

Evaluation of Multi-Carrier Burst Contention and IEEE 802.11 with Fading During Channel Sensing

Bogdan A. Roman[†], Ioannis Chatzigeorgiou, Ian J. Wassell, Frank Stajano
Digital Technology Group, Computer Laboratory, University of Cambridge, United Kingdom
Email: {abr28, ic231, ijw24}@cam.ac.uk

Abstract—With the prevalence of wireless capability in mobile devices and the increasing number of wireless network deployments where protocols using carrier sensing (such as IEEE 802.11) are employed, it is important to understand the impact of fading during the channel sensing period and its direct effects on performance in realistic environments. In this paper we evaluate the performance of IEEE 802.11 and that of our proposed cross-layer Multi-Carrier Burst Contention (*MCBC*) protocol, which we have implemented in hardware, in realistic indoor fading environments with no direct line of sight. We present thorough simulation results, backed up by hardware and real-world measurements, performance issues for each protocol and describe methods for increasing resilience to fading during channel sensing in order to improve performance.

I. INTRODUCTION

The IEEE 802.11 protocols family is now the dominant standard in Wireless Local Area Networks. The IEEE 802.11 Medium Access Control (MAC) has received much attention and a wealth of papers have proposed analytical models or presented simulation results concerning its performance. However, in much of this work, ideal channel conditions are assumed [1], [2], whilst papers that consider more realistic scenarios such as fading [3]–[5], do not always take into account channel fading during carrier sensing and the subsequent effects on protocol performance.

Channel or carrier sensing in IEEE 802.11 is achieved using one of the available Clear Channel Assessment (CCA) techniques, the most common of which are *energy detection* (ED) [6]–[8] and *preamble detection* (PD) [6], [7]. In ED the instantaneous power of the received signal is measured over a sensing time and compared to a threshold that is usually predefined. The measured energy level sets the channel status flag to either *busy* (above threshold) or *idle* (below threshold). On the other hand, PD uses the additional physical (PHY) layer resources to constantly process the receiver input and declares the channel busy when features of the protocol – most commonly the preamble – are detected in the data stream. Compared to ED, PD can detect signals of lower power rendering it a more reliable CCA technique. However, ED is of much lower complexity and is considerably less power-hungry than PD which made it an attractive choice for most manufacturers of mobile devices where extending battery life is of prime concern. In general, during the scanning period, when nodes are searching for Access Points (APs), PD is used

in order to accurately discover all reachable APs. However, datasheets for several IEEE 802.11a/g OFDM Radio-on-Chip devices have confirmed that at least during the backoff period, ED is employed.

In this paper we concentrate on wireless networks of power constrained terminals, hence we assume carrier sensing for IEEE 802.11 during backoff is performed using ED. We consider the OFDM-based IEEE 802.11 protocols (e.g. a/g) and our recently proposed Multi-Carrier Burst Contention (*MCBC*) [9] protocol that we have implemented in hardware. *MCBC*, which makes use of a novel node elimination algorithm and exploits both time and frequency domains of the underlying OFDM PHY layer, imposes some synchronization requirements but yields very good performance in ideal channel conditions. The objective of this paper is to evaluate the performance of both protocols in practical scenarios (i.e., with fading), where sensing is performed using ED during contention. In particular, we consider Rayleigh fading which impairs the link quality, introducing hidden nodes, and can thus markedly degrade the protocols' performance.

II. CHANNEL MODEL

We consider a wireless network of M nodes plus an AP and that the links between nodes are subject to Rayleigh fading, modeling indoors environments with no direct line of sight.

Let us consider that node i transmits to the destination and node j senses the channel, with $i, j = 1, \dots, N$ and $i < j$. We use $\gamma_{i,j}$ to denote a random variable that represents the instantaneous signal-to-noise ratio (SNR) at node j , when node i is transmitting. For simplicity, we assume that the links between nodes are statistically similar, therefore they can be described by an average SNR, $\bar{\gamma}$, given by $\bar{\gamma} = \mathbb{E}[\gamma_{i,j}]$ for all values of i, j , where $\mathbb{E}[\cdot]$ denotes the expectation operation.

In our fading model, the amplitude of the received frame, represented by $\alpha_{i,j}$, is Rayleigh distributed. Its probability density function (PDF) is given by:

$$f(\alpha_{i,j}) = \frac{\alpha_{i,j}}{\sigma^2} e^{-\alpha_{i,j}^2/2\sigma^2},$$

where $2\sigma^2$ is the mean power of the received frame and can be obtained [10] from the average SNR:

$$\bar{\gamma} = \frac{2\sigma^2}{N_0},$$

[†]The first author is supported by Girton College at University of Cambridge and Cambridge European Trust.

where $N_0/2$ is the noise power spectral density per complex dimension. The instantaneous SNR, $\gamma_{i,j}$, which is exponentially distributed with mean $\bar{\gamma}$, and its PDF are given by:

$$\gamma_{i,j} = \frac{\alpha_{i,j}^2}{N_0} \quad \text{and} \quad f(\gamma_{i,j}) = \frac{1}{\bar{\gamma}} e^{-\gamma_{i,j}/\bar{\gamma}}.$$

Node j can correctly identify a busy channel if $\gamma_{i,j}$ is at least equal to an SNR threshold γ_0 , usually predefined in the standard. However, in the event that $\gamma_{i,j} < \gamma_0$, node j erroneously detects the channel as idle and yields a false negative. In that case, node j is *hidden* from the transmitting node i and could thus cause a collision. The probability that node i will not be overheard by node j is also known as *outage probability* and can be calculated as follows [10]:

$$P_{\text{out}} = \text{Prob}\{\gamma_{i,j} < \gamma_0\} = \int_0^{\gamma_0} f(\gamma_{i,j}) d\gamma_{i,j} = 1 - e^{-\gamma_0/\bar{\gamma}}.$$

When $n > 1$ nodes are transmitting simultaneously then the resulting received power at a receiver follows a chi-squared distribution with $2n$ degrees of freedom [10]. In this case, the outage probability becomes:

$$\xi(n) = 1 - e^{-\gamma_0/\bar{\gamma}} \sum_{k=0}^{n-1} \frac{1}{k!} \left(\frac{\gamma_0}{\bar{\gamma}}\right)^k.$$

As intuition suggests, poor link conditions markedly increase the outage probability, leading to a potentially high number of hidden nodes which in turn, reduce the probability of a node successfully completing its transmission.

III. SYSTEM MODEL

A. IEEE 802.11

The core functionality of the IEEE 802.11 protocol has become well known and the usage of the Distributed Coordination Function (DCF) is predominant among IEEE 802.11 network deployments. A *binary exponential backoff* procedure is the collision resolution employed by DCF (Fig. 1).




Fig. 1. Collision avoidance in IEEE 802.11 using binary exponential backoff.

Nodes with data to send first sense the channel for a duration of a Distributed Interframe Space ($34 \mu\text{s}$ for the OFDM PHY) and if it is idle, they initiate a counter to a random value between 0 and a contention window (CW) size less one. The counter is decremented every backoff slot ($9 \mu\text{s}$ for the OFDM PHY). While decrementing, nodes continue to sense the channel and will freeze their counter in case they detect the channel as *busy*. The procedure is resumed in the same manner after the channel becomes idle again and the counter is decremented until it reaches 0, at which point the node transmits its frame. In case the frame collides then the CW is doubled and the backoff procedure is reinitiated. A collision is determined if an acknowledgement (ACK) frame is not

received when using the basic access mode or the CTS-to-self (Clear-To-Send) mode, or if a CTS frame is not received when using the RTS/CTS (Request-To-Send/Clear-To-Send) mode. CTS-to-self is only employed when there are mixed nodes (e.g. 802.11b/g) and is otherwise equivalent to the basic access mode (i.e., the data frame is transmitted directly after the backoff period) although it adds some overhead.

Fading during the backoff procedure can affect the node's ability to correctly sense the channel state and hence can increase the probability of collision.

We collected wireless data in the William Gates Building of the Computer Laboratory where there are over 25 access points managed by the department and various other wireless networks. We have used a multitude of PDA and laptop devices with wireless capability from different manufacturers to observe the sensing performance during scanning, when the more effective and power hungry PD technique is used to detect APs. Many devices were able to decode AP beacons at power levels as low as -95 dBm . We also observed that, with one exception, all devices we tested had RTS/CTS disabled, which is also the default setting. It is thus of high importance to study the performance of the basic access mode which is used by default by almost all devices.

Taking the above into account, we believe that, *as far as channel sensing is concerned*, ideal channel conditions between the AP and the nodes is generally not a non-realistic assumption for simulations because a) the AP is plugged into the mains and can afford to transmit signals at much higher energy levels than mobile nodes and b) the AP can afford to use PD as opposed to ED hence the additional processing gain allows it to decode very weak signals – provided they are from a unique transmitter – even at received powers of -97 dBm . Thus, collisions occur mainly between nodes. Collisions between nodes and the AP can only occur at the ACK (CTS) frames, in the case where the colliding node(s) did not hear the previous data (RTS) frame when using the basic (RTS/CTS) access mode and their transmission starts at the same time as the AP's ACK (CTS). This however accounts for less than 1% of the collisions and it can be ignored.

B. Multi-Carrier Burst Contention (MCBC)

MCBC is a cross-layer protocol which uses a contention mechanism that employs a rapidly converging node elimination algorithm based on synchronized rounds. In this section we briefly summarize the main features of *MCBC*. A more detailed description can be found in [9].




Fig. 2. Contention session example; c , r and n are the number of contenders, referees and nominees respectively. The shaded cells denote the contention subcarrier activated by referees in the feedback slots.

The contention algorithm will be explained using an example. Fig. 2 depicts a contention session for a network of $M = 250$ competing nodes when $N = 3$ contention rounds are considered and $F = 4$ contention subcarriers¹ are used. All 250 nodes are initially *contenders* ($c = 250$). During the contention slot of the first round, denoted as c_1 in Fig. 2, contenders become *nominees* with a predefined probability p . Let's assume that $n = 132$ nodes are elevated to nominee status; subsequently, each one of them randomly activates a single contention subcarrier² f_i , where $i = 1, \dots, F$. The remaining $r = 118$ contenders sense the channel in order to identify the active subcarrier with the highest index, in this case f_4 , and become *referees*. During the feedback slot of the first round, namely b_1 , the referees activate subcarrier f_4 and the 34 nominees that had used it in c_1 are selected as round winners. Winners are promoted to contenders ($c = 34$) for the next round and the contention continues similarly until the end of the third round. In b_3 , referees select the active subcarrier with the highest index (in this case f_3), and the final winner initiates its data transmission at the end of the session. If more than one winner survive contention, a collision will occur.

The algorithm maximizes the probability that there is a single winner in the last round. On average, roughly $\lceil M (p/F)^N \rceil$ contenders are expected to reach the end of the contention session. The protocol also guarantees that there will be at least one winner at the end of the contention while fairness is statistically ensured since each node uses the same parameters for every contention session.

MCBC uses a slightly modified version of the OFDM PHY that IEEE 802.11a/g is based on. It relies on sensing the energy level of individual contention subcarriers at the output of the Fast Fourier Transform (FFT) block located at the receiver. The effect of frequency drift is alleviated using windowing techniques and proper spacing of the contention subcarriers [9]. Accurate slot synchronization is needed to maintain a low overhead. We have designed a sync mechanism that maintains a sync accuracy of less than $1 \mu s$ and the protocol was implemented in hardware (see Section IV).

The collision probability of the *MCBC* algorithm can be affected by fading, which causes each referee to detect different energy levels at each contention frequency. However, if an AP exists, it can be used as referee to decrease collision probability. It is important to notice that since *MCBC* nodes are synchronized, a collision can occur only if two or more nodes begin transmission at the same time, irrespective of the fading conditions or data frame size.

IV. RESULTS AND DISCUSSION

To evaluate the fading effects during channel sensing for IEEE 802.11 and *MCBC*, we used simulations which accurately implement the protocols. We used our simulator³ that we validated against OPNET, results from literature [2], [3] and real-world measurements. For IEEE 802.11, we built 2 to

10 node networks using Atheros AR5002X IEEE 802.11a/b/g cards and the collision probability and throughput results matched the simulations with $\pm 0.2\%$ error. We implemented the *MCBC* protocol in hardware on an Altera EP2C35F FPGA platform using mainly Verilog for the FPGA part. Multiple *MCBC* cores (up to 15, the 16 available clock buffers being the limiting factor) were implemented on the same chip. External, independent quartz oscillators from different manufacturers were used to clock each core to validate the *MCBC* synchronization mechanism. We ran proof of concept contention sessions with tight timings [9] and the collision probability was within 0.5% error compared to the simulations and theoretic model. The *MCBC* hardware implementation will be presented in detail in a subsequent paper.




Fig. 3. *MCBC* hardware implementation.

In simulations, the PHY used for both protocols was OFDM with a bit-rate of 6 Mbps and a frame size of 1024 bytes unless otherwise stated. We assumed saturation conditions, i.e., all nodes always have data to send to the AP. We used the standard OFDM parameters for IEEE 802.11, $W = 16, m = 6$ for the minimum contention window size and maximum backoff stage respectively. Unless otherwise stated, *MCBC* uses the basic parameter set $N = 3, F = 8, p = [0.35, 0.8, 0.85]$ for the number of contention rounds, number of contention subcarriers and the per-round probability of transmission respectively as seen in Section III-B.

For fading, the received signal amplitude for each link was independently subjected to Rayleigh fading and we considered two scenarios: *quasi-static fading* where the fading statistics do not change during the transmission of a frame and *fast fading* where the fading statistics change at random time intervals of duration $9 \mu s < t < 100 \mu s$. The results of both scenarios were on average the same. This is because in IEEE 802.11, the backoff algorithm makes nodes attempt transmissions at random times while in *MCBC* the collision probability does not depend on the data frame size.

The average received power was varied to yield a link outage probability range of $0 \leq P_{out} \leq 0.92$. We chose a detection threshold of -82 dBm [11] and varied the average received power between -86 dBm and 0 dBm. However, using different values do not affect the results as long as the above P_{out} range is satisfied. This allowed us to observe system behavior ranging from ideal conditions, when all nodes can hear each other, to scenarios where almost all nodes

¹A small subset of the 52 subcarriers [11], part of the same OFDM channel.

²No data is being sent or modulated on contention subcarriers [9].

³Available online at <http://www.cl.cam.ac.uk/research/dtg/wisim>

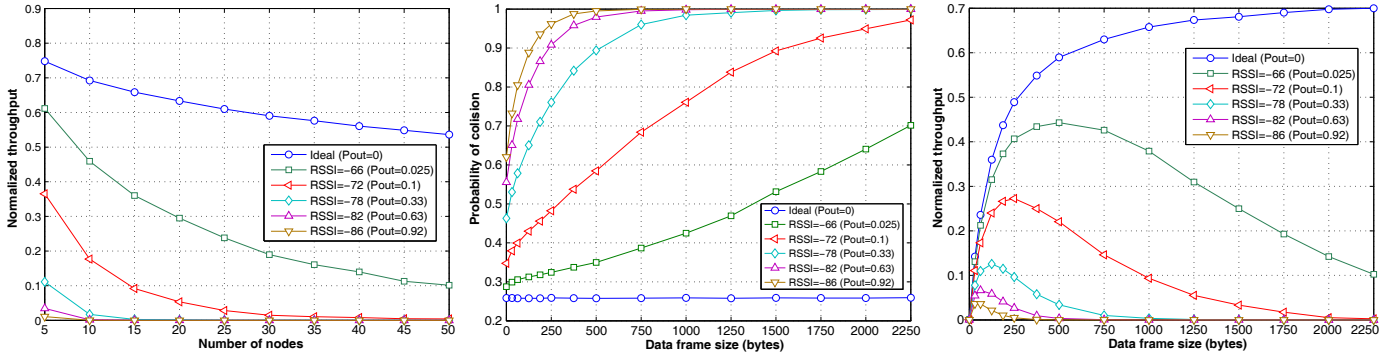


Fig. 4. IEEE 802.11 basic access throughput (left), collision probability for $M=15$ nodes (center), variable data frame (right). RSSI values are in dBm.

are hidden from each other. The minimum received power considered was -99 dBm (background noise).

The *MCBC* nodes did not use the AP as referee during contention so they ran the decentralized algorithm, i.e., nodes become referees as they lose in a round. We considered *a)* each subcarrier from each link independently subjected to Rayleigh fading (i.e., *frequency selective fading*) but also *b)* grouped subcarriers subjected to the same level of fading (i.e., *semi-flat fading*). The results for *a)* and *b)* were very close owing to the spacing of the F contention subcarriers.

For IEEE 802.11 basic access mode, the number and duration of collisions are greatly increased by fading and throughput decreases considerably as shown in Fig. 4 (left). With fading, a competing node may not detect a transmission during backoff and will continue to decrement its counter rather than freezing it, eventually starting a colliding transmission and so increasing the average duration of collisions and decreasing throughput. As opposed to many models used in literature, collisions can occur at any point in the data frame when fading is considered. As a result, and when increasing the number of nodes, the performance penalty becomes important even at a very low link outage probability of $P_{out} = 0.025$.

This effect worsens as the data frame size increases as shown in Fig. 4 (center). For large data frames, collisions keep forming one after the other, merging into one long collision seen by the network so a successful frame rarely gets through. When short frames are used, however, even when fading is pronounced, the increasing size of the contention window of colliding nodes will eventually accommodate successful transmissions as the duration of a successful transmission decreases while colliding nodes are still backing off. The effect of this on throughput is represented in Fig. 4 (right) where we see that, in case of fading, there exists an optimum data frame size which maximizes the throughput. In general, in the presence of fading, large data frames sizes should be avoided.

Fig. 5 shows the throughput of *MCBC* with the basic parameter set. As expected, fading affects this mode of *MCBC* which otherwise is able to maintain a very low collision probability. However, the random choice between multiple contention subcarriers, which are spaced in frequency, helps combat fading overall. Since the protocol does not care who the contention winner is, as long as at least one subcarrier can be sensed during contention, the numbers of round winners is




Fig. 5. *MCBC* throughput with basic parameters (no fading resilience).

exponentially decreased and thus collisions are reduced.

A somewhat non-intuitive throughput increase is observed in Fig. 5. At low number of nodes and with fading considered, the few activated subcarriers are less likely to be detected by the listening nodes so there is a high chance that all nodes will select themselves as winners, thus causing a collision. A similar effect exists even without fading ($P_{out} = 0$) where the probability that either all or no nodes transmit in the contention slot is high. These effects are reduced when increasing the number of nodes. With more nodes, the received power at listening nodes increases and so does the probability that some of them turn into referees and reply in the feedback slot.

Both IEEE 802.11 and *MCBC* are able to reduce the collision probability in the presence of fading. IEEE 802.11 can use shorter data frames or activate RTS/CTS. *MCBC* can activate two or more subcarriers at the same time, which are spaced in frequency with F (i.e., activate subcarriers i and $F + i$ simultaneously so no overlapping occurs), as opposed to just one, and use scrambling (i.e., reordering) of the subcarrier indexes in the feedback slot. A node thus obtains a higher diversity hence a higher chance of more subcarriers being heard. The cost associated with enabling such features ultimately translates into using more power: sending and receiving of an RTS/CTS exchange in IEEE 802.11 and activating more than one subcarrier in *MCBC*. However, they are worth enabling in the presence of fading since the throughput gain usually makes up for the extra power, power that would otherwise be wasted resolving collisions.




Fig. 6. IEEE 802.11 RTS/CTS throughput. dBm values represent RSSI.

The short RTS frames considerably reduces IEEE 802.11's collision probability compared to the basic access mode. Assuming a CTS frame from the AP is usually heard by all nodes (updating their network allocation vector), a successful RTS frame means a successful data frame. Thus, longer data frames can be used to increase channel efficiency. This is reflected in Fig. 6 where a much reduced dependence of the throughput on the fading level is observed.

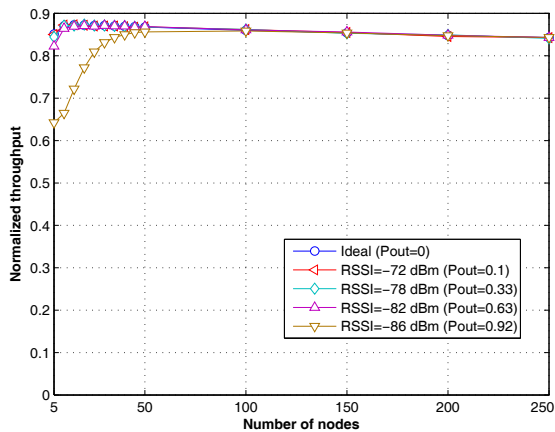


Fig. 7. *MCBC* throughput with $F = 5$ subcarriers repeated 3 times.

Allowing *MCBC* to activate more than one subcarrier simultaneously, visibly improves the overall throughput under fading as can be seen in Fig. 7. We observe the same increasing throughput effect as discussed above, although it is much improved owing to the diversity introduced by activating multiple simultaneous subcarriers.

Fig. 8 shows the throughput for all OFDM bit-rates [11] in moderate fading ($P_{out} = 0.5$). For IEEE 802.11 basic access, a high bit-rate results in shorter data frames, hence it has a similar beneficial effect we observed in Fig. 4. It is important to realize that a high bit-rate usually requires a higher received power and many devices do not usually attain 54 Mbps at -80.5 dBm received power. However, the values in Fig. 8 hold since the -80.5 dBm value was only chosen to give a link outage probability of $P_{out} = 0.5$ and, in our model, nodes transmit to the AP which many times has better sensitivity.

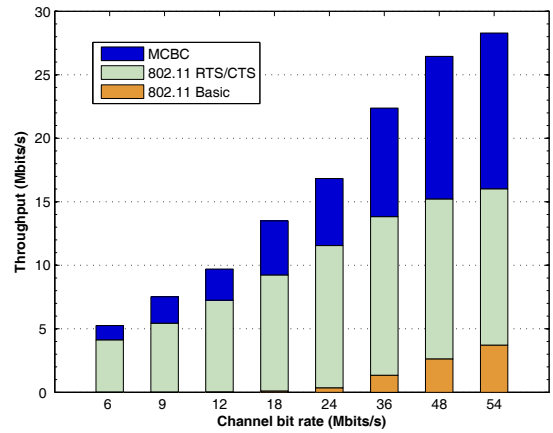


Fig. 8. Protocols throughput for all OFDM bit-rates in fading channels at -80.5 dBm average received power for $M = 25$ nodes.

In general, in fading conditions it is highly recommended to activate RTS/CTS for IEEE 802.11. Overall, *MCBC* offers 50% more throughput on average- and 75% more throughput at high bit-rates than IEEE 802.11 with RTS/CTS.

V. CONCLUSION AND FUTURE WORK

By incorporating Rayleigh fading statistics during channel sensing, we presented a performance evaluation of the standard IEEE 802.11 and our proposed *MCBC* protocol. We showed IEEE 802.11 is heavily affected by fading when using the basic access mode which is the mode used by default in more than 95% of devices. We outlined methods for increasing fading resilience for both protocols: shorter (optimal) frame sizes, higher bit-rates or RTS/CTS for IEEE 802.11 and subcarrier repeating and index scrambling for *MCBC*. We are currently developing Quality of Service extensions and a complete analytical model for *MCBC* to also account for fading.

REFERENCES

- [1] G. Bianchi, "Performance Analysis of the IEEE 802.11 Distributed Coordination Function," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 3, 2000.
- [2] C. Foh and J. Tantra, "Comments on IEEE 802.11 Saturation Throughput Analysis with Freezing of Backoff Counters," *IEEE Commun. Lett.*, vol. 9, 2005.
- [3] Z. Hadzi-Velkov and B. Spasenovski, "Saturation throughput-Delay analysis of IEEE 802.11 DCF in Fading Channel," in *IEEE ICC*, May 2003.
- [4] X. Zhou and J. C. Jr., "Cross-Layer Analysis of IEEE 802-11 DCF in Burst-Error Fading Channels with Diversity," in *IWCMC*, Jun. 2005.
- [5] X. Dong and P. Varaiya, "Saturation Throughput Analysis of IEEE 802.11 Wireless LANs for a Lossy Channel," *IEEE Commun. Lett.*, vol. 9, 2005.
- [6] S. Y. Shin, I. Ramachandran, S. Roy, and W. H. Kwon, "Cascaded Clear Channel Assessment: Enhanced Carrier Sensing for Cognitive Radios," in *IEEE ICC*, Jun. 2007.
- [7] I. Ramachandran and S. Roy, "On the Impact of Clear Channel Assessment on MAC Performance," in *IEEE GLOBECOM*, Dec. 2006.
- [8] K. Jamieson, B. Hull, A. Miu, and H. Balakrishnan, "Understanding the Read-World Performance of Carrier Sense," in *SIGCOMM*, Aug. 2005.
- [9] B. Roman, F. Stajano, I. Wassell, and D. Cottingham, "Multi-Carrier Burst Contention (MCBC): Scalable Medium Access Control for Wireless Networks," in *IEEE WCNC*, Mar. 2008.
- [10] A. Goldsmith, *Wireless Communications*. Cambridge Univ. Press, 2005.
- [11] 802.11a-1999 (R2003), *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications - High-speed Physical Layer in the 5 GHz Band*, IEEE Std., 1999, Reaffirmed June 2003.