication systems since joining the University of Cambridge in May 1990, and is also a fellow of Churchill College. His research interests include broadband fixed wireless networks, wireless sensor networks, radio propagation and modelling, coding and communication signal processing. He is a member of the Institution of Engineering and Technology.



Ioannis Chatzigeorgiou received the Dipl.-Ing. degree in Electrical Engineering from Democritus University of Thrace, Greece, in 1997 and the MSc. degree in Satellite Communication Engineering from the University of Surrey, UK, in 2000. From 2000 to 2002, he held positions at Marconi Communications and Inmarsat Ltd. In 2002 he joined the Cambridge University Engineering Department and completed his PhD studies in 2006. From 2007 to 2009 he was a Research Associate at the Computer Laboratory, University of Cambridge, UK. In 2010 he received

an Alain Bensoussan fellowship by the European Research Consortium for Informatics and Mathematics (ERCIM) to pursue research in wireless communications at the Norwegian University of Science and Technology (NTNU). Ioannis' research interests are in coding, modulation and detection techniques, multiple antenna systems and cooperative networks.