An Overview of WLAN Performance, Some Important Case-Scenarios and their Associated Models

Federico Larroca and Fernanda Rodríguez

Abstract—The study of the 802.11 standard has been very intense for more than a decade now. Several works have striven at understanding its performance, even in the simplest topology of a wireless local area network (WLAN) with a single access point (AP). The present survey is an effort to classify and present the enormous literature on the subject into several important casescenarios, and summarizes the current understanding of WLAN performance. The resulting performance and associated models are discussed (and sometimes extended) and simulation results are used to illustrate them. We also highlight interesting open research problems that we believe the community should address.

I. INTRODUCTION

Wireless is probably today's most used access technology. Among the many possibilities, 802.11 [1] (commercially known as WiFi) has gained the majority of the market, mainly because of its low cost and reasonable performance. From refrigerators [2] to the more traditional laptops or tablets, virtually all modern electric devices are equipped with (at least) an 802.11 interface.

It should come at no surprise then that academia has invested an enormous effort on studying this kind of network. There exist studies based on simulations or experimentation; considering different flavors of 802.11 (e.g. a, b, g, e, n, ac); some propose improvements to the standard, while some propose new algorithms altogether.

At the most basic level, there exist two kinds of network configuration: WLAN (Wireless Local Area Network) and adhoc, both of them contemplated by the standard. In the former, all hosts send and receive data to and from a centralized node (known as the access point, AP). In the latter, no such node exists, and communication may be carried between all hosts.

The present survey is concentrated on the WLAN case, illustrated in Fig. 1. In particular, we present a summary of the current understanding of WLAN performance. To this end, we discuss what we consider the most important case-scenarios, and present the resulting 802.11 performance. In addition, we present models pertinent to each case-scenario, which will help us explain some of the observed phenomena. Moreover, results are verified (and sometimes extended) by means of simulation, although we refer to the literature for experimental campaigns when possible.

Given the sheer amount of papers on the subject, our intention is by no means to be exhaustive, but we will highlight only those papers we consider to be most representative of each scenario. Precisely, because of the enormous existing



Fig. 1: The network considered in the present article: all stations (in this case two laptops and two smartphones) communicate exclusively with the AP, which in turn may act as a gateway to the internet or any wired infrastructure.

literature, we believe the present paper to be a very important contribution and, surprisingly and to the best of our knowledge, the first overview of the subject.

The rest of the article is structured as follows. The next section gives a brief overview of the 802.11 standard. We then begin with the discussion of the considered case-scenarios. Section III presents one in which all stations can "hear" every other station, and they all have packets ready to be sent constantly. We consider unsaturated stations in Sec. IV, in particular when the stations are sending and receiving TCP traffic. The models and results we will present up to this last section are considering a simplistic receiver's PHY model, where if more than two transmissions share a (however small) portion of time at the receiver, neither one can be correctly received (i.e. no capture effect). We dedicate Sec. V to the capture phenomena and its effects over the performance of the WLAN. The infamous hidden terminal problem is discussed in Sec. VI. 802.11a/b/g are still by far the most widely used versions of the standard. As such, most of the discussion up to this last section focus on them. Relatively recent versions of the standard, such as 802.11e and 802.11n, are discussed in Secs. VIII and VII. In particular we will discuss how most of the analysis we will present for 802.11a/b/g is still valid for 802.11n. In Sec. X we concisely comment some important case-scenarios out of the scope of the present overview. We briefly discuss the very important and challenging problem

of metrology in an uncontrolled WLAN scenario in Sec. IX. Finally, we conclude the article in Sec. XI.

II. THE IEEE 802.11 STANDARD

We now briefly present the 802.11 standard, whose objective is to specify the physical (PHY) layer and the medium access control (MAC) mechanism for implementing a wireless local area network. The next two subsections are devoted to these two specifications. It should be noted however, that we will limit ourselves to those aspects that are most important from a global performance perspective. For a complete description, the interested reader should consult [1].

A. Physical Layer

Regarding the PHY layer, the standard notably specifies the working frequency bands and available modulation schemes (which in turn dictates the available data rates; e.g. DSSS with DQPSK encoding results in a 2Mbps rate). Another important aspect worth noting for our purposes is the presence of the Physical Layer Convergence Protocol (PLCP) sublayer, which is responsible of interfacing between the MAC layer and the "actual" PHY layer (or more correctly stated, the Physical Medium Dependent, PMD, sublayer). Its basic objective is to make the MAC layer as agnostic as possible to the different physical specifications we will discuss later. It should be noted that the PLCP sublayer appends a PHY-specific preamble and header to each packet, modulated at a specific rate, and that they must be taken into account in all computations and are not to be neglected.

These PLCP preamble and header are used by the receiver as follows. After detecting and measuring the frame's preamble power, the receiver adapts its gain to the measured power. It then uses a set of training symbols included in the preamble so as to synchronize with the transmitter. After successfully synchronizing, the receiver decodes the PLCP header that follows (which for instance includes information on the modulation used for the payload). If the header is correctly received (which is verified by a parity bit) the receiver goes to a receiving state and decodes the rest of the frame. To verify the correct reception of the complete frame, a CRC is appended at the end of it. Such detailed description of the receiving process will be helpful to understand the capture phenomena, discussed in Sec. V.

In any case, the most important differences between 802.11 flavors lies precisely in this layer. For instance, 802.11a works in the 5.8Ghz frequency range with OFDM (Orthogonal Frequency Division Multiplexing), and possible data rates of 6, 9, 12, 18, 24, 36, 48 and 54 Mbps. On the other hand, 802.11b operates at the 2.4 Ghz band and has three possibles physical layers: IR (Infrared), FHSS (Frequency Hopping Spread Spectrum) and DSSS (Direct Sequence Spread Spectrum), although only the latter became widely implemented. The raw data transfer rates supported by this PHY are 1, 2, 5.5 and 11 Mbps. Naturally, as the rate increases, so does the Signal to Interference and Noise Ratio (SINR) required for a correct

reception (although recent studies show that this is not entirely true [3]).

Although the 802.11a flavor offers bigger data rates than the 802.11b counterpart, the latter is still more widely used mostly due to the fact that at the moment of their release, only the 2.4Ghz band was unlicensed worldwide (notably in Europe). 802.11g was proposed so as to obtain the benefits of both 802.11 a and b. It works using OFDM in the 2.4Ghz band and is backwards-compatible with 802.11b to facilitate its market incorporation.

However, the ever increasing traffic demands called for even higher data rates. The answer was the 802.11n standard [4], whose objective was to obtain at least a 100Mbps throughput measured at the MAC level. Multiple-Input Multiple-Output (MIMO), and the option of operating using a channel with twice as much bandwidth as the legacy 802.11 (resulting in either 20MHz or 40MHz channels) are the key to the greater physical data rates in 802.11n. Further optional features are considered in the standard with the objective of improving robustness and thus range. For instance, two variants of transmit beamforming are included, as well as Space Time Block Coding (STBC). In a nutshell, while the former adapts the spatial streams to the current MIMO channel (by means of the so-called steering matrix), the latter sends redundant information over the available streams. When all possibilities are considered (spatial streams, modulation method, coding rate, channel bandwidth, and guard interval), the number of possible raw data rates has enormously increased. While in 802.11a we had 8 different rates, the new amendment offers 128, ranging from 6.5Mbps to 600Mbps.

B. Medium Access Control Mechanism

This subsection describes the main MAC mechanism defined by 802.11a/b/g, the so-called Distributed Coordination Function (DCF). It then discusses the modifications implemented in 802.11e and 802.11n.

DCF is a CSMA/CA (Carrier Sense Multiple Access with Collision Avoidance) mechanism, that is executed by each node in a distributed way. More precisely, it is a variation of the so-called *p-persistent* CSMA proposed by Kleinrock and Tobagi in [5]. The standard also specifies a centralized MAC mechanism (the Point Coordination Function, PCF). Cards implementing it are almost nonexistent, and as such, we shall not discuss PCF any further. In any case, the foremost objective for a MAC mechanism is to avoid potentially interfering concurrent transmissions to occur. In DCF this objective is achieved roughly in two steps.

As in every CSMA protocol, in DCF the first step that a station that wants to access the medium must take is to decide whether the channel is idle or not (i.e. perform the carrier sense). This task is carried out in two ways: physically and virtually. The former is a service provided by the PHY layer, and is called Clear Channel Assessment (CCA). According to the standard, the CCA implementation should report a busy medium at least for one of the following conditions. Either when the detected energy level is above a certain threshold,

when a signal with the same PHY characteristics is detected, or a combination of both. The virtual carrier sensing is performed by the MAC layer, and is based on its channel reservation capabilities. More of this will be discussed when explaining the Request To Send/Clear To Send procedure.

The second step is the collision avoidance (CA) one. This is implemented by the backoff mechanism, which randomly delays the transmission of frames. This mechanism helps in preventing the synchronization of transmissions, since all nodes that are waiting to transmit observe the channel becoming idle at the same time, and would collide if the CA would not be implemented.

More in detail, the complete MAC mechanism is as follows. Assume the station has just successfully sent a packet and received its corresponding acknowledgement. The station then senses the medium and waits until it is detected idle for a period longer than the so-called DCF Interframe Space (DIFS). Then a random number with uniform distribution between (0, W - 1) is chosen. For this first transmission attempt, W takes the value of W_{\min} and it is multiplied by 2 for each retransmission (i.e. not receiving the corresponding acknowledgement) up to a maximum value of $W_{\max} = 2^m W_{\min}$. The packet will be dropped when a certain number of retransmissions is reached. On the other hand, with each successful transmission, W returns to its minimum value.

Once the backoff counter is selected, it is decremented by one each time the medium is detected as idle for a period of aTimeSlot. If the backoff counter has not reached zero and the medium becomes busy, the station freezes its counter. That is to say, the value of the backoff counter is saved and is decremented again when the channel is idle for a period longer than the so-called Extended Interframe Space (EIFS) if the station detected and synchronized to the transmission using the channel, but either the PLCP header or the CRC check failed. In any other case, it waits for a DIFS period. It is important to highlight however, that most simulators (in particular ns-3 [6]) always wait a EIFS period after a failed reception, even if it did not synchronize with the PLCP preamble.

When the backoff reaches zero, the station will check whether it has a packet waiting to be sent in its queue. In such case, it will immediately send it. Else, it will wait for a packet to arrive to the queue. When a packet arrives, the station checks whether the channel has been idle for more than DIFS or EIFS. If so, it will send the packet right away. In any other case, it will wait until the channel is idle for more than DIFS or EIFS, and performs again the backoff procedure (i.e. draws a new backoff counter and counts down to zero).

After the destination node receives the data frame correctly, it must send an acknowledgement (ACK) frame to inform it. This is necessary since the transmitter cannot tell if its frame has been correctly received by listening to the medium (as in wired mediums, such as Ethernet). This control frame is sent within a time called Short Interframe Space (SIFS), shorter than both DIFS and EIFS so as to provide a higher priority to the ACK frame. If the sender does not receive the ACK frame within a time called ACK Timeout after the original transmission, the frame is scheduled for retransmission.

What we just described is called Basic Access. Figure 2 illustrates the frames and times involved after two frames collide. In addition to the basic access, there is a four-way handshaking access method called Request to Send/Clear to Send (RTS/CTS). In this mode, when a station is granted access to the channel, instead of immediately sending the data frame, it first sends a control frame called RTS. Once the destination node correctly receives the RTS frame, it waits for a SIFS period and responds with another control frame, this one called CTS. The transmitter station then sends the data frame (after waiting for another SIFS). Both control frames carry the information of the time remaining to complete the data exchange, including the ACK frame. The stations that detect either the RTS or CTS frames read this information, and update their so-called Network Allocator Vector (NAV) accordingly. The NAV indicates the time periods when the station should not transmit, even if its CCA senses the channel as idle. During the periods the channel is reserved, and if the station was executing the backoff procedure, the backoff counter is frozen. The objective of this mechanism is to solve the hidden station problem, a problem we will discuss later in the article.

To attain the 100Mbps throughput objective in 802.11n, and for reasons that will become clear later in the article, not only modifications in the physical layer were necessary, but the MAC layer efficiency had to be improved. This was achieved basically by aggregating frames. That is to say, several higher-layer packets are sent together in a relatively big frame. This way, new and legacy stations can co-exist, since DCF's basic operation is still the same. The amendment contemplates two different aggregation schemes, which may be combined: aggregate MAC service data unit (A-MSDU) and aggregate MAC protocol data unit (A-MPDU), the last of which includes selective acknowledgements by means of the so-called BlockAck frame. A deeper discussion on the specific characteristics of 802.11n and their impact on performance is included in Sec. VII.

In order to provide the MAC layer with Quality of Service (QoS), 802.11e was developed [7]. Its basic idea is to differentiate between types of traffic at the MAC level, and to provide priorities between them. This objective is performed by the Enhanced Distributed Channel Access (EDCA) mechanism, a backwards-compatible variant of DCF, of which we now briefly present its operation. It is important to note however that usage of EDCA has been marginal, therefore its discussion will be limited to Sec. VIII.

The standard considers four traffic classes (or access classes, AC): voice, video, best effort and background. Each station now includes a queue for each AC, and the backoff procedure is executed in each queue independently. Collisions between packets of the same station (or virtual collisions) are resolved in favor of the AC with the biggest priority. Moreover, three "knobs" are used to differentiate between packets, as each AC now presents its own value of: (i) the amount of time it has to wait before resuming the backoff procedure after



Fig. 2: Example of Basic access mechanism after two frames have collided. All stations have packets in their queue at all times.

| MAC Header size | 272 bits |
|-------------------------------|-------------|
| PLCP Header size | 48 bits |
| PLCP Preamble size | 144 bits |
| ACK size | 112 bits |
| CTS size | 112 bits |
| RTS size | 160 bits |
| aSlotTime | $20\mu s$ |
| SIFS | $10 \mu s$ |
| DIFS | $50 \mu s$ |
| EIFS | $364 \mu s$ |
| PLCP header and preamble rate | 1 Mbps |
| W _{min} | 31 |
| Wmax | 1023 |

TABLE I: Parameters of 802.11b Standard

the medium is sensed as idle (basically a per-AC value of EIFS and DIFS)¹; (ii) the maximum and minimum contention window; (iii) the maximum amount of time each AC may use the channel during each access opportunity (the so-called Transmission Opportunity, TXOP).

The simulation results we will present were obtained by means of the ns-3 simulator version 3.14, using its default parameters except when specifically stated. In particular, we have used the 802.11b standard, the parameters' value of which are shown in Table I. Moreover, in Table II we show other parameters we used in our simulations that are important in the computations. Of these, some were our choice (like the control rate) and other were automatically set by the simulator (like the rate at which the ACK frames are sent).

III. FULL CONNECTIVITY CASE-SCENARIO

In this subsection we discuss the simplest case-scenario. A fixed number n of stations are disposed around the AP in such

¹More precisely, for all practical purposes DIFS is replaced by the Arbitration Interframe Space (AIFS), the value of which depends on the AC as follows: $AIFS[AC] = AIFSN[AC] \times aSlotTime + aSIFSTime, where AIFSN[AC]$ is an integer. Similarly, EIFS is now "replaced" by EIFS – DIFS + AIFS[AC].

| Propagation Delay | 7 ns |
|-------------------|-----------|
| RTS rate | 2Mbps |
| ACK rate | 1 Mbps |
| CTS rate | 1 Mbps |
| Data rate | 11Mbps |
| Data payload | 988 Bytes |

TABLE II: Other parameters used

way that the CCA of any node always returns busy if any other node is transmitting. In the literature this is typically expressed as "every station can hear every other station in the network", or simply that there are no hidden stations. Moreover, this case-scenario further assumes that there does not exist the socalled channel capture effect. That is to say, no receiver is capable of decoding a packet if there is any other concurrent transmission, even if the former and the latter share a very small portion of air-time. Such event is called *collision*.

A. Bianchi's Model

The first articles that analyzed this case scenario focused primarily on estimating the *saturation throughput*. That is to say, calculating the per-station or the system's total throughput when all stations always have a packet ready to be sent (which may be the case if all stations send UDP flows with a high enough rate). This is a very important performance indicator, since it may be proved that for certain systems, stability of the queues is guaranteed as long as new packets are generated at a rate that is less than the saturation throughput (see for instance [8], or [9] for a result pertaining wireless networks specifically).

Among these early works, we may cite the simulation campaign carried out in [10] or the analytical studies of [11], [12]. However, it was not until Bianchi's seminal paper [13] that the study of 802.11 really took off. We briefly discuss this now classic model next.

In order to estimate the saturation throughput Bianchi makes

a number of assumptions. In addition to the ones mentioned above, we may also enumerate assuming a perfect channel (i.e. no losses due to noise in the channel), symmetry (i.e. all stations with the same configuration), but most importantly he assumes what is now known as the *decoupling assumption*. In a nutshell, this means that for every station, the event of a transmission resulting in a collision is Bernoulli distributed with probability p, independent of the history of collisions so far, and of the rest of the stations. Moreover, p is the same for all stations.

With these assumptions he can analyze each station separately. With this end, he constructs a discrete-time markov chain (noted $\{s(t), b(t)\}$), where each state is a twodimensional vector containing the backoff stage and the backoff time counter of the tagged station. Discrete times t and t + 1 correspond to the beginning of two consecutive time slots. It is important to highlight that in this model, a timeslot is defined as the period at the end of which the station can modify its backoff time counter, and it is not the time-slot as defined in the standard. This means that not all time-slots have the same duration, as for instance it may well include a packet transmission.

Assuming p is known, he proceeds to calculate the probability of sending a packet at any given time-slot, noted τ , which is simply the probability of the chain being in a state with the backoff time counter at zero (i.e. $\tau = P(b(t) = 0)$). Surprisingly, the resulting markov chain is amenable to an explicit calculation of its steady-state distribution, resulting in:

$$\tau = \tau(p) = \frac{2(1-2p)}{(1-2p)(W_{\min}+1) + pW_{\min}(1-(2p)^m)}, \quad (1)$$

where W_{\min} and m were defined in the previous section (i.e. W_{\min} is the minimum contention window and $W_{\max} = 2^m W_{\min}$).

Moreover, p may be calculated as the probability that none of the other n-1 stations transmit at the same time-slot, which, since stations are assumed to evolve independently from each other, results

$$p = p(\tau) = 1 - (1 - \tau)^{n-1}.$$
 (2)

The unique pair p^* and τ^* that complies with Eqs. (1) and (2) are then used to calculate the probability of a time-slot containing either a successful transmission (P_s) , a collision (P_c) , or simply being idle (P_i) . This fixed-point analysis has been further simplified and generalized in [14], where for instance the authors show that considering only the backoff time counter is enough to model the system.

It is important to highlight that in all the above calculations, no specific rate nor access mechanism is assumed. These aspects, as explained in the following paragraph, are just considered in the calculation of the slot's mean duration, and thus in the saturation throughput computation. This means for instance that the ratio between the number of collisions and transmission attempts (equal to p^*) is independent of the chosen rate or access mechanism.



Fig. 3: Collision probability (p^*) as a function of the number of stations: simulation and Bianchi's estimation. Simulation results correspond to a data rate of 11Mbps, as all rates yield the same p^* .

With P_i , P_s and P_c in hand, the system throughput can be calculated as the ratio between the mean amount of data transmitted per time-slot and the slot's mean duration. The former is $P_s E$, where E is the payload size. The latter is calculated as follows. Let T_i , T_s and T_c be respectively the duration of an idle time-slot, one containing a successful transmission, and one containing a collision. The first one is fixed by the standard as aSlotTime. The other two depend on the access mechanism. For instance, a time-slot containing a successful transmission under the basic access is equal to the time it takes to send a complete packet, plus aSIFSTime, plus the time it takes to transmit an acknowledgement, plus aDIFSTime. On the other hand, a time-slot containing a collision is equal to the time it takes to send a complete packet, plus aEIFSTime. After establishing T_i , T_s and T_c , the slot mean duration is simply $P_iT_i + P_sT_s + P_cT_c$. Thus, the system total throughput results in:

$$S = \frac{P_s E}{P_i T_i + P_s T_s + P_c T_c}.$$

Interestingly and quiet surprisingly, the predicted collision probability and saturation throughput obtained with these simplistic assumptions is remarkably precise. For instance, in Figs. 3 and 4 we show the comparison between Bianchi's estimation and simulations as we vary the number of stations for both access methods and different data rates. Each point in the simulation curve is the average of ten 60 second-long simulations. Unless otherwise stated, this shall be the case for all simulation results we show in the rest of the article.

Bianchi's original intent was to estimate the saturation throughput. However, the reason behind the apparent validity of the decoupling assumption, as the precision of the estimation so testifies, has been the object of intense debate over these last years. For instance, the authors of [15] analyze, among other, whether the sequence of successful transmissions and collisions form a stochastically independent sequence, and



Fig. 4: Saturation throughput as a function of the number of stations: simulation and Bianchi's estimation. Simulation results are shown with marks, and Bianchi's estimation with lines. Solid lines correspond to basic access, and dashed ones to RTS/CTS.

if this sequence is identically distributed and independent of the collision history. To this end they conduct a simulative and experimental campaign, and conclude that both are reasonably correct in this scenario.

There are other, somewhat more theoretical, studies of this hypothesis. Among them, we may highlight [16], which studies a markov chain where each state is a *m*-dimensional vector indicating the number of stations at each backoff stage. The authors then prove that as time goes to infinity, the system remains close to a typical (or equilibrium) state, the more so as the number of stations increases. As a consequence, each station will "see" the rest of the stations as if they were in this typical state, thus proving the validity of the decoupling assumption.

This last paper uses ideas drawn from the so-called *fluid limits* (see for instance [17], [18]). Instead of studying a complex stochastic process directly, the basic idea behind these techniques is to study an appropriately scaled version of it. If some technical conditions are met, this process will converge to a limit, typically the solution to a differential equation. In the particular example we just discussed, the process is scaled on the number of stations. However, the scaling may be over other parameters, such as time or the intensity of the input process.

There is a related technique, which is used to study complex stochastic systems resulting from the interaction of (infinitely) many agents, which is called *mean field approximations* (see for instance [19]). This involves basically two steps. Firstly, the calculation of the asymptotic proportion of individuals in each state (what is known as the occupation measure limit). One may think as this step as a fluid limit, where the scaling is on the number of agents. Secondly, proving that all agents behave independently from each other and, from the perspective of any given agent, the rest of the system is well approximated by the fluid limit calculated in the previous step. This fact is called decoupling assumption, and is the origin of the term we discussed before in the context of Bianchi's work. The interested reader may consult [20] for a discussion on necessary and sufficient conditions for the decoupling assumption to hold in our particular context based on this technique.

It should be noted that all these theoretical justifications of the decoupling assumption require infinitely many stations. However, as it can be appreciated in Figs. 3 and 4, the estimations obtained by Bianchi's formula are accurate even for a small number of them. This fact, together with its selfcontained nature (no input are required, except for the operative parameters, such as the minimum contention window), has made this model very popular, and several subsequent studies are based on its ideas.

In any case, there are certain qualitative aspects to be learnt from Bianchi's model and Fig. 4. Firstly, as congestion increases, the smaller time wasted on collisions obtained by the RTS/CTS scheme (a RTS packet and a whole data packet are wasted during a collision in RTS/CTS and basic access respectively) compensates the time spent on channel reservation. This results in a relatively constant saturation throughput for RTS/CTS as the number of stations increases, whereas the basic access presents a decreasing one. Moreover, if congestion is high, it may be the case that the RTS/CTS access obtains better results than the basic one. This is an interesting side-effect of a scheme that was originally conceived to mitigate the hidden terminal problem (more on this issue will be discussed later in this article).

Secondly, since the collision probability (p^*) and the attempt rate (τ^*) of all nodes is the same, we may conclude that 802.11 fairly distributes resources in this scenario. It is important to highlight that this is only true on the long run. That is to say, as time goes to infinity, the ratio between the number of packets successfully sent by each station and time tends to be the same for all stations. However, the model provides no guarantee regarding the short-term behaviour of the protocol. In fact, as studied for instance in [21], DCF presents short-term unfairness, specially so as the number of stations increases. The interested reader may also refer to [22], [23] for some recent discussion regarding fairness and its definition.

Lastly, the efficiency of the control access (defined as the ratio between the saturation throughput and the nominal maximum throughput) decreases as the modulation rate increases. For instance, for 2Mbps the efficiency is roughly 0.8, whereas for 11Mbps it is 0.4. This is due to timing aspects that do not change with the data rate. As we discussed before, control packets (such as ACK or RTS) are always transmitted at the same rate, generally either 1 or 2 Mbps. Moreover, stations always wait the same amount of time for events. For instance, a slot-time has always the same duration independently of the data rate.

B. Multi-Rate

In the last section we assumed a perfect channel with no hidden stations. Additionally, we implicitly assumed that the multi-rate feature is disabled in all stations. The algorithm used by stations to choose the rate at which to send its packets is not specified in the standard, and as such, is left vendor specific. The objective of such algorithm is rather straightforward: operate at the maximum possible rate, given the current interference and noise levels.

One of the first such algorithms to be published is *Auto Rate Fallback* (ARF) [24], first proposed in the context of Wavelan-II, a compatible alternative to 802.11. The idea is simple and easy to implement, which is the reason behind its wide adoption by WiFi manufacturers. If after sending a data packet an ACK is not received, a loss event occurs. Else, a success event occurs. When the number of consecutive loss events reaches a certain threshold (say 3), the station falls to the rate immediately below the current one (if any). Conversely, when the number of successive success events reaches another threshold (say 10) or a timer expires, the station sends the next packet at the rate immediately above the current one (if any). If the ACK corresponding to this packet is not correctly received, it falls back to the original rate, else it stays.

In the scenario we are now considering, where no losses are due to noise or interference, but only to collisions, a well designed rate-adaptation algorithm should operate at the maximum possible rate. However, in the ARF algorithm, all ACK timeouts are assumed to be due to interference, and as such, collisions may also trigger a downgrade in the rate. It should come as no surprise then that as congestion reaches a moderate level, the saturation throughput decreases rapidly under ARF in the basic access. On the other hand, the RTS/CTS scheme is almost not affected by ARF since collision do not trigger an ACK timeout, but rather a CTS timeout.

Figure 5 precisely illustrates this fact. In it, we compare the saturation throughput for ARF, under both basic and RTS/CTS access, for the same scenario as before (i.e. UDP traffic saturating the uplink in all stations). Note how, under the basic access, when the number of stations is 20, ARF uses almost exclusively the 1Mbps rate.

This steep throughput degradation for basic access (which, to the best of our knowledge, was first thoroughly studied in [25]) and the qualitatevly different behaviour of the RTS/CTS counterpart has confused some researchers. For instance, the authors of [26] carried out an experimental campaign where they measured the total throughput of a 802.11b WLAN. For example, they concluded that throughput under the scenario considered so far is fairly distributed among stations. However, they report a saturation throughput of nearly 1Mbps for as little as 14 stations, wrongfully blaming DCF for this performance degradation.

Another example is the more recent paper [27], which studies precisely this effect (which they call *rate avalanche*). In it, the authors claim that "*because RTS frame is much shorter*



Fig. 5: Saturation throughput as a function of the number of stations for the ARF multi-rate algorithm.



Fig. 6: Collision probability (p^*) as a function of the number of stations for ARF: Bianchi's model and simulations.

compared to most data frames, thus the collision probability of RTS frames is much lower.". Although, as shown in Figs. 4 and 5, in some scenarios RTS/CTS leads to more throughput than its basic counterpart, the ratio between collisions and transmission attempts is the same with or without RTS/CTS. This fact, which as we discussed was already proved by Bianchi, is illustrated in Fig. 6, where we show this ratio for the simulations corresponding to Fig. 5. Please note that, as expected, it is exactly the same as Fig. 3. Actually, the difference between the basic access and RTS/CTS under ARF is that (as observed and further exploited in [28]) the loss counter is incremented only when an ACK timeout occurs (which may only happen if a data packet was sent) and not after a CTS timeout. This is yet another unforeseen consequence and benefit of the RTS/CTS four-way handshake.

We now briefly discuss how to predict the saturation throughput in this case. It should be clear from the discussion presented in the last paragraph that Bianchi's analysis can be extended to this case. Picking up the notation of Sec. III-A, we may safely assume that even under ARF (or any other rate selection algorithm), p^* and τ^* are the ones calculated by Eqs. 1 and 2. Then, the probability that a given time-slot contains a collision (P_c), a successful transmission (P_s), or is idle (P_i), are the same as in Sec. III-A. The only quantities we need to calculate in this case are then the duration of each of these time-slots (i.e. T_c , T_s , T_i).

An idle time-slot always lasts aSlotTime, independently of the modulation rate (i.e. T_i =aSlotTime). Moreover, the control rate is not modified by ARF, so the duration of ACKs, RTS and CTS packets are the same as before too. The same goes for aDIFSTime, aSIFSTime and aEIFSTime. The only aspect of the calculation of T_s and T_c that has to be modified from the one carried out in Sec. III-A is the time it takes to send a complete packet, both under a successful transmission or a collision.

If, given the packet loss probability (in this case p^*), it was possible to calculate the mean of the inverse of the modulation rate chosen by the algorithm for each packet transmission ($\mathbb{E}\{1/R\}$), the time it takes to send a complete packet may be calculated as the physical preamble and header plus ($H_{\text{mac}} + E$) $\mathbb{E}\{1/R\}$ (where H_{mac} is the size of the MAC header and recall that E is the payload size).

The calculation of the mean duration of a complete packet under a collision requires more information. Indeed, the duration of such time-slot is dictated by the station using the least rate. Let N(t) be a random variable indicating the number of stations attempting to transmit at time-slot t. Then, $\mathbb{E}\{1/R_{\min}^{col}\}$, the mean of the inverse of the minimum rate among this N(t) colliding stations, is:

$$\mathbb{E}\left\{1/R_{\min}^{\text{col}}\right\} = \mathbb{E}\left\{\max_{k=1,\dots,N(t)} \frac{1}{R_k} \middle| N(t) \ge 2\right\} = \sum_{i=2}^n \mathbb{E}\left\{\max_{k=1,\dots,i} \frac{1}{R_k}\right\} \frac{P(N(t)=i)}{P(N(t)\ge 2)}.$$

Please note that N(t) is a random variable that follows a binomial distribution with parameters n and τ^* . If the distribution of the rate chosen by each colliding station (i.e. R_k , which is distributed as R) is known, it is then straightforward to calculate the mean of the maximum inside the addition, and thus $\mathbb{E}\{1/R_{\min}^{col}\}$. Finally, the time it takes to transmit a complete packet under a collision results in $(H_{\max}+E)\mathbb{E}\{1/R_{\min}^{col}\}$ plus the physical preamble and header.

In the case of ARF, since packets collide with probability p^* , we may model its behaviour with a markov chain, and we may thus calculate the distribution of the chosen rate at any given time-slot. The interested reader may consult [29] for the details, although it is important to highlight that the model does not consider the timer used to probe higher rates, nor our simulations, even if its impact is not significant at all. The result of the analysis we just described applied to ARF is shown in Fig. 5. Please note that we obtain remarkably good predictions.

All in all, unless very carefully designed, using automatic

rate selection algorithms is generally not a good idea. The next section shows more arguments in this sense for UDP traffic, and Sec. IV-A1 for TCP traffic.

C. Asymmetry: Performance Anomaly

Until now we have considered perfect symmetry among stations. That is to say, the generated traffic, the DCF parameters, the multi-rate algorithm, and the propagation conditions, are the same for all stations. However, we have also discussed that in order to calculate the probability that a tagged station attempts a transmission at any given time-slot (i.e. τ^*), and the probability that such transmission collides (i.e. p^*), we may use Eqs. 1 and 2, independently of the duration of the slots. This means that these values will be the same whatever the packet size or modulation rates chosen by stations, even if they are different between stations. Bianchi's model allows us then to analyze asymmetric situations, that although somewhat limited, are nevertheless very interesting as we shall now discuss.

1) Different Rates: Let us first consider a case-scenario where stations transmit at fixed, but different modulation rates. In particular, all stations will transmit at a high modulation rate (R, which may for instance be 11Mbps), except for one which transmits at a low rate (r, e.g. 2Mbps). Such situation may arise when a SNR-based multi-rate algorithm is enabled, and one of the stations is exposed to an important level of interference or the path-loss between this particular station and the AP is much bigger than for the rest. Moreover, we will assume that, at least on the long-run, transmissions only fail due to collisions (meaning that the multi-rate algorithm is "perfect" in the sense that it finds a suitable rate considering only noise levels).

To calculate the saturation throughput in this case we proceed analogously to the previous section. Thus, we only need to calculate the mean of the time it takes to send a complete packet, both during a successful transmission and a collision (only required for the basic access, since under RTS/CTS no packet is sent during a collision). To this end, we will calculate the mean of the inverse of the smallest rate under both scenarios. The mean of the time it takes to send a packet is then the physical preamble and header, plus the size of the packet payload times this mean.

Since all stations transmit equiprobably, the inverse of the mean rate under a successful transmission is equal to 1/rn + (n-1)/Rn. The mean of the inverse of the minimum rate among colliding stations may be calculated as in the last section as follows:

$$\mathbb{E}\{1/R_{\min}^{\text{col}}\} = \sum_{i=2}^{n} \mathbb{E}\left\{\max_{k=1,\dots,i} \frac{1}{R_k}\right\} \frac{P(N(t)=i)}{P(N(t)\ge 2)} = \sum_{i=2}^{n} \left(\frac{i}{n}\frac{1}{r} + \left(1 - \frac{i}{n}\right)\frac{1}{R}\right) \frac{P(N(t)=i)}{P(N(t)\ge 2)}.$$

The comparison between the simulations and our analysis is shown in Fig. 7. Again, the estimation obtained with the means described before is very accurate. In addition to comparing



Fig. 7: Modulation rate asymmetry case-scenario. All stations transmit at a rate of 11Mbps, except one that does it at 2Mbps.

simulation to analysis, we have also included the analytical results obtained when all stations transmit at either 11Mbps or 2Mbps. It should be clear that as n increases the results tend to be the same under this scenario and when all stations transmit at 11Mbps. The graph also displays the throughput obtained by an arbitrary station at 11Mbps, and the station at 2Mbps. Please note that, as expected, both throughputs are the same.

This particular case-scenario, first studied by [30], has gained much attention in the last years, being termed as performance anomaly in the literature. This anomaly is the important performance degradation that occurs due to a single station choosing a very small modulation rate, as illustrated in Fig. 7. For instance, at 10 stations, the saturation throughput obtained when one station uses the 2Mbps modulation is 20% smaller than when all stations use the 11Mbps one. The reason behind this degradation is simply the fairness imposed by DCF under this scenario, i.e. all stations have roughly the same number of transmission attempts. Consequently, alternative long-term fairness objectives have been proposed in the literature, being time-based fairness the most prominent one [31], [32]. That is to say, all stations having roughly the same time usage of the channel. Methods to achieve this fairness without drastic changes to DCF are for instance discussed in [33].

The general conclusion of this case-scenario is similar to the one we made in Sec. III-B. It is generally a good idea to avoid the co-existence of several different modulation rates in the same WLAN. Unless forced to (if for instance the coexistence of 802.11g and b stations is necessary), probably the best way to proceed is to choose a single modulation rate, and those stations whose channel conditions are such that they cannot communicate with the AP, should not be allowed in the WLAN (effectively using the chosen rate as an admission control mechanism).

2) Different Packet Sizes: We now discuss a case-scenario where stations transmit payloads of different sizes. Such situ-

ation is very common, specially with multimedia applications, as VoIP, where using the maximum possible payload is not always possible. In particular, we will consider two opposite case-scenarios. In both, all stations except for one transmit using a payload equal to 988 bytes. However, in the first casescenario, the last station uses a payload of size 2028 bytes, and in the second one, the station uses a payload of size 328 bytes.

In order to predict the saturation throughput in this case, we may proceed similarly to Bianchi and the previous subsections, except for some necessary modifications. Firstly, the mean amount of data sent per time-slot is P_s times the mean of the transmitted payload size. We have already mentioned that all involved probabilities (i.e. P_s , P_c and P_i) are again calculated with Eqs. (1) and (2), since they do not depend on the duration of the slots, and are thus independent of the packet size. The mean of the transmitted payload size (which we shall note as $\mathbb{E}\{E\}$) is simply $E^1/n + E^{n-1}(n-1)/n$, since all stations still have the same access opportunities (where $E^{n-1} = 988B$ and $E^1 = 2028B$ or 328B).

The time consumed by an idle time-slot is still aSlotTime. The mean of the time consumed by a successfully sent packet is now $T_s = H_{\text{mac}}/R + (E^1/n + E^{n-1}(n-1)/n)/R$, plus the physical preamble and header. The mean of the time it takes to send a complete packet under a collision is the mean of the maximum packet size among the colliding stations, divided by the rate. Proceeding very similarly to the previous sections, we may calculate this mean as:

$$\mathbb{E}\left\{\frac{E_{\max}^{\text{col}}}{R}\right\} = \frac{1}{R}\mathbb{E}\left\{\max_{k=1,\dots,N(t)} E_k \middle| N(t) \ge 2\right\} = \frac{1}{R}\sum_{i=2}^n \mathbb{E}\left\{\max_{k=1,\dots,i} E_k\right\} \frac{P(N(t)=i)}{P(N(t)\ge 2)} = \frac{1}{R}\sum_{i=2}^n \left(\frac{i}{n}\max\{E^1, E^{n-1}\} + \left(1 - \frac{i}{n}\right)E^{n-1}\right)\frac{P(N(t)=i)}{P(N(t)\ge 2)}$$

Note that we have omitted the MAC header for the sake of clarity of the presentation, as it simply adds a constant in the above calculation. We now turn our attention to calculating the throughput each of the stations obtain out of this total. We already know that the number of successful transmissions is roughly the same for each station. However, the "lonely" station will transmit E^1 bytes during each of these transmissions, whereas the rest of the stations shall transmit E^{n-1} . It is then straightforward to see that the ratio between the throughput obtained by the "lonely" station and the one obtained by any other one is E^1/E^{n-1} .

A comparison between simulation and analysis is shown in Fig. 8. Again, the predictions yielded by the analysis provide very accurate results. It should also be noted that a station using big packets may obtain a very important advantage over the rest of the stations. There is thus an incentive for stations to use the biggest possible packets. On the other hand, when one of the stations uses a small packet size, the overall throughput decreases, just like in the case of a station using a smaller rate. However, there are two very important differences. Firstly,



Fig. 8: Different packet size asymmetry case-scenario. All stations transmit a payload of size 988 bytes, except one that does it either at 2028 bytes (above) or 328 bytes (below).

the difference in the system throughput between this and the completely symmetric case decreases very rapidly as the number of stations increases. Secondly, the rest of the stations do not see their performance hindered by the presence of this station (actually, they obtain a small improvement).

IV. UNSATURATED STATIONS: TCP OVER WLANS

We have so far considered saturated stations; i.e. a casescenario where stations always have packets to send to the AP. Although as we discussed before, it is interesting in its own right, the saturated condition somewhat simplifies the analysis in the sense that traffic is readily modelled. The question that arises when studying the unsaturated case-scenario is then what traffic should the stations inject to the network.

A very interesting case is TCP. Although real-time traffic over UDP has increased its presence in the last years, TCP is yet by far the predominant transport layer protocol in any network, and WLANs are no exception. Consequently, academia has studied its performance over wireless networks, and in particular over 802.11. The first articles to discuss the performance of TCP over wireless networks (and over 802.11 in particular), were rather pessimistic [34], [35]. Indeed, at the time the wireless medium was considered to be imperfect and prone to errors. Moreover, since TCP throughput depends on the packet loss probability roughly like $1/(RTT\sqrt{p})$ [36] (where RTT is the roundtrip time of the flow), it was only natural to assume (and measurements so indicated) that TCP performance would be very poor. However, we have so far assumed the opposite; i.e. that the wireless channel is perfect and no packets are lost due to path-loss or interference. These assumptions hold for instance in an indoor scenario where the AP is close to the stations, and if the AP is relatively close to other APs, they choose different channels (a feature present in most APs, where the least congested channel is chosen periodically).

However, noise and path-loss are not the only possible reason behind TCP's throughput degradation. As Fig. 3 illustrates, the probability that a transmitted packet collides with another concurrent transmission is not negligible (at least in the up-link saturated case-scenario). The resulting re-transmissions would mean bigger delays and thus bigger round-trip times.

A. Down-link Transfers

Let us first consider the down-link traffic case-scenario. That is to say, n stations are each downloading a file of infinite size from a server behind the AP. We will further assume that all stations use the same TCP flavor (in particular, we have used NewReno [37]), and that SACK or delayed ACK are disabled. Moreover, we will also assume that no packets are dropped of the AP's buffer. This holds for instance if the AP's buffer size is bigger than the sum of the maximum TCP window of all concurrent connections.

Since all data packets are transferred from the AP to the stations, a hasty analysis may lead us to conclude that no collisions occur because the AP competes only with itself. However, although small in size, stations do send a TCP acknowledgement for every data packet they receive. We have discussed in previous sections how packet collisions depend exclusively on timing, rather than size. Moreover, even if no packets are lost, numerous collisions will result in bigger delays, which negatively impact TCP performance too.

Let us then make a careful qualitative analysis of this casescenario. Firstly, if no packets are lost in the AP's buffer, and assuming that the probability of a station reaching the MAC retry limit for any given packet is very small (as we have done so far), we can safely assume that the AP always has a packet ready to be sent, and constantly contends for the channel. Secondly, the buffer of any given station only increases when it receives a data packet (since it has to answer with a TCP acknowledgement). That is to say, when the AP gains access to the channel and does not collide with any other station. Now, let us assume that several stations have acknowledgements ready to be sent. Since DCF fairly distributes access opportunities, the probability that specifically the AP sends a new data packet (and potentially increases the number of "active" stations) is very low. Thus, the system will have a



Fig. 9: Ratio between the successful transmissions and the transmissions attemps for the stations (i.e. TCP acknowledgements packest) and the AP (i.e. data packets). Transmissions are modulated at 11Mbps (similar results are obtained for other modulation rates).

very strong drift towards decrementing the number of stations that contend for the channel.

The above analysis (which, to the best of our knowledge was first presented in [25]) leads us to conclude that the actual number of stations contending for the channel remains very low. This will in turn result in a relatively high throughput and very few collisions. This analysis is validated by the simulation results we show in Figs. 9 and 10. The former shows the ratio between the number of successful transmissions and total transmission attemps, for each individual station and the AP, as the number of stations varies. The latter shows the resulting throughput, measured as the ratio between the number of bytes successfully received by all stations (including TCP and IP headers, and a payload of 1460 bytes) and the total simulation time.

Interestingly enough, not only do collisions and throughput remain very low and relatively high respectively, even when the number of stations is high, but they remain roughly constant. This means that TCP over WLANs scales very well with the number of stations (more precisely, the per-station throughput scales as 1/n, very much like in wired mediums). Moreover, the very low collision rate results in ARF having little or no effect in this case-scenario. Finally, note that the stations and the AP obtain roughly the same collision probability. We could not observe the unfairness discussed in [38], which reports that in this same case-scenario the proportion of collided packets are more than twice for the stations than for the AP.

Predicting the performance of this case-scenario is not an easy task. As we mentioned before, [25] has been one of the first studies to highlight and explain quantitatevly the very graceful performance of TCP over WLANs (although prior reports of such behaviour exist, such as [39], [40]). An ana-



Fig. 10: Total throughput when n stations are downloading an infinitely long file from the same server for different values of the (fixed) modulation rate used by all stations.

lytical model (*à la* Bianchi) by the same authors was presented in [29], although in a more general setting than in our case. An alternative analysis is presented in [41], that uses the now classic network utility maximization framework [42] to predict TCP throughput in this case. This allows the proposed model to be very flexible, and for instance can accommodate different flavours of TCP or connections with different round-trip times.

Here we will limit ourselves to presenting a very simple analysis, although it will help explaining certain important aspects of TCP performance. As we mentioned before, for every data packet sent by the AP, there is a corresponding TCP acknowledgement, so the total time to transmit the payload must consider both packets. Since the AP always has a packet ready to be sent, it will run the backoff mechanism for every data packet it sends. Assuming no collisions, the mean time spent in backoff will be approximately $(W_{\min} - 1)/2$ slots. Due to the symmetry of the considered case-scenario, the destination of this data packet may be any of the stations, with equal probability. As we discussed before, the number of stations that are contending for the channel (i.e. "active") is very small. This means that the probability that the receiving station performs the backoff mechanism before sending the corresponding TCP acknowledgement is very small. This results in the following estimation of the system throughput:

$$S = \frac{E}{\text{aSlotTime} \times (W_{\min} - 1)/2 + T_{\text{data}} + T_{\text{tcp-ack}}}, \quad (3)$$

where T_{data} and $T_{\text{tcp-ack}}$ is the time it takes to send a complete packet (data and TCP acknowledgement respectively), including the MAC acknowledgement, and is calculated as in Sec. III-A.

The prediction obtained by this analysis is compared against simulations in Fig. 10. It should be clear that Eq. (3) constitutes an upper bound to the system throughput. However, the gap between analysis and simulation is remarkably small.



Fig. 11: System and stations throughput for the TCP downlink case-scenario. Modulation rate is controlled by ARF, and the channel conditions are such that all stations transmit mostly at 11Mbps, except for one that transmits mostly at 1Mbps.

Moreover, Eq. (3) highlights the influence of TCP acknowledgement packets on the total throughput. For instance, for 11Mbps, T_{data} is approximately 1.7ms, whereas $T_{tcp-ack}$ is 0.6ms. In order to improve the obtained performance, some works have proposed TCP variants tailored for wireless networks [43], [44], where the receiver does not answer each data packet with an acknowledgement, but rather delays it until receiving a certain number of data packets (a scheme known as *delayed ack*, already present, although in a simpler form, in the original RFC [45]).

1) TCP's Performance Anomaly: To conclude the downlink case-scenario we will discuss the performance anomaly presented in Sec. III-C1 in this case. We let the stations and the AP use ARF to control the modulation rate used for each transmission. All stations count with excellent propagation conditions towards (and from) the AP, except for one, whose path-loss from the AP is such that only packets modulated at 1Mbps are correctly received. Due to the very few collisions that TCP over WLAN experiences, this scenario will result in all stations using (most of the time) the 11Mbps modulation rate, except for one which will (again, most of the time) use the 1Mbps one.

Simulation results for this case-scenario are presented in Fig. 11. We have also included, as a reference, the results of the analysis we discussed before for the case where all stations modulate at 11Mbps. The performance penalty when a single station is modulating at a lower rate than the others is very significant, even more than in the UDP case considered in Sec. III-C1. Moreover, it should be noted that all stations receive the same throughput, and as before this fairness is the reason behind the performance degradation.

Regarding analytical means to predict this performance, the original paper that coined the term *performance anomaly* [30] included an analysis of both the TCP and UDP case, similar

in spirit to the one we discussed in Sec. III-C1. Another possibility is presented in [41], which as we mentioned before uses the network utility framework and allows to study this case-scenario.

We will again limit ourselves to presenting a very simple model, which will nevertheless serve for illustrative purposes. Let R_i be the modulation rate used by station i (i = 1, ..., n). Then, S_i , the throughput station i would obtain if it was alone, may be calculated as:

$$S_i = \frac{E}{\text{aSlotTime} \times (W_{\min} - 1)/2 + T_{\text{data}}(R_i) + T_{\text{tcp-ack}}(R_i)}$$

where $T_{\text{data}}(R_i)$ and $T_{\text{tcp-ack}}(R_i)$ are defined as before, except that the dependence on the used rate is made explicit. Then the total system's throughput will be:

$$S = \frac{1}{\sum_{i=1}^{n} \frac{1}{n} 1/S_i}.$$
(4)

That is to say, the packet size divided by the mean amount of time it takes to send a complete packet. Since, as observed, all stations have the same access opportunities and there are negligible collisions, this is simply the mean of the inverse of the S_i (i.e. the harmonic mean of the S_i). Again, although an upper-bound, Eq. (4) is a very tight one, as shown in Fig. 11.

B. Up-Link Transfers

Let us now turn our attention to the up-link case-scenario. That is to say, *n* stations are transferring an infinitely long file to a server behind the AP. In this case then, the AP sends TCP acknowledgements, and stations send data packets. If we make the same assumptions as before, the analysis we performed is still valid in this case. A station is allowed to send a new data packet only after having received a TCP acknowledgement, which may happen only if the AP successfully gains access to the channel and the particular acknowledgement has it as its destination. Thus, the strong drift towards few "active" stations is still valid, and we obtain the same results as before, as Fig. 12 so shows.

From all the assumptions we made before, there is one in particular that has received much attention from the academia: the infinitely big AP buffer. Let us consider for instance a casescenario with 10 stations, and an AP with a buffer capacity of only 20 packets. Figure 13 shows the throughput obtained by each station for a particular simulation run. The unfairness in this case is remarkable, as the number of articles referring to this situation ([46], [47], [48], [49] just to name a few).

The reason behind this unfairness lies on the asymmetry of the path from the stations to the server. Whereas a data packet is rarely dropped (the only possibility of such event occurrence is that the maximum MAC retry limit is reached by the station), acknowledgements that transit the reverse path are more frequently dropped by the AP due to its limited buffer and saturated condition. Persistent loss of TCP acknowledgements hinder the performance of TCP flows and result in such unfairness phenomena, an effect previously studied for wired networks (see for instance [50], [51]). A qualitative explanation is the following. TCP flows with a small congestion



Fig. 12: Total throughput when n stations are uploading an infinitely long file to the same server for different values of the (fixed) modulation rate used by all stations.



Fig. 13: Throughput obtained by each station. Up-link TCP traffic, AP's buffer size equal to 20 packets.

window (i.e. that have recently started or suffered a timeout) may only send very few packets. Thus, acknowledgements corresponding to these flows being dropped by the AP will most probably result in time-outs. On the other hand, the same number of losses for a flow with a big congestion window (i.e. many packets in flight) will most probably result in a triple-ACK event, and will have a much smaller impact on the flow's throughput. This situation is only exacerbated in time, resulting in a few lucky flows capturing the channel for themselves.

C. Up-Link and Down-Link Transfers

The next case-scenario we will consider is a mix of the previous two. From a total of n stations, n_{up} are transferring an infinitely long file to a server behind the AP, whereas the rest is downloading an infinitely long file from the same server.



Fig. 14: Throughput obtained by each station. Down-link and up-link TCP traffic, AP's buffer size equal to 20000 packets. The total number of stations is 20, and we vary the number of up-link stations from 0 to 20. The lower indexes correspond to these stations.

If we still assume no losses in the AP's buffer, then we will obtain fairness between flows, since both uplink and downlink stations have to receive a packet from the AP in order to send a new one themselves. That is to say, uplink stations will send a new data packet when they receive a TCP acknowledgement, and downlink stations will send a TCP acknowledgement when they receive a data packet.

This fairness is illustrated by Fig. 14. In it we have fixed the number of stations to 20, and have varied the number of up-link stations from 0 to 20. The graph shows the individual throughput obtained by each station in a single simulation, where the lower indexes correspond to the up-link stations (if any). Note that all stations, independently of being up-link or down-link, obtain roughly the same throughput, and that it does not depend on the amount of up-link flows neither.

However, if the AP were to have a small buffer, then the same unfairness we observed before would occur. In addition to a few uplink flows capturing most of the channel (for the same reasons we discussed before), downlink flows are almost starved, a phenomenon we now qualitatively explain. Firstly, it should be noted that uplink flows may only loose acknowledgement packets. If one such packet is lost, but another one reaches the sender fast enough, it will acknowledge all previous data packets. However, downlink flows are prone to loosing data packets. Such event will trigger either a triple ACK event or a time out, both of which result in a drop in the congestion window and consequently in the flow's throughput. Thus, packets discards at the AP are much more serious for down-link flows than up-link ones.

Figure 15 is an example of such situation. In this case, the AP's buffer is equal to only 20 packets, and there are a total of 20 stations, 10 of which are generating up-link TCP traffic, and the rest down-link. Figure 15 shows the throughput obtained



Fig. 15: Throughput obtained by each station. Down-link and up-link TCP traffic, AP's buffer size equal to 20 packets. The total number of stations is 20. The first 10 indexes correspond to up-link stations.

by each flow, where the first 10 indexes correspond to the up-link ones. Note how four flows obtain all of the system's throughput, whereas the rest are totally starved.

Several papers discuss this unfairness issue, of which we have already cited a very small list [46], [47], [48], [49]. In addition to presenting the problem, they also propose solutions other than augmenting the AP's buffer, which does not scale and may be unfeasible for low-end equipment. For instance, [46] proposes a mechanisms where the AP manipulates the receiver window of all ongoing flows, so as to throttle the uplink flows' throughput. Since per-packet header parsing and modification may be prohibitively costly or simply not possible (for instance, if end-to-end encryption is used), the authors of [49] propose a rate limiting scheme at the AP based on token buckets. A somewhat simpler alternative is proposed in [47], [48], whose authors use the QoS differentiation provided by 802.11e in order to guarantee unrestricted access to the medium for TCP acknowledgement packets.

D. TCP and UDP traffic

To finish this section we briefly discuss the performance of co-existing TCP and UDP flows in a WLAN. The first thing that should be clear is that down-link UDP flows are perceived as a single flow by DCF due to the per-station fairness it enforces. We will then only consider that there are either one or no down-link UDP flows. Moreover, it should be clear that, as long as no packets are dropped at the AP, up-link and downlink TCP flows are equivalent. We shall then only consider down-link TCP flows.

1) No down-link UDP flows: Let us first consider the case with no down-link UDP flows. In particular, we will study the interaction between a fixed number of stations downloading an infinitely long file from a server behind the AP, and an increasing number of stations sending UDP traffic towards the



Fig. 16: Throughput when ten stations are downloading an infinitely long file from a server behind the AP, and n_{up} stations are sending UDP traffic in saturation conditions. Results for the case when no TCP flows are present are included as a reference.



Fig. 17: Collision probability when ten stations are downloading an infinitely long file from a server behind the AP, and n_{up} stations are sending UDP traffic in saturation conditions. Results for the case when no TCP flows are present are included as a reference.

AP. Figures 16 and 17 show the results obtained for ten TCP stations and n_{up} stations sending UDP traffic (with n_{up} varying from 0 to 40).

As expected, the throughput obtained by the TCP flows decreases drastically with the presence of UDP flows. As we mentioned before, due to the DCF per-station fairness, the AP has to share the channel fairly with all the other competing stations. Thus, in this case it will obtain a fraction smaller than $1/(n_{\rm up} + 1)$ of the total channel access opportunities (since stations downloading the file also compete for the channel when they respond with a TCP acknowledgement

the actual fraction will be smaller). It is logical then that, as n_{up} increases, the throughput obtained by the UDP stations approaches the one obtained when no TCP stations are present.

On the other hand, the results corresponding to the collision probability are surprising. We have computed the ratio between the number of collisions and transmission attempts, both for the stations and the AP. Moreover, in the former we have distinguished between TCP stations (that transmit TCP acknowledgements) and UDP stations. Figure 17 shows the simulations results, in addition to the results obtained when no TCP station is present, which we will use as a reference.

Firstly, it should be noted that the presence of the TCP stations increases the UDP stations' collision probability only slightly. The explanation for this is the same as in the previous section. TCP stations only become active (i.e. compete for the channel) when they receive a data packet from the AP. Since there are 10 TCP stations, then there will an important drift towards few active TCP stations. Thus, the UDP stations will compete mostly only with the AP (and naturally among themselves too).

The second interesting thing to note is that, differently to the results we obtained until now, not all stations obtain the same collision probability. In particular, the AP obtains a slightly (but persistently) bigger collision probability than the UDP stations, whereas TCP stations obtain a much smaller one. The reason behind these differences is subtle but important. Let us consider a time-slot where the AP gains access to the channel. The data packet may be destined to a station that is either waiting to send a previous acknowledgement (i.e. already executing the backoff mechanism) or not. In the latter case, immediately after having received the data packet from the MAC layer, the TCP layer of the tagged station will send in return the TCP acknowledgement to the MAC layer. Since only one out of ten AP transmission are data packets destined to the tagged station, it is not improbable that the latter's backoff counter has already expired (the more so as congestion and retransmissions increase, as it takes more time to the AP to send each data packet). In such case, the station senses the channel for a DIFS period, and if idle, sends the TCP acknowledgement. It is precisely at this moment that all other stations are allowed to "defreeze" their backoff counter, or draw a new one in the case of the AP. Now, the backoff counter of all other stations must be bigger than zero (else, they would have collided with the AP in the first place) and thus will not transmit until after the TCP acknowledgement is sent. However, if the AP draws a backoff equal to zero, it will transmit at the same moment than the tagged station, and their packets will collide.

All in all, TCP stations may, in most of the cases, only collide with AP, and this will happen with a probability of $1/W_{\min}$ (i.e. when the drawn backoff counter is equal to zero). If the data packet was destined for a station that is already contending for the channel, it will obtain the same collision probability as the UDP stations. Moreover, it should be clear by now that any station that draws a backoff counter equal to zero after a successful transmission will be able to send

another packet immediately, and that this transmission will not collide². This is not the case for the AP, for which drawing a backoff counter of zero after a successful transmission will most probably result in a collision due to the receiver responding with a TCP acknowledgement. This in turn explains the slight difference between the collision probability of the AP and the UDP stations.

Regarding more quantitative analysis of this case-scenario, to the best of our knowledge there does not exist a model to predict the results obtained in this case. Such model should be carefully constructed to consider the details we discussed before, which result in the invalidity of one of the most important hypothesis in all the analysis we have carried out until now.

2) One down-link UDP flow: Let us now consider the same case-scenario, but with the addition of a station that is receiving an UDP stream from a server behind the AP. The presence of such flow will result in an increase in the drop rate at the AP's buffer, whose terrible effects on TCP's performance we have already discussed. We will consider an extreme case here, where the down-link UDP flow has the same rate as the up-link ones (i.e. if alone, it would saturate the AP). It should come at no surprise that the resulting throughput distribution is such that TCP flows are starved, and the UDP ones obtain a throughput similar to the one they would obtain if TCP flows were not present. Figure 18 shows the throughput obtained by each flow (in a single simulation run) in the case of 10 TCP flows, 3 up-link UDP flows, and a single down-link UDP flow (indexes in the abscissa correspond to this order). Note how the results are similar to the case where only four UDP up-link flows are present and saturating the channel.

Although somewhat extreme, we have included these examples of UDP and TCP co-existence in a WLAN to illustrate on the unfairness the flows are prone to suffer. In the case where a UDP flow exists in the down-link direction, the situation is not very different from the wired case, and a single UDP flow can starve all TCP flows. However, if UDP flows only exist in the up-link direction, then they share the channel access opportunities fairly with the AP. Although TCP flows still are significantly affected in their performance, they are not totally starved.

V. A MORE REALISTIC PHY: PHYSICAL CAPTURE

Up to this point, we have considered a very simple model for the physical layer. In particular, the receiving station's PHY has been abstracted to two cases. Either a single transmitter uses the channel, in which case the frame is correctly received, or more than one transmitter use the channel, in which case all frames are incorrectly received.

Although simple, this model has allowed us to gain much insight of DCF. However, the real receiving process is much more complicated. In particular, it is not difficult to imagine a situation in which a station has a path-loss to the AP that

²Please note that such situation is not considered in Bianchi's original model. The interested reader is referred to [52] for details on how to include it in the model.



Fig. 18: Throughput obtained by each station. 10 down-link TCP flows, 3 up-link UDP flows and 1 down-link UDP flow. AP's buffer size equal to 20000 packets. Flow indexes correspond to this order.

is much bigger than the rest. If this station transmits at the same time as any other one, the AP may be able to correctly decode the other transmission, as the tagged station's signal will appear to the AP receiver as weak noise.

The situation we just described, where even if more than one station transmits, the receiver may still decode the stronger frame, is called the *Physical Layer Capture*, or channel capture, or simply capture. This problem is particularly serious in multiple access mechanisms based on spread-spectrum, where it is called the *near-far problem* [53]. In such systems, where all stations share both time and frequency, a single station may effectively capture the channel for itself if its transmission power is high enough, making power-control a necessity.

In DCF, and in every other CSMA/CA system, the capture phenomena may manifest itself only when collisions occur. It is important to highlight that the term collision now refers to the event of two or more transmissions overlapping in time. This may lead to the receiver dropping all frames, or correctly decoding one. In the latter case we will say that the frame captured the channel. In the following subsections we will discuss under which conditions capture is possible in 802.11 systems, and what are its effects on the upper layers.

A. Modelling Physical Capture in 802.11 Systems

We now briefly discuss how to model the capture phenomena; i.e. under which conditions a frame is correctly decoded even if it overlaps with one or more concurrent transmissions. Most of the earlier works on the subject are based on a so called capture ratio. That is to say, a specific frame will be correctly decoded at the receiver if its power divided by the sum of all other concurrent transmissions is bigger than a certain threshold (see for instance [54]). Other works additionally considered that it was enough for this condition to hold during a (capture) time window [55]. These earlier models have been used to expand Bianchi's model so that it considers this phenomena. For instance, the authors of [56] re-calculates the probability that a given time-slot contains a successful transmission by including the probability that, although more than one station transmitted, the tagged station's frame captured the channel (i.e. the AP receives the tagged station's transmission with a power such that its division by the total power of all other concurrent transmissions is bigger than a certain threshold).

However, and somewhat in parallel with these theoretical developments, some of the first experimental campaigns were showing an unfair behaviour of DCF. For instance, the one reported in [57] showed that when two stations transmitting to the AP do not sense each other, a difference of a few dB in their received power will result in the specific station capturing the channel when their transmissions collide, and thus in different successful probability for the two stations. This would limit the applicability of Bianchi's model (and most extensions based on it), since a difference of a few dB in path-loss is not rare and would break its symmetric assumption.

One could argue that two stations hidden from each other was out of Bianchi's model hypothesis from the start. However, the experiments conducted in [58] extended these results by showing that it was not the fact that the stations were hidden from each other that generated the unfairness, but that the few dB in receiving power were enough to generate this imbalance.

Two very important observation are made on reference [58]. Firstly, that even if the stronger frame arrives second, it is able to capture the channel. However, if the receiver was able to correctly detect and synchronize to the physical preamble of the first weaker frame, the second stronger frame is dropped. Secondly, they observed that the propagation delay does not dictate the order of arrival of the two frames at the receiver, even if they are sent at the same time. Actually, they observed time differences of several microseconds between two frames, where the arrival order is random and is influenced by clock desynchronizations, RX/TX turnaround times or processing times in the stations.

To date, and to the best of our knowledge, the most complete and detailed measurement and analysis of the capture effect for 802.11 is carried out by the authors of [59], later extended in [60], the results of which we now briefly discuss. Capture scenarios are classified in three categories depending on the timing between the frame of interest (FoI) and an interfering frame: either the FoI arrives first, or second and the receiver was already receiving the interfering frame, or second but the receiver was not receiving the interfering frame (i.e. the receiver has not detected nor synchronized to the interfering frame's preamble).

The Signal to Interference Ratio (SIR, in this case defined as the difference between the FoI's and the interferer frame's power at the receiver) necessary for the FoI to capture the channel with a big probability (say, more than 0.9) is very different in the three scenarios. In the first one, it further depends on the precise timing between frames. If the interfering frame arrives before the receiver has synchronized to the preamble of the FoI, the necessary SIR is approximately 4dB. After having synchronized, it is as little as 1dB. In the second casescenario, where the FoI arrives after the interfering frame, the former captures the channel as long as the SIR is more than roughly 10dB. Naturally, this threshold is bigger than in the previous scenario. Moreover, and differently to [58], it was noted that the chipset used in their experiments was able to capture the FoI even if it arrived after the interfering frame's preamble. The last scenario's results are similar to the one where the interfering frame arrives within the FoI preamble. It is important to highlight that all these results were obtained when using the 6Mbps modulation rate in 802.11a. Naturally, as the modulation rate used to transmit the FoI increases, so does the SIR required to capture it. According to the results presented in [60], this growth is roughly linear.

Most simulators do not consider these three possible capture scenarios and their different SIR thresholds. In particular, the ns-3 reception procedure is as follows [61], [62]. The PHY layer can be in three possible states: transmitting, receiving or idle. If the first bit of a new packet is received while the PHY is not in the idle state, it is immediately dropped. Otherwise, and if the frame's power is higher than a certain sensitivity threshold, the PHY switches to the receiving state, and schedules an event at the moment the last bit of the packet is expected to arrive. To decide whether or not the frame is correctly received, a random uniform number is drawn and compared to a packet error rate. The latter is calculated based on the bit error rate of each received bit, which in turn depends on the modulation used and possible interfering frames. The most important consequence of this simplified PHY model is that the simulations will show less capture events than reality. More precisely, of the three possible scenarios discussed before, a ns-3 simulation may only include the first and third one, and in a relatively simple form. The simulation results shown hereafter will then be presented from a qualitative perspective, and should only be considered as a conservative bound on the effects of physical capture.

It is important to highlight that the results shown up to this point had the capture effect disabled. This was achieved by artificially increasing the packet error probability to 1 when a collision was detected. Such simple modification is not only the result of a well-designed software, but also of its opensource nature.

B. Effects of the Physical Capture on the Upper Layers

1) MAC layer: We will now discuss the effects of this physical capture on the layers above, in particular the MAC and transport layers. Let us first consider DCF only (and thus use UDP at the transport layer). As we mentioned before, the capture phenomena will result in unfairness among stations with different propagation conditions towards the AP. Indeed, in the event of a collision, the frame that is received with the highest power may still be correctly decoded, the probability of which depends on the timing between frames, the modulation in use, and the received power of all the frames involved.

Let us first consider a scenario with only two stations, whose distance to the AP is the same. One of these stations (which we will note as n_1) transmits at a fixed power (10dBm) while the other one (n_2) transmits at different power levels, simulating a difference in the propagation conditions towards the AP. We show the results for each of these levels in Fig. 19. More in particular, both stations are transmitting up-link UDP traffic in saturation conditions using the 2Mbps modulation rate, and we present the results corresponding to both access methods. Although somewhat simple, the example will nevertheless illustrate some interesting aspects of the effects of capture on DCF.

Figure 19a shows the ACK-timeout probability for each station. That is to say, the ratio between the number of failed transmissions and the total transmission attempts. Naturally, when the power of both stations is the same, their results are also the same. However, note the important difference between the access methods. Whereas under the basic access both stations obtain an ACK-timeout probability of roughly 0.05, under RTS/CTS they obtain a very small 0.015. This happens despite the fact that both methods use the same modulation rate in this scenario. The reason behind the difference actually lies on the packet size. Naturally, the longer a packet, the bigger the probability of it being dropped due to interference, and in this case in particular, of not capturing the channel. All in all, the capture phenomena, and its associated unfairness, is exacerbated by the RTS/CTS mechanism due to the combined effects of a (generally) lower modulation rate and smaller packets involved in the collisions.

As we mentioned before, ns-3 may only simulate a bound on the capture effect. Indeed, as the transmit power of n_2 increases, the ACK-timeout probability does not goes to zero as expected [58], but converges to a fixed value. However, we may still observe important differences in throughput between both stations (see Fig. 19b). In addition to most collisions being resolved to its favour, and precisely due to this fact, n_2 will operate most of the time with a contention window equal to W_{\min} , resulting in more frequent transmission attempts than n_1 . We have thus two levels of unfairness at two different layers: the physical layer capture which arbitrates most collisions favourably to n_2 , and a resulting bigger contention window in the MAC layer of n_1 which diminishes its transmission attempt rate.

To finish the analysis of the MAC layer we will discuss a somewhat more realistic case-scenario. Let us consider ten stations randomly distributed inside a circle of radius 50m around the AP, where each of them saturates the channel with UDP traffic, all transmitting at the same power and with the same modulation rate (in this case 11Mbps). Figure 20 shows the individual throughput obtained by each station as a function of their mean received power at the AP for both access methods. Please note that the results obtained by each station in each of the 10 simulation runs are displayed.

It is important to highlight that we have verified that a station alone obtains the same throughput independently of its distance to the AP (the verification was made up to 100m, the



Fig. 19: ACK-timeout probability and individual throughput when two stations are at the same distance of the AP, one transmitting with a power of 10dBm and the other with the power indicated at the x-axis. Traffic is saturating UDP up-link, modulation rate is 2Mbps, and results for both access methods are shown.



Fig. 20: Individual throughput as a function of the received power. Ten stations are randomly deployed in a circle of radius 50m around the AP and are sending uplink UDP traffic in saturation conditions. The data rate is fixed at 11Mbps, and transmission power is also fixed. Results correspond to ten different simulation runs.

maximum distance between any two stations in this case). This means that the difference in the station's individual throughput in this case-scenario is due only to channel capture (and in particular, not due to hidden stations).

Naturally, as the received power increases, so does the resulting throughput, as those stations capture the channel more easily in a collision. Furthermore, the difference between the minimum and the maximum throughput can be very important. This is specially so for the RTS/CTS access, where some few lucky stations may receive more than three times throughput than the (not so few) unlucky ones. Finally, note that for the basic access, most of the stations obtain a throughput of roughly 500kbps, similar to the results we showed in the case of no capture (cf. Fig. 4). This similarity is because at 11Mbps, the difference in power necessary to capture the channel is very important and capture events are fewer. On the other hand, in the case of RTS/CTS access, the results are very different from the 400kbps we obtained before (again, see Fig. 4). This, as we mentioned before, is due to the tendency that the (small and modulated at lower rates) RTS packets have to capture the channel.



Fig. 21: Individual throughput as a function of the received power. Ten stations are randomly deployed in a circle of radius 50m around the AP and are all downloading an infinitely long file from a server behind the AP. The transmission power is fixed. Results correspond to ten different simulation runs.

2) Transport layer (TCP): We will now discuss the effects of the capture phenomena over TCP. As we showed in Sec. IV, if all stations are using TCP as their transport protocol (either downlink or uplink) the number of "active" stations (i.e. those with a packet in their buffer waiting to be sent) at any given time is generally very small. This in turn results in a very small collision probability. It is then to be expected that under this scenario the capture phenomena is not influential, since it simply does not manifest itself.

Figure 21 verifies the above analysis. It shows the individual throughput of each station corresponding to the same scenario as Fig. 20; i.e. 10 stations randomly deployed in a circle of radius 50m around the AP. Note how the results are almost the same as when no capture was allowed to take place (cf. Fig. 10). This is a good example that illustrates that extrapolating results from different case-scenarios is not necessarily correct. In this case, an unforeseen behaviour of the upper layers negates a potentially harmful effect of the lower ones, without necessarily violating the layering principle.

Naturally, several models to predict the resulting throughput considering the capture effects have been proposed in the literature, specially for the UDP saturated case. We have already cited [56], that corrects the success and failure probabilities of Bianchi's model to include the possibility of a capture. Other works exists in this same line with different levels of complexity (see for instance [63], [64], [65]). However, most of them assume that stations may control their transmission power so that the received power at the AP is the same for all of them. Then, based on assumptions on fading calculate the probability of a given station to capture the channel. Power control is not a common feature in standard 802.11 cards, and even if implemented, the transmitting power may be chosen from a discrete set of values. To the best of our knowledge, the



Fig. 22: A scenario where two stations are hidden from each other: although communications with the AP work properly, an opaque wall between them results in neither being able to sense the other.

only work that explicitly considers different receiving powers and is interested in estimating the unfairness among stations is [66]. However, this estimation is provided indirectly by what they call the average inter-success backoff delay (i.e. the mean number of slots between two successful transmissions). Moreover, the effects of channel capture are blurred by the presence of multiple rates (and an unspecified rate adaptation algorithm). All in all, a model to analyze this case-scenario, that for instance considers a realistic capture model such as the one proposed in [60], is still a challenging open problem.

To finish this section, let us remark that there exist some works that strive at minimizing this potential unfairness. For instance, the authors of [67] study the usefulness of several MAC and PHY layer parameters, such as the W_{\min} , to mitigate the difference in individual throughput. A more recent article has extended this work and proposes a mechanism to adapt these parameters on the fly based on some local measurements [68].

VI. THE HIDDEN STATION PROBLEM

We will now discuss one of the most "popular" problems in CSMA: the hidden station. It should be clear that CSMA relies heavily on the ability of each station to sense the transmission of every other station. However, it is easy to imagine a situation where this assumption does not hold. Figure 22 illustrates such case-scenario. Two stations are transmitting to a single AP, but a wall between them, opaque to the RF signals, results in neither stations sensing the other. Another case scenario is one where stations are disposed around the AP, all being able to communicate with the AP, but those stations far away from each other do not sense each other. Such situation may arise for instance in a residence with a single AP, where the signal of those stations far away from the AP has to traverse several walls in order to reach the stations at the other end of the building. Let us now define more formally this case-scenario. Several stations are disposed around the AP, and the signal from and to it is excellent (i.e. no losses due to errors in the channel are experienced in the exchanges with the AP). We will say that station j is hidden from station i when a transmission from station j reaches station i with a power such that the former's CCA will respond with idle. In the symmetric scenario we were considering so far (i.e. all stations have equivalent equipment and apply the same algorithms), if station j is hidden from station i, then the reciprocal is also true (i.e. station i is hidden from station j).

A. Performance Penalties

The most important consequence for a station that is hidden from other stations is that its transmission is simply ignored by the latter. Up to this point, collisions could only occur when the backoff counter of more than one station reached zero at the same time-slot. However, under this scenario, collisions may occur in the middle of a transmission, because the tagged station's frame took so long to be transmitted, that the backoff counter of one of its hidden stations reached zero before it could end the transmission. In fact, if all stations are hidden from each other, then CSMA behaves similarly to a simple Aloha [69], where stations transmit without sensing the channel first, and if the packet is not acknowledgement, they re-transmit it a random time later.

Figure 23 shows the results corresponding to the simplest case-scenario: only two stations are sending traffic to the AP. Different modulation rates and access methods are considered, as well as whether the stations are hidden from each other or not. It is important to highlight that all simulations we show from now on allow channel capture. We believe it contradictory to consider a more complete PHY model that includes the possibility of hidden stations but ignores channel capture, which as we discussed before requires much smaller power differences in order to manifest itself.

The first thing that should be noted in Fig. 23a is that naturally the system throughput is always smaller when the stations are hidden from each other than otherwise. However, note that the impact is much smaller with RTS/CTS access than with its basic counterpart. This is precisely the goal of this access method. The basic idea is to decrease the duration over which a hidden station may inadvertently interfere the tagged station's ongoing transmission (the so-called *vulnerable period* [5]). The RTS packet is small in size, and the CTS packet signals all stations in the AP's range (in this case, all stations) of the impending transmission, as well as its duration. The (generally) longer packet transmission may then occur without collisions.

The second aspect that deserves attention is the fact that as the modulation rate increases, the saturation throughput of the RTS/CTS access method becomes comparable to the one obtained by the basic one. Two elements act in conjunction here. Firstly, as the modulation rate increases, the vulnerable period for the basic access decreases (in this case the airtime of a frame). This is very clear in Fig. 23b, where we show the ratio between the number of transmission and ACKtimeout events. Note that as the modulation rate increases, this ratio decreases. This is despite the fact that capture phenomena becomes rarer as the rate increases. The second aspect in play here is the fact that the overhead produced by the RTS/CTS handshake becomes relatively more important as the modulation rate increases. This overhead is such that, although the ACK-timeout probability is almost twice for the basic access than for RTS/CTS, the throughput in both cases is very similar.

We have mentioned the vulnerable period as a very important factor in this case-scenario, and how the RTS/CTS access mechanism strives at minimizing it. Figure 24 further illustrates this point. In it, we show the results obtained in the same case-scenario as before (i.e. two stations hidden from each other sending up-link UDP traffic in saturation conditions), although in this case we vary the payload size used by both stations.

As we discussed in Sec. III-C2, in Bianchi's context, where stations were not hidden from each other and frames were not lost except for collisions, throughput was a linear function of the payload size used by all stations, and the collision probability was independent of the payload size. This is precisely the case for the RTS/CTS access method even in this case, because its vulnerable period is independent of the payload size. On the other hand, for the basic access the vulnerable period is almost proportional (albeit the headers, which are fixed) to the payload size. As shown in Fig. 24b, this results in an ACK-timeout probability that increases with the payload size, and a saturation throughput that is not strictly increasing with the payload size.

To finish this subsection, we will consider a case-scenario where some stations are hidden from each other, whereas some other are visible to all. As we mentioned before, such situation may arise in an indoor scenario, with very poor propagation conditions, where those stations far away from the AP may be hidden from each other, but those close to the AP can "listen" to all transmissions (and are thus "listened" by all other stations).

In particular, and similarly to what we did in Fig. 20, we will randomly position 10 stations in a circle of radius 50 meters with the AP at its center. However, this time the path-loss function is such that stations that are more than 60 meters apart are hidden from each other (although we have verified that each station is capable of perfectly communicating with the AP if alone, independently of its position in the circle).³ The results for both access methods (with a data modulation fixed at 11Mbps) are shown in Fig. 25. Each point corresponds to the throughput obtained by each station in each of the 10 simulation runs.

Results pertaining to the basic access (Fig. 25a) are in sharp contrast with the ones we obtained when no hidden stations

³More in particular, we have used the *ThreeLogDistancePropagationLoss-Model* model of ns-3, with parameters *Distance0*, *Distance1*, *ReferenceLoss*, *Exponent0* and *Exponent1* equal to 1, d, 2, 4.9 and α respectively (where d varied linearly from 50 to 100 and α from 24 to 13).



Fig. 23: Results corresponding to two stations sending UDP traffic towards the AP in saturation conditions for different modulation rates. The case-scenario where the stations are hidden from each other and that they "see" each other are included.



Fig. 24: Results corresponding to two stations hidden from each other sending UDP traffic towards the AP in saturation conditions for different payload sizes. The modulation rate is fixed at 11Mbps.

were allowed (Fig. 20a). Whereas before we obtained that most stations obtained roughly 500kbps, the level of unfairness is now very important. Stations near the AP may lose a packet only if their backoff counter reaches zero simultaneously with another station. Even in such case, the AP may respond with an ACK if the tagged station captures the channel, an event very likely to happen if for instance the competing station is one of the farthest away from the AP. On the other hand, those stations far away from the AP may further lose packets because a hidden station interrupts its transmission. This asymmetry in packet losses results in some stations receiving almost no throughput, whereas the other luckier stations obtain a throughput of roughly 1200kbps. The unfairness is still very important even when using the RTS/CTS access mechanism (Fig. 25b). In addition to the capture phenomena we observed and discussed before (see Fig. 20b), the presence of stations hidden from those farthest away from the AP results in these stations obtaining a throughput close to zero. Those close to the AP obtain a throughput of roughly 8000kbps, smaller than the one obtained by the basic access mostly due to its overhead.

The simulation results we just presented indicate that the RTS/CTS mechanism is far from alleviating the performance penalties generated by hidden stations. Let us briefly consider the TCP case now. The same results as before, corresponding now to this transport protocol, both for downlink as well



Fig. 25: Individual throughput as a function of the received power. Ten stations are randomly deployed in a circle of radius 50m around the AP and are sending uplink UDP traffic in saturation conditions. The data rate is fixed at 11Mbps, and stations that are more than 60 meters apart are hidden from each other. Results correspond to ten different simulation runs.



Fig. 27: Ten stations are randomly deployed in a circle of radius 50m around the AP. Individual throughput as a function of the carrier sense for all case-scenarios considered so far. Results correspond to the median and the 0.95 and 0.05 quantiles of the individual throughput obtained by all stations.

as uplink traffic, are shown in Fig. 26 (since similar results are obtained for the RTS/CTS mechanism, we omit them). Moreover, we compare all the case-scenarios considered so far in Fig. 27, which additionally shows the effect of the carrier sense (i.e. at which distance stations are hidden from each other), which was fixed at 60 meters up to now. More in particular, the graph displays the median and the 0.95 and 0.05 quantiles of the individual throughput obtained by all stations.

Once again, the results for TCP are remarkably good, specially for downlink traffic where no unfairness is visible

(Fig. 26a and the squares in Fig. 27). The reasons behind this excellent performance are the same as before: TCP operates in such way, that the number of stations actually competing for the channel at any given moment are very few, and thus few collisions occur. When the packets generated by the stations are small TCP acknowledgements (and thus the vulnerable period is small), collisions between hidden stations are rare. However, when the stations send bigger data packets, the possibility of collisions increases, and we obtain that some few stations are starved (Fig. 26b and the diamonds in Fig. 27). However, this is true only in the case where most stations are hidden from each other. With a carrier sense of just 70 meters, this negative effects are much less visible. Finally, note how no real benefits are obtained from the RTS/CTS scheme in the complete range of carrier sense considered.

B. Previous Studies

As we mentioned before, this case-scenario has been identified very early on as a very problematic one. In particular, Tobagi and Kleinrock analyzed it in the context of their then newly proposed CSMA [70]. To alleviate it, they propose the so-called *Busy Tone Multiple-Access* (BTMA) mechanism, where whenever a station senses an ongoing transmission, it sends a busy-tone over an independent signalling channel.

However, although unanimously recognized as grave, BTMA was not widely adopted as the solution to the hidden station problem. The reason being that permanently dedicating a (however small) portion of the total channel to signalling was considered by most as too big a waste (although further refinements of BTMA were proposed, e.g. [71]). Instead, "inband" signalling was preferred as a solution to this problem. In particular, the RTS/CTS mechanism was proposed in [72], and further extended in [73].



Fig. 26: Individual throughput as a function of the received power. Ten stations are randomly deployed in a circle of radius 50m around the AP and are sending downlink and uplink TCP traffic. The data rate is fixed at 11Mbps, and stations that are more than 60 meters apart are hidden from each other. Results correspond to ten different simulation runs.

We have already discussed some short-comings of the RTS/CTS mechanism even in our WLAN scenario, and more arise in the context of ad-hoc networks (specially mobile ones, like vehicular networks). It should come as no surprise then that several works have striven at further improving it, or replacing it altogether (see for instance [74], [75], [76], [77] just to name a few). Instead of modifying the MAC layer, some works have studied methods that allow the receiver to decode colliding frames when they arrive with a relatively significant time offsets as in this case (e.g. [78], [79]).

An approach that has received little attention in the literature is to force the system to operate at a regime where the effective number of competing stations is low. This way, collisions would be few, and the negative effects of hidden stations (as well as all problems stemming from collisions, like capture and wasted resources) would be minimized. We have shown how TCP indirectly drives the system to such behaviour. Further studies in this direction would be interesting.

We now turn our attention to modelization efforts of this phenomena. In addition to presenting the problem, [70] also analytically studies the performance of CSMA in the presence of hidden terminals. To this end, they propose a graph, where an edge between nodes i and j indicates "visibility" between them, and its absence that they are hidden from each other. Their analysis, and all further extensions (e.g. [80], [81]), are however not directly applicable to our case scenario, since it pertains to a general CSMA algorithm and does not consider any of the particularities of DCF.

Anyway, a similar approach was taken by other authors when studying 802.11 in this scenario. For instance, the authors of [82], [83] also consider a model where stations may "see" each other or not if they are closer than a certain *carrier range* (which may be viewed as a way of constructing the graph we mentioned before). However, they further consider two other ranges: the transmission and interference range. The former is the maximum distance at which two stations can communicate without errors in the absence of other concurrent transmissions. The latter is the distance within which an interfering transmitting station can jam a receiver.

Although the above model, based on one-on-one relationships between nodes, simplifies the analysis and has proved somewhat popular (see for instance [84] for a work based on these thresholds which studies WLANs in particular), it is not the best way to proceed. Actually, depending on the CCA mode in use, a node may judge the channel as busy based on the *sum* of all detected signals, and not on each individual station signal. Moreover, the interference range lacks any real significance, since again the interference is the sum of all ongoing transmissions other than the one of interest. Finally, let us highlight that, as analyzed in [85], this oversimplified model may lead to contradictory results when compared to more realistic ones.

Naturally, extensions to Bianchi's model for this casescenario have been proposed (for instance [86], [87], [88]). However, as discussed in [89], the variable length timeslot, central in Bianchi's analysis, is ambiguous in this casescenario, where coordination is lost between stations hidden from each other. Moreover, the authors of [90] show that the collision probability is not independent of the backoff stage in this case. This invalidates the formula resulting from the markov chain, and thus all analysis based on it. The authors of [90] consider instead a simple case where only two stations are present in the WLAN, construct an alternative markov chain, and derive the total throughput. They then generalize the model to consider an arbitrary number of stations, which allows them to calculate the station's individual throughput. However, the decision of whether one station senses the channel as idle or not in their model is based on one-on-one relations between them.

Let us finally highlight that all these models ignore the effects of physical capture. For instance, in the case of RTS/CTS access, with RTS frames modulated at 2Mbps, a difference of only 2dBm would result in one of the colliding frames capturing the channel. To consider that the path loss is such that two stations, which are in the range of the AP, suffer losses of tens of dB between them, but that their signal reaches the AP with almost exactly the same power, strikes us as somewhat contradictory. To the best of our knowledge, a model that considers both hidden stations and the capture effect does not exist.

VII. HIGH THROUGHPUT WI-FI: 802.11N

In this section we discuss 802.11n in more detail. We begin by highlighting research that the main new features included in the standard has sparked. We then argue that the results we had presented so far may be adapted to the new standard and how to do this. The last subsection briefly presents the recent amendments that would permit Gigabit WLAN.

A. New features: new research opportunities

1) 40MHz bandwidth channel: The basic idea behind doubling the channel's bandwidth is that, in principle, it results in twice the rate (see [91] for a more precise estimation of this gain). However, this seemingly simple method has its associated costs. For instance, in the 2.4GHz band, where there are only three non-overlapping 20MHz channels available (compare this to the twelve 40MHz non-overlapping channels available in the 5.8GHz band), interference from other WLANs is not to be neglected and may have a very important impact on the obtained throughput [92], [93].

Moreover, receivers with twice the bandwidth also capture twice the noise, and since the total transmitted power is still the same, SNR is reduced by 3dB. This may be specially important in the 2.4GHz band, which is shared with cordless phones, microwave ovens or bluetooth.

It should be clear by now that 802.11n should be used in the 5.8GHz band, and that the decision of whether to use a 40MHz or 20MHz channel has no simple answer. The interested reader should consult [94], [95] for an in-depth study on this question.

2) *MIMO:* As for instance demonstrated in [96], MIMO technology introduced in 802.11n is beneficial under all circumstances when compared to traditional SISO (Single Input Single Output) used in legacy a/b/g devices. In 802.11n up to four antennas may be used both at the transmitter and the receiver, resulting in up to four spatial streams. However they may be used to transmit several different streams (i.e. spatial multiplexing) or several redundant streams (i.e. spatial diversity). Naturally, although the former is capable of obtaining bigger data rates, the latter is more robust in face of interference, and a tradeoff exists between both benefits [97]. Moreover, spatial diversity requires a propagation environment where the samples obtained in the different space dimensions

are actually independent. This means that, even with an excellent SINR, the biggest rates may not work [93].

It should be clear that choosing the appropriate data rate (out of the 128 available) is no easy task. This is specially so when one realizes that the same rate (or very similar ones) may be obtained by more than one combination of MCS (Modulation and Coding Scheme) and number of spatial streams. For instance, under the 800ns GI and 40MHz channel, 108Mbps may be obtained by either using a single spatial stream in combination with 64-QAM and a coding rate of 2/3, or two different spatial streams with 16-QAM and a coding rate of 1/2. Naturally, in the former case transmission will operate under the spatial diversity regime if both receiver and transmitter have more than one antenna.

The observation above means that smaller rates do not necessarily result in more robust transmission schemes, breaking the assumption that underlies most rate-selection algorithms (such as ARF, which we discussed in Sec. III-B). As presented for instance in [98], [99], a correctly designed algorithm should explore both dimensions: MCS and number of spatial streams.

3) Frame Aggregation: Although a simple idea, aggregation poses several fundamental questions that should be answered. For instance, the standard assumes that each subframe in either an A-MPDU or A-MSDU should correspond to a higher layer packet (e.g. TCP). Moreover a simple scheme could be to form an A-MPDU or A-MSDU when either a certain size is reached or a timer expires. But, are these decisions optimal in any sense?

The authors of [100] present a very interesting study on frame aggregation and provide answers to many of the questions we asked before. In a nutshell, they conclude that A-MPDU is a very good aggregation strategy, mainly due to its selective retransmission capability (an unsurprising result, discussed for instance in [101]). However, they also show that two features should be implemented in order to fully exploit A-MPDU: a carefully designed fragmentation strategy and the so-called zero-waiting mechanism (i.e. accumulate packets only until the MAC obtains a transmission opportunity).

B. Extrapolating results to 802.11n

We now briefly discuss how virtually all we have presented in previous sections, mostly pertaining to 802.11 a/b/g, is still valid for 802.11n. To begin with, it should be clear that Bianchi's model is perfectly applicable to the new standard. Relatively small modifications are necessary to include the new MAC layer features (see for instance [100] for an example).

The key to this simple adaptation of Bianchi's model is that DCF's basic operation has not changed. This in turn means that some of its "negativities" are still present. Notably, the performance anomaly we discussed in Sec. III-C1, stemming from the access opportunities fairness enforced by DCF, still manifests itself [102]. However, the new features included in the standard, particularly the aggregation, may be used to

control this performance degradation to some degree (see for instance [92], [103]).

Moreover, TCP still performs very well, even more so when aggregation is enabled both at the AP and the stations [103]. To illustrate why aggregation in both senses is necessary, let us consider the case of a station downloading a file. If aggregation is allowed only on the AP, the station will only be able to answer with a single acknowledgement per access opportunity.

Naturally, both the physical layer capture (cf. Sec. V) and the hidden station problem (cf. Sec. VI) are still present. Regarding the latter, the large frame sizes may result in larger vulnerable periods (and thus more collisions) if the fastest modulation and coding schemes are not used. However, if the colliding frame arrives to the receiver in the middle of an A-MPDU transmission, most probably part of the A-MPDU will be correctly received and a BlockAck acknowledging the first sub-frames will be sent. This may be further exploited to detect hidden stations as discussed in [104].

However, and to the best of our knowledge, a thorough study of the benefits of the new features of 802.11n in the context of hidden stations has not been performed. Similarly, the effects of MIMO on the capture phenomena have not been reported so far. Although at first glance one would tend to believe that such high rates do not permit capture to happen, the usage of robust coding schemes (like STBC) may prove this wrong.

C. Further rate extensions: Gigabit WLAN and beyond

With the 100 Mbps MAC level throughput objective accomplished, it was only natural to target the Gigabit WLAN. In 2008 two Task Groups were formed to pursuit such accomplishment: 802.11ac and 802.11ad, corresponding to the bands below 6GHz and the 60GHz one respectively.

802.11ac is still not standardized, being 5.0 its latest draft and having very recently been forwarded to sponsor ballot. Overall, the new amendment will further transit the path of 802.11n, in the sense that the channel bandwidth may now be as large as 80 or 160MHz, 256 QAM is now supported, as well as 8×8 MIMO allowing up to 8 spatial streams (resulting in a maximum throughput of 6.9Gbps).Moreover, A-MSDU and A-MPDU are now allowed to be even larger than before [105].

However, the main novelty in this amendment is the use of the so-called Downlink Multi-User MIMO (DL MU-MIMO). In a nutshell, this is a technique by which different spatial streams may be sent to different receivers. If up to eight antennas may be present in the AP, this feature is particularly appealing for communicating with devices that for different reasons (e.g. cost or space) may not include so many antennas. In such case, DL MU-MIMO allows the AP to multiplex both spatial streams. The interested reader may consult [106] for an introduction on the subject, or [107] for a study on the benefits that this scheme brings to CSMA.

On the other hand, the standardization process of 802.11ad was finished in late 2012 [108]. The basic objective of this amendment was to take advantage of the relatively underused and unlicensed 60GHz band. In addition to presenting little interference, it has more bandwidth available than the 2.4GHZ

or 5.8GHz band, and thus permits much wider channels. However, such high frequency present the disadvantage of resulting in a much bigger path-loss (due to smaller antenna effective area, atmospheric absorption and higher loss in building materials), thus requiring highly directive antennas. Such characteristics have made this band the most appealing candidate for wireless replacement of wired digital interfaces (e.g. HDMI). An example of industrial efforts in this sense is the WirelessHD specification [109], whose first version was published in 2008.

The interested reader should consult the recent survey [110] for further references regarding these two new standards. Finally, let us highlight that this order-of-magnitude increase in the raw data rate has re-sparked interest in more efficient MAC algorithms. Examples of recent articles in this sense include [111], [112], [113], [114].

VIII. A MAC WITH QUALITY OF SERVICE: 802.11E

In this section we briefly discuss 802.11e (i.e. EDCA), its performance and the models present in the literature. Naturally, this scheme grasped much research attention very early on. For instance, one of the first articles to study EDCA's performance was published as early as 2001 [115]. The most natural question is how each of the "knobs" we mentioned in Sec. II-B influences performance. The effects of a big TXOP are relatively straightforward to understand, at least in the saturated case-scenario. A very interesting discussion of the effect on prioritization and overall performance of the other two parameters is presented in [116]. The article demonstrates the superiority of AIFS differentiation due to, among other reasons, the important number of time-slots wasted in collisions when several ACs with a small $W_{\rm min}$ contend for the channel.

Regarding the models proposed for this MAC mechanism, extensions to Bianchi's model were also rapidly proposed. See for instance [117] (which reaches conclusions similar to [116] regarding the usefulness of AIFS and $W_{\rm min}$ differentiation mechanisms) or the extensions of [14] to this case presented in [118].

Although the first experiments using the new standard recognized its ability to differentiate traffic, they also identified the need of an access control scheme if QoS was to be guaranteed [119]. Consequently, several such schemes were proposed in the literature, a overview of which may be consulted in [120], [121].

There is another important aspect of EDCA that we have not mentioned so far, which is the specific values of the AIFSN, TXOP, W_{\min} and W_{\max} for each AC. The standard only dictates the existence of four different ACs, and suggest default values for each of them. However, they may be changed on the fly, and this ability has attracted many researchers' interest. A recent discussion on this topic is provided in [122], [123].

However, and despite all the research effort invested in it, as we mentioned in Sec. II-B, the implementation and usage of EDCA has been very limited. This situation responds mostly to three facts. Firstly, the Wi-Fi Alliance provides certification testing of its own *Wi-Fi Multimedia* (WMM) specification, based on a subset of functions from the 802.11e draft 6. In this sense, it is interesting to highlight that, since the final version of 802.11e and WMM differ, a Study Group was formed in 2007 so as to adapt 802.11e to WMM, which was shortly after dissolved [124].

The second reason behind 802.11e's (or even WMM's) limited usage is the decision of which frames are assigned to each AC. This is up to the specific application, which for instance may mark each packet with the pertinent IEEE 802.1D priority tag [125], and the conversion to the corresponding AC is defined in the 802.11e standard (or automatically performed by WMM-enabled APs). However, such marking is rarely performed by real multimedia applications or devices. For instance, the interested reader may consult [126] for an account of the frustrations of a user trying to exploit WMM on a domestic network.

The third and last reason stems from the fact that even certified equipment is not guaranteed to be totally interoperable [127]. In such case, stations generating traffic of the same AC may be unintentionally prioritized over each other. For instance, the authors of [128] report on the usage of 802.11e cards constructed by different vendors. Results indicate that prioritization takes place (or not) depending on the card that sends the highest priority traffic. This is in addition to the implicit prioritization provided by the PHY layer when the capture phenomena manifests.

IX. 802.11 ON THE WILD: SOME MEASUREMENTS STUDIES ON OPERATIVE WLANS

We will now briefly overview one very important aspect of WLAN's performance: measuring operative and uncontrolled networks. One of the first works to report such measurement campaign was [129], which presented the measurement methodology and analysis of a twelve-week trace of a WLAN spanning a complete building (in particular, the Gates Computer Science Building of the Stanford University, comprising 12 APs and 74 wireless users). The authors were particularly interested in user behaviour and traffic characteristics. That is to say, they studied aspects like the number of active users, sessions' length, users' mobility, traffic distribution among APs, or traffic per user and per application. All these aspects are important and should be used as an input when planning the deployment of a new WLAN.

Regarding the measurement methodology, the authors of [129] used a combination of tcpdump, SNMP queries and the WLAN's authentication system log. These same authors used a similar approach in [130], where they analyzed network usage and user's mobility in a metropolitan network by means of analyzing the logs of the poletops (equivalent to the APs in Metricom's Ricochet multi-hop network).

This kind of "wired-side" analysis of the WLAN has been carried out by several works in different contexts: conferences [131], corporate buildings [132], university campuses [133], [134], [135], [136], metropolitan networks [137], [138] or even restaurants [139]. The interested reader should consult each of them for they sometimes reach contradictory conclusions, which is natural since these studies span almost a decade. A study of the evolution of wireless data traffic based on some of these studies may be found in [140].

Although interesting in its own right, the measuring method cited above is not useful if we are interested in the wireless medium. For instance, if a station sends a packet while another sender has already captured the AP, this event will not be logged except by that particular station. Differently to a testbed, in an uncontrolled network it is not possible to perform measurements at every station. In this context, the idea of using several *wireless monitors* disposed around the WLAN emerged [141], [142], [143].

There are two main difficulties that one must overcome when using this "wireless-side" analysis. The first one, is that the vision of each individual monitor is intrinsically incomplete and different from each other, as we illustrated above. The second one, is that we would ideally want a single view of the wireless medium (i.e. a time-line with the events and their duration), so a method to merge all these different views has to be devised [144]. For instance, the authors of [141], [142] rely on the AP's beacon packets in order to align all traces, whereas those of [143] study the possibility of inferring missing packets by considering the possible frame exchanges as established by the standard (a method later refined in [145]).

Several works have been published since that propose different measurement techniques or software. For instance, Jigsaw [146] is a wireless-side measurement software based on [141], [142] that stresses scalability. More recently, WiMed [147] and PIE [148], both designed to measure interference, have been made available. Several other tools exists, some of which are classified and briefly discussed in [149]. Practically speaking, the most natural use for this kind of methods is security. There are several security breaches which cannot be detected at the wired-side of the network. Examples of such usage are [150], [151], [152], [153].

Quiet surprisingly, the articles that focus on the actual medium usage are relatively few. For instance, the authors of [154] use commercial wireless sniffers to study a relatively big WLAN (97 APs and 6775 users). However, their analysis of the wireless side of the network is very small. An exception are [155], [156], which provide significant insight into the operation of DCF in real deployments. Among the conclusions they reach are some of which we already discussed in this article: the usage of lower rates by some stations is detrimental to the network as a whole, the tendency of most rate control algorithms to confuse congestion with poor propagation conditions exacerbates this phenomena and should then be avoided, and the use of RTS/CTS only increases the unfairness among stations.

Finally, let us remark that conducting this kind of measurements is not an easy task at all, even in a controlled environment like a test-bed. To illustrate these difficulties, let us cite the experiments reported in [157], where the author obtain a very poor performance of 802.11g in an outdoor scenario. This was later explained to be due to a buggy behaviour of the driver in use [158], [159].

X. OTHER CASE-SCENARIOS

Although our survey is limited to the single AP WLAN case-scenario, there are other equally important scenarios the reader should be aware of, and as such we now very briefly describe.

The most clear extension to our scenario is the multi-AP case, where interference and handover play a major role in the performance of the access network as a whole. Moreover, and regarding the former, it is not rare at all for terminals to be equipped with several radio access technologies (e.g. most smartphones include both cellular and WiFi capabilities). Consequently, there has been much research and industry interest in enabling interworking and vertical (or inter-technology) handover. Examples of such efforts include the Interworking Wireless LAN (I-WLAN) [160], the Generic Access Network (GAN) [161] and the Media Independent Handover (MIH) [162] architectures. A very accessible introduction to this standard may be consulted in [163] and an interesting categorization of all these architectures is discussed in [164].

It is important to highlight that neither of these architectures include definitions regarding the decision by which a specific network is selected or when the handoff is effectively performed. Naturally, several such proposals exist in the literature, an overview of which may be found in [165], [166]. Moreover, the interested reader should consult [167] and references therein for analytical performance evaluation techniques regarding these handoff mechanisms. Finally, it is interesting to highlight the situation where not all these different networks are exploited by the same operator, and the challenging revenue sharing problem such situation poses [168].

To conclude this section, we wish to briefly comment on the 802.11af WG, formed on 2009, and in charge of adapting 802.11 to operate in the so-called TV White Spaces (TVWS). Since the FCC unlicensed the use of several bands corresponding to TV channels in 2008 [169] much effort has been put into designing communication systems that comply with the stringent requirements demanded by the regulation (e.g. less than -100dBm sensitivity in detecting primary users) [170].

As usual, the IEEE is no exception, creating the 802.11af SG and the 802.22 WG. In particular, the former is concentrated on wireless regional area networks (WRAN), and it has already published its standard [171]. Arguably the main difficulty for these systems is developing the cognitive [172] capabilities required for their operation. In this sense, the conjunction of cognitive networking (see for instance [173]) and the WLAN universe we just presented promises a very interesting and fruitful research area.

XI. CONCLUSIONS

Our survey provided what we consider a very complete view of the current understanding of WLAN performance. From it, we have learnt several lessons, some of which we would like to highlight. Firstly, and from a theoretical point of view, Bianchi's model was shown to be very versatile. For instance, although it was initially conceived with the completely symmetric casescenario in mind, we have discussed simple extensions that allow us to analyze different levels of asymmetry (e.g. different rates or packet sizes among stations). Fundamental in these extensions was the fact that all stations obtain the same access opportunities and collision probability. If this is not the case (which may arise for instance if the capture phenomena manifests) then the analysis is not so simple.

Moreover, and regarding the capture phenomena, we have shown that ns-3 does not correctly model the receiving process of an 802.11 card. The consequence of this oversimplistic model is that the quantitative results we obtain from the simulations may not be correct, and thus cannot be relied on. Although more detailed extensions exist (e.g. PhySim [174]), they are not widely used, and more importantly, do not scale very well. Obtaining any of the graphs we show in this paper may easily take a complete day if performed with PhySim on a standard PC.

From a practical perspective we have shown some detrimental side-effects of RTS/CTS and multi-rate algorithms. We believe that the optimal configuration for a WLAN is one in which basic access is used exclusively (which is generally the default behaviour of most cards) and that a single rate is allowed by the AP. This rate should be set to the minimum acceptable by the network manager, and it will operate as an access control, allowing stations to transmit only if their channel quality is good enough.

We have also illustrated the excellent performance attained by TCP in most cases, effectively avoiding many of the problematic situations we have presented. This was due to the protocol's ability of limiting collisions. A very interesting open problem is designing a scheme that emulates this desirable side-effect, without the potential unfairness issues of TCP when, for instance, the AP's buffer is not big enough to accommodate all connections.

Yet another open problem is the analysis of the WLAN's performance when the capture effect is considered. The lack of such model is surprising considering how easily this phenomena manifests itself. An even more complex task is to analyze the case-scenario where hidden stations and capture effect are jointly considered.

Finally, let us remark that at least one important aspect of WLAN's performance was mostly left out of this overview: security. In particular, a very interesting problem is investigating the possible performance benefits of a station that does not respect the standard, and how to detect it (see for instance [175]).

ACKNOWLEDGEMENTS

This work was partially funded by the Universidad de la República's CSIC i+d project "Algoritmos de Control de Acceso al Medio en Redes Inalámbricas".

REFERENCES

- "802.11-2012 IEEE Standard for Information technology– Telecommunications and information exchange between systems Local and metropolitan area networks–Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," 2012.
- [2] Samsung, "30 cu. ft. Side by Side Refrigerator and 8" LCD Digital Display with Apps." [Online]. Available: http://www.samsung.com/us/appliances/refrigerators/RSG309AARS/XAA
- [3] K. Huang, D. Malone, and K. Duffy, "The 802.11g 11 Mb/s Rate is More Robust than 6 Mb/s," Wireless Communications, IEEE Transactions on, vol. 10, no. 4, pp. 1015 –1020, april 2011.
- [4] "IEEE Standard for Information technology– Local and metropolitan area networks– Specific requirements– Part 11: Wireless LAN Medium Access Control (MAC)and Physical Layer (PHY) Specifications Amendment 5: Enhancements for Higher Throughput," *IEEE Std* 802.11n-2009 (Amendment to IEEE Std 802.11-2007 as amended by IEEE Std 802.11k-2008, IEEE Std 802.11r-2008, IEEE Std 802.11y-2008, and IEEE Std 802.11w-2009), 2009.
- [5] L. Kleinrock and F. Tobagi, "Packet Switching in Radio Channels: Part I-Carrier Sense Multiple-Access Modes and Their Throughput-Delay Characteristics," *Communications, IEEE Transactions on*, vol. 23, no. 12, pp. 1400 – 1416, dec 1975.
- [6] "ns-3." [Online]. Available: http://www.nsnam.org/
- [7] "IEEE Standard for Information Technology Telecommunications and Information Exchange Between Systems - Local and Metropolitan Area Networks - Specific Requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 8: Medium Access Control (MAC) Quality of Service Enhancements," *IEEE Std 802.11e-2005 (Amendment to IEEE Std 802.11, 1999 Edition (Reaff 2003))*, 2005.
- [8] F. Baccelli and S. Foss, "On the Saturation Rule for the Stability of Queues," *Journal of Applied Probability*, vol. 32, no. 2, pp. 494–507, 1995.
- [9] A. Kumar and D. Patil, "Stability and throughput analysis of unslotted CDMA-ALOHA with finite number of users and code sharing," *Telecommunication Systems*, vol. 8, pp. 257–275, 1997.
- [10] B. Crow, I. Widjaja, J. Kim, and P. Sakai, "Investigation of the IEEE 802.11 medium access control (MAC) sublayer functions," in *INFO-COM* '97. Sixteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE, vol. 1, apr 1997, pp. 126 –133 vol.1.
- [11] T.-S. Ho and K.-C. Chen, "Performance analysis of IEEE 802.11 CSMA/CA medium access control protocol," in *Personal, Indoor* and Mobile Radio Communications, 1996. PIMRC'96., Seventh IEEE International Symposium on, vol. 2, oct 1996, pp. 407 –411 vol.2.
- [12] F. Cali, M. Conti, and E. Gregori, "IEEE 802.11 wireless LAN: capacity analysis and protocol enhancement," in *INFOCOM '98. Seventeenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 1, mar-2 apr 1998, pp. 142 –149 vol.1.
- [13] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *Selected Areas in Communications, IEEE Journal on*, vol. 18, no. 3, pp. 535 –547, mar 2000.
- [14] A. Kumar, E. Altman, D. Miorandi, and M. Goyal, "New Insights From a Fixed-Point Analysis of Single Cell IEEE 802.11 WLANs," *Networking, IEEE/ACM Transactions on*, vol. 15, no. 3, pp. 588–601, june 2007.
- [15] K. Huang, K. Duffy, and D. Malone, "On the Validity of IEEE 802.11 MAC Modeling Hypotheses," *Networking, IEEE/ACM Transactions on*, vol. 18, no. 6, pp. 1935 –1948, dec. 2010.
- [16] G. Sharma, A. Ganesh, and P. Key, "Performance Analysis of Contention Based Medium Access Control Protocols," *Information Theory*, *IEEE Transactions on*, vol. 55, no. 4, pp. 1665 –1682, april 2009.
- [17] P. Robert, *Stochastic Networks and Queues*, ser. Stochastic Modelling and Applied Probability Series, Springer-Verlag, Ed., 2003.
- [18] R. W. Darling and J. R. Norris, "Differential equation approximations for Markov chains," *Probability Surveys*, vol. 5, pp. 37–79, 2008.
- [19] M. Benaim and J.-Y. L. Boudec, "A class of mean field interaction models for computer and communication systems," *Performance Evaluation*, vol. 65, no. 1112, pp. 823–838, 2008.

- [20] J.-w. Cho, J.-Y. Le Boudec, and Y. Jiang, "On the validity of the fixed point equation and decoupling assumption for analyzing the 802.11 MAC protocol," *SIGMETRICS Perform. Eval. Rev.*, vol. 38, no. 2, pp. 36–38, oct. 2010.
- [21] M. Bredel and M. Fidler, "Understanding Fairness and its Impact on Quality of Service in IEEE 802.11," in *INFOCOM 2009, IEEE*, april 2009, pp. 1098 –1106.
- [22] T. Lan, D. Kao, M. Chiang, and A. Sabharwal, "An Axiomatic Theory of Fairness in Network Resource Allocation," in *INFOCOM*, 2010 *Proceedings IEEE*, march 2010, pp. 1 –9.
- [23] E. Altman, K. Avrachenkov, and S. Ramanath, "Multiscale fairness and its application to resource allocation in wireless networks," *Computer Communications*, vol. 35, no. 7, pp. 820 – 828, 2012.
- [24] A. Kamerman and L. Monteban, "WaveLAN-II: a high-performance wireless LAN for the unlicensed band," *Bell Labs Technical Journal*, vol. 2, no. 3, pp. 118–133, 1997.
- [25] S. Choi, K. Park, and C.-K. Kim, "On the performance characteristics of WLANs: revisited," in *Proceedings of the 2005 ACM SIGMETRICS international conference on Measurement and modeling of computer systems*, ser. SIGMETRICS '05, 2005, pp. 97–108.
- [26] A. Vasan and A. U. Shankar, "An Empirical Characterization of Instantaneous Throughput in 802.11b WLANs," Tech. Rep. CS-TR-4389, UMIACS-TR-2002-69, Oct 2002. [Online]. Available: http://hdl.handle.net/1903/1216
- [27] L. Zhang, Y.-J. Cheng, and X. Zhou, "Rate avalanche: Effects on the performance of multi-rate 802.11 wireless networks," *Simulation Modelling Practice and Theory*, vol. 17, no. 3, pp. 487 – 503, 2009.
- [28] J. Kim, S. Kim, S. Choi, and D. Qiao, "CARA: Collision-Aware Rate Adaptation for IEEE 802.11 WLANs," in *INFOCOM 2006. 25th IEEE International Conference on Computer Communications. Proceedings*, april 2006, pp. 1–11.
- [29] J. Choi, K. Park, and C.-K. Kim, "Analysis of Cross-Layer Interaction in Multirate 802.11 WLANs," *Mobile Computing, IEEE Transactions* on, vol. 8, no. 5, pp. 682–693, may 2009.
- [30] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11b," in *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications. IEEE Societies*, vol. 2, march 2003, pp. 836 – 843 vol.2.
- [31] G. Tan and J. Guttag, "Time-based fairness improves performance in multi-rate WLANs," in *Proceedings of the annual conference on* USENIX Annual Technical Conference, ser. ATEC '04, 2004, pp. 23– 23.
- [32] D. Pong and T. Moors, "Fairness and capacity trade-off in IEEE 802.11 WLANs," in *Local Computer Networks*, 2004. 29th Annual IEEE International Conference on, nov. 2004, pp. 310 – 317.
- [33] A. Babu and L. Jacob, "Fairness Analysis of IEEE 802.11 Multirate Wireless LANs," *Vehicular Technology, IEEE Transactions on*, vol. 56, no. 5, pp. 3073 –3088, sept. 2007.
- [34] G. Xylomenos, G. Polyzos, P. Mahonen, and M. Saaranen, "TCP performance issues over wireless links," *Communications Magazine*, *IEEE*, vol. 39, no. 4, pp. 52 –58, apr 2001.
- [35] K. Pentikousis, "TCP in wired-cum-wireless environments," Communications Surveys Tutorials, IEEE, vol. 3, no. 4, pp. 2 –14, quarter 2000.
- [36] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose, "Modeling TCP Reno performance: a simple model and its empirical validation," *Networking*, *IEEE/ACM Transactions on*, vol. 8, no. 2, pp. 133–145, apr 2000.
- [37] S. Floyd, T. Henderson, and A. Gurtov, "The NewReno Modification to TCP's Fast Recovery Algorithm," RFC 3782 (Standard), Internet Engineering Task Force, april 2004. [Online]. Available: http://www.ietf.org/rfc/rfc3782.txt
- [38] R. Pries, D. Staehle, S. Oechsner, M. Menth, S. Menth, and P. Tran-Gia, "On the unfair channel access phenomenon in wireless LANs," in *Teletraffic Congress, 2009. ITC 21 2009. 21st International*, sept. 2009, pp. 1–8.
- [39] A. Kamerman and G. Aben, "Net throughput with IEEE 802.11 wireless LANs," in Wireless Communications and Networking Confernce, 2000. WCNC. 2000 IEEE, vol. 2, 2000, pp. 747 –752 vol.2.
- [40] R. Bruno, M. Conti, and E. Gregori, "Throughput Evaluation and Enhancement of TCP Clients in Wi-Fi Hot Spots," in *Wireless On-Demand Network Systems*, ser. Lecture Notes in Computer Science, R. Battiti, M. Conti, and R. Cigno, Eds. Springer Berlin / Heidelberg, 2004, vol. 2928, pp. 119–125.

- [41] A. Ferragut and F. Paganini, "Resource allocation over multirate wireless networks: A network utility maximization perspective," *Computer Networks*, vol. 55, no. 11, pp. 2658 – 2674, 2011.
- [42] F. Kelly, A. Maulloo, and D. Tan, "Rate control for communication networks: shadow prices, proportional fairness and stability," *Journal* of the Operational Research Society, vol. 49, no. 3, pp. 237–252, 1998.
- [43] E. Altman and T. Jimnez, "Novel Delayed ACK Techniques for Improving TCP Performance in Multihop Wireless Networks," in *Personal Wireless Communications*, ser. Lecture Notes in Computer Science, M. Conti, S. Giordano, E. Gregori, and S. Olariu, Eds. Springer Berlin / Heidelberg, 2003, vol. 2775, pp. 237–250.
- [44] J. Chen, M. Gerla, Y. Z. Lee, and M. Sanadidi, "TCP with delayed ack for wireless networks," *Ad Hoc Networks*, vol. 6, no. 7, pp. 1098 – 1116, 2008.
- [45] R. Braden, "Requirements for Internet Hosts Communication Layers," RFC 1122 (Standard), Internet Engineering Task Force, Oct. 1989, updated by RFCs 1349, 4379, 5884, 6093, 6298, 6633. [Online]. Available: http://www.ietf.org/rfc/rfc1122.txt
- [46] S. Pilosof, R. Ramjee, D. Raz, Y. Shavitt, and P. Sinha, "Understanding TCP fairness over wireless LAN," in *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications. IEEE Societies*, vol. 2, march 2003, pp. 863 – 872 vol.2.
- [47] D. Leith and P. Clifford, "Using the 802.11e EDCF to achieve TCP upload fairness over WLAN links," in *Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks, 2005. WIOPT 2005. Third International Symposium on*, april 2005, pp. 109 – 118.
- [48] D. Leith, P. Clifford, D. Malone, and A. Ng, "TCP fairness in. 802.11e WLANs," *Communications Letters, IEEE*, vol. 9, no. 11, pp. 964 – 966, nov. 2005.
- [49] N. Blefari-Melazzi, A. Detti, I. Habib, A. Ordine, and S. Salsano, "TCP Fairness Issues in IEEE 802.11 Networks: Problem Analysis and Solutions Based on Rate Control," *Wireless Communications, IEEE Transactions on*, vol. 6, no. 4, pp. 1346 –1355, april 2007.
- [50] H. Balakrishnan and V. Padmanabhan, "How network asymmetry affects TCP," *Communications Magazine, IEEE*, vol. 39, no. 4, pp. 60 –67, april 2001.
- [51] S. Shioda, H. Iijima, T. Nakamura, S. Sakata, Y. Hirano, and T. Murase, "ACK pushout to achieve TCP fairness under the existence of bandwidth asymmetry," in *Proceedings of the 5th ACM workshop on Performance monitoring and measurement of heterogeneous wireless* and wired networks, ser. PM2HW2N '10, 2010, pp. 39–47.
- [52] G. Bianchi and I. Tinnirello, "Remarks on IEEE 802.11 DCF performance analysis," *Communications Letters, IEEE*, vol. 9, no. 8, pp. 765 – 767, aug 2005.
- [53] R. Pickholtz, D. Schilling, and L. Milstein, "Theory of Spread-Spectrum Communications–A Tutorial," *Communications, IEEE Transactions on*, vol. 30, no. 5, pp. 855 – 884, may 1982.
- [54] J. H. Kim and J. K. Lee, "Capture effects of wireless CSMA/CA protocols in Rayleigh and shadow fading channels," *Vehicular Technology*, *IEEE Transactions on*, vol. 48, no. 4, pp. 1277 –1286, jul 1999.
- [55] J. Arnbak and W. van Blitterswijk, "Capacity of Slotted ALOHA in Rayleigh-Fading Channels," *Selected Areas in Communications, IEEE Journal on*, vol. 5, no. 2, pp. 261 – 269, feb 1987.
- [56] Z. Hadzi-Velkov and B. Spasenovski, "Capture effect in IEEE 802.11 basic service area under influence of Rayleigh fading and near/far effect," in *Personal, Indoor and Mobile Radio Communications, 2002. The 13th IEEE International Symposium on*, vol. 1, sept. 2002, pp. 172 – 176 vol.1.
- [57] C. Ware, J. Judge, J. Chicharo, and E. Dutkiewicz, "Unfairness and capture behaviour in 802.11 adhoc networks," in *Communications*, 2000. ICC 2000. 2000 IEEE International Conference on, vol. 1, 2000, pp. 159 –163 vol.1.
- [58] A. Kochut, A. Vasan, A. Shankar, and A. Agrawala, "Sniffing out the correct physical layer capture model in 802.11b," in *Network Proto*cols, 2004. ICNP 2004. Proceedings of the 12th IEEE International Conference on, oct. 2004, pp. 252 – 261.
- [59] J. Lee, W. Kim, S.-J. Lee, D. Jo, J. Ryu, T. Kwon, and Y. Choi, "An experimental study on the capture effect in 802.11a networks," in *Proceedings of the second ACM international workshop on Wireless* network testbeds, experimental evaluation and characterization, ser. WinTECH '07, 2007, pp. 19–26.
- [60] J. Lee, J. Ryu, S.-J. Lee, and T. T. Kwon, "Improved modeling of IEEE 802.11a PHY through fine-grained measurements," *Computer Networks*, vol. 54, no. 4, pp. 641 – 657, 2010.

- [61] "The WifiChannel and WifiPhy models in ns-3." [Online]. Available: https://www.nsnam.org/docs/release/3.15/models/singlehtml/index.html#thewifichannel-and-wifiphy-models
- [62] M. Lacage and T. R. Henderson, "Yet another network simulator," in Proceeding from the 2006 workshop on ns-2: the IP network simulator, ser. WNS2 '06, 2006.
- [63] F. Daneshgaran, M. Laddomada, F. Mesiti, M. Mondin, and M. Zanolo, "Saturation throughput analysis of IEEE 802.11 in the presence of non ideal transmission channel and capture effects," *Communications, IEEE Transactions on*, vol. 56, no. 7, pp. 1178 –1188, july 2008.
- [64] F. Daneshgaran, M. Laddomada, F. Mesiti, and M. Mondin, "Unsaturated Throughput Analysis of IEEE 802.11 in Presence of Non Ideal Transmission Channel and Capture Effects," *Wireless Communications, IEEE Transactions on*, vol. 7, no. 4, pp. 1276 –1286, april 2008.
- [65] G. Sutton, R. P. Liu, X. Yang, and I. Collings, "Modelling Capture Effect for 802.11 DCF under Rayleigh Fading," in *Communications* (ICC), 2010 IEEE International Conference on, may 2010, pp. 1 –6.
- [66] S.-J. Han, T. Nandagopal, Y. Bejerano, and H.-G. Choi, "Analysis of Spatial Unfairness in Wireless LANs," in *INFOCOM 2009, IEEE*, april 2009, pp. 2043 –2051.
- [67] S. Ganu, K. Ramachandran, M. Gruteser, I. Seskar, and J. Deng, "Methods for restoring MAC layer fairness in IEEE 802.11 networks with physical layer capture," in *Proceedings of the 2nd international* workshop on Multi-hop ad hoc networks: from theory to reality, ser. REALMAN '06, 2006, pp. 7–14.
- [68] J. Jeong, S. Choi, J. Yoo, S. Lee, and C.-K. Kim, "Physical layer capture aware MAC for WLANs," *Wireless Networks*, pp. 1–14, 2012.
- [69] N. Abramson, "THE ALOHA SYSTEM: another alternative for computer communications," in *Proceedings of the November 17-19, 1970, fall joint computer conference*, ser. AFIPS '70 (Fall), 1970, pp. 281– 285.
- [70] F. Tobagi and L. Kleinrock, "Packet Switching in Radio Channels: Part II–The Hidden Terminal Problem in Carrier Sense Multiple-Access and the Busy-Tone Solution," *Communications, IEEE Transactions on*, vol. 23, no. 12, pp. 1417 – 1433, dec 1975.
- [71] Z. Haas and J. Deng, "Dual busy tone multiple access (DBTMA)-a multiple access control scheme for ad hoc networks," *Communications, IEEE Transactions on*, vol. 50, no. 6, pp. 975 –985, jun 2002.
- [72] P. Karn, "MACA A New Channel Access Method for Packet Radio," in ARRL/CRRL Amateur Radio 9th Computer Networking Conference, September 1990.
- [73] V. Bharghavan, A. Demers, S. Shenker, and L. Zhang, "MACAW: a media access protocol for wireless LAN's," in *Proceedings of the conference on Communications architectures, protocols and applications*, ser. SIGCOMM '94, 1994, pp. 212–225.
- [74] C. L. Fullmer and J. J. Garcia-Luna-Aceves, "Floor acquisition multiple access (FAMA) for packet-radio networks," in *Proceedings of the* conference on Applications, technologies, architectures, and protocols for computer communication, ser. SIGCOMM '95, 1995, pp. 262–273.
- [75] —, "Solutions to hidden terminal problems in wireless networks," in Proceedings of the ACM SIGCOMM '97 conference on Applications, technologies, architectures, and protocols for computer communication, ser. SIGCOMM '97, 1997, pp. 39–49.
- [76] F. Borgonovo, A. Capone, M. Cesana, and L. Fratta, "RR-ALOHA, a Reliable R-ALOHA broadcast channel for ad-hoc inter-vehicle communication networks," in *Proceedings of Med-Hoc-Net 2002*, 2002.
- [77] K. A. Rahman and K. E. Tepe, "Extended sliding frame R-Aloha: Medium access control (MAC) protocol for mobile networks," Ad Hoc Networks, vol. 10, no. 6, pp. 1017 – 1027, 2012.
- [78] S. Gollakota and D. Katabi, "Zigzag decoding: combating hidden terminals in wireless networks," in *Proceedings of the ACM SIGCOMM* 2008 conference on Data communication, ser. SIGCOMM '08, 2008, pp. 159–170.
- [79] D. Halperin, T. Anderson, and D. Wetherall, "Taking the sting out of carrier sense: interference cancellation for wireless LANs," in *Proceedings of the 14th ACM international conference on Mobile computing and networking*, ser. MobiCom '08, 2008, pp. 339–350.
- [80] S. Khurana, A. Kahol, and A. Jayasumana, "Effect of hidden terminals on the performance of IEEE 802.11 MAC protocol," in *Local Computer Networks*, 1998. LCN '98. Proceedings., 23rd Annual Conference on, oct 1998, pp. 12 –20.
- [81] S. Khurana, A. Kahol, S. Gupta, and P. Srimani, "Performance evaluation of distributed co-ordination function for IEEE 802.11 wireless LAN protocol in presence of mobile and hidden terminals," in *Mod-*

eling, Analysis and Simulation of Computer and Telecommunication Systems, 1999. Proceedings. 7th International Symposium on, 1999, pp. 40–47.

- [82] K. Xu, M. Gerla, and S. Bae, "How effective is the IEEE 802.11 RTS/CTS handshake in ad hoc networks," in *Global Telecommunications Conference*, 2002. GLOBECOM '02. IEEE, vol. 1, nov. 2002, pp. 72 – 76 vol.1.
- [83] —, "Effectiveness of RTS/CTS handshake in IEEE 802.11 based ad hoc networks," Ad Hoc Networks, vol. 1, no. 1, pp. 107 – 123, 2003.
- [84] F. Liu, J. Lin, Z. Tao, T. Korakis, E. Erkip, and S. Panwar, "The Hidden Cost of Hidden Terminals," in *Communications (ICC), 2010 IEEE International Conference on*, may 2010, pp. 1–6.
- [85] A. Iyer, C. Rosenberg, and A. Karnik, "What is the right model for wireless channel interference?" Wireless Communications, IEEE Transactions on, vol. 8, no. 5, pp. 2662 –2671, may 2009.
- [86] A. Tsertou, D. Laurenson, and J. Thompson, "A new approach for the throughput analysis of ieee 802.11 in networks with hidden terminals," in *Int. Workshop Wireless Ad-hoc Netw.*, 2005.
- [87] H. Wu, F. Zhu, Q. Zhang, and Z. Niu, "Analysis of IEEE 802.11 DCF with Hidden Terminals," in *Global Telecommunications Conference*, 2006. GLOBECOM '06. IEEE, dec. 2006, pp. 1 –5.
- [88] O. Ekici and A. Yongacoglu, "IEEE 802.11a Throughput Performance with Hidden Nodes," *Communications Letters, IEEE*, vol. 12, no. 6, pp. 465 –467, june 2008.
- [89] A. Tsertou and D. Laurenson, "Revisiting the Hidden Terminal Problem in a CSMA/CA Wireless Network," *Mobile Computing, IEEE Transactions on*, vol. 7, no. 7, pp. 817–831, july 2008.
- [90] B. Jang and M. Sichitiu, "IEEE 802.11 Saturation Throughput Analysis in the Presence of Hidden Terminals," *Networking, IEEE/ACM Transactions on*, vol. 20, no. 2, pp. 557 –570, april 2012.
- [91] R. Chandra, R. Mahajan, T. Moscibroda, R. Raghavendra, and P. Bahl, "A case for adapting channel width in wireless networks," in *Proceed*ings of the ACM SIGCOMM 2008 conference on Data communication, ser. SIGCOMM '08, Seattle, WA, USA, 2008, pp. 135–146.
- [92] V. Shrivastava, S. Rayanchu, J. Yoonj, and S. Banerjee, "802.11n under the microscope," in *Proceedings of the 8th ACM SIGCOMM conference* on Internet measurement, ser. IMC '08, Vouliagmeni, Greece, 2008, pp. 105–110.
- [93] K. Pelechrinis, T. Salonidis, H. Lundgren, and N. Vaidya, "Experimental characterization of 802.11n link quality at high rates," in *Proceedings of the fifth ACM international workshop on Wireless network testbeds, experimental evaluation and characterization*, ser. WiNTECH '10, Chicago, Illinois, USA, 2010, pp. 39–46.
- [94] L. Deek, E. Garcia-Villegas, E. Belding, S.-J. Lee, and K. Almeroth, "The impact of channel bonding on 802.11n network management," in *Proceedings of the Seventh Conference on emerging Networking EXperiments and Technologies*, ser. CoNEXT '11, Tokyo, Japan, 2011, pp. 11:1–11:12.
- [95] M. Y. Arslan, K. Pelechrinis, I. Broustis, S. Singh, S. V. Krishnamurthy, S. Addepalli, and K. Papagiannaki, "ACORN: An Auto-Configuration Framework for 802.11 n WLANs," *Networking, IEEE/ACM Transactions on*, to appear.
- [96] K. Pelechrinis, I. Broustis, T. Salonidis, S. V. Krishnamurthy, and P. Mohapatra, "Design and deployment considerations for high performance MIMO testbeds," in *Proceedings of the 4th Annual International Conference on Wireless Internet*, ser. WICON '08, Maui, Hawaii, 2008, pp. 53:1–53:9.
- [97] L. Zheng and D. Tse, "Diversity and multiplexing: a fundamental tradeoff in multiple-antenna channels," *Information Theory, IEEE Transactions on*, vol. 49, no. 5, pp. 1073–1096, May 2003.
- [98] I. Pefkianakis, Y. Hu, S. H. Wong, H. Yang, and S. Lu, "MIMO rate adaptation in 802.11n wireless networks," in *Proceedings of the sixteenth annual international conference on Mobile computing and networking*, ser. MobiCom '10, Chicago, Illinois, USA, 2010, pp. 257– 268.
- [99] I. Pefkianakis, S.-B. Lee, and S. Lu, "Towards MIMO-Aware 802.11n Rate Adaptation," *Networking, IEEE/ACM Transactions on*, to appear.
- [100] T. Li, Q. Ni, D. Malone, D. Leith, Y. Xiao, and T. Turletti, "Aggregation With Fragment Retransmission for Very High-Speed WLANs," *Networking, IEEE/ACM Transactions on*, vol. 17, no. 2, pp. 591–604, April 2009.
- [101] Y. Lin and V. Wong, "Frame Aggregation and Optimal Frame Size Adaptation for IEEE 802.11n WLANs," in *Global Telecommunications Conference*, 2006. GLOBECOM '06. IEEE, 2006, pp. 1–6.

- [102] O. Abu-Sharkh and M. Abdelhadi, "The impact of multi-rate operation on A-MSDU, A-MPDU and block acknowledgment in greenfield IEEE802.11n wireless LANs," in *Wireless Advanced (WiAd)*, 2011, pp. 116–121.
- [103] M. Zubelda, A. Ferragut, and F. Paganini, "Averting speed inefficiency in multirate wireless networks through queueing and aggregation," in *IEEE Global Communications Conference (GLOBECOM 2012)*, Anaheim, USA, Dec. 2012.
- [104] M. Kim and C. Choi, "Hidden Node Detection in IEEE 802.11n Wireless LANs," Vehicular Technology, IEEE Transactions on, To appear.
- [105] E. H. Ong, J. Kneckt, O. Alanen, Z. Chang, T. Huovinen, and T. Nihtila, "IEEE 802.11ac: Enhancements for very high throughput WLANs," in Personal Indoor and Mobile Radio Communications (PIMRC), 2011 IEEE 22nd International Symposium on, 2011, pp. 849–853.
- [106] Q. Spencer, C. Peel, A. Swindlehurst, and M. Haardt, "An introduction to the multi-user MIMO downlink," *Communications Magazine, IEEE*, vol. 42, no. 10, pp. 60–67, 2004.
- [107] D. Chan, T. Berger, and L. Tong, "Carrier Sense Multiple Access Communications on Multipacket Reception Channels: Theory and Applications to IEEE 802.11 Wireless Networks," *Communications, IEEE Transactions on*, vol. 61, no. 1, pp. 266–278, 2013.
- [108] "IEEE Standard for Information technology–Telecommunications and information exchange between systems–Local and metropolitan area networks–Specific requirements-Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band," *IEEE Std 802.11ad-2012 (Amendment to IEEE Std 802.11-2012, as amended by IEEE Std 802.11ae-2012 and IEEE Std 802.11aa-2012)*, 2012.
- [109] "WirelessHD." [Online]. Available: http://www.wirelesshd.org/
- [110] E. Charfi, L. Chaari, and L. Kamoun, "PHY/MAC Enhancements and QoS Mechanisms for Very High Throughput WLANs: A Survey," *Communications Surveys Tutorials, IEEE*, To appear.
- [111] J. Fang, K. Tan, Y. Zhang, S. Chen, L. Shi, J. Zhang, Y. Zhang, and Z. Tan, "Fine-Grained Channel Access in Wireless LAN," *Networking*, *IEEE/ACM Transactions on*, To appear.
- [112] K. C.-J. Lin, S. Gollakota, and D. Katabi, "Random access heterogeneous MIMO networks," in *Proceedings of the ACM SIGCOMM 2011 conference*, Toronto, Ontario, Canada, 2011, pp. 146–157.
- [113] S. Sen, R. Roy Choudhury, and S. Nelakuditi, "No time to countdown: migrating backoff to the frequency domain," in *Proceedings of the 17th* annual international conference on Mobile computing and networking, ser. MobiCom '11, Las Vegas, Nevada, USA, 2011, pp. 241–252.
- [114] X. Feng, J. Zhang, Q. Zhang, and B. Li, "Use your frequency wisely: Explore frequency domain for channel contention and ACK," in *INFOCOM*, 2012 Proceedings IEEE, 2012, pp. 549–557.
- [115] A. Lindgren, A. Almquist, and O. Schelen, "Evaluation of quality of service schemes for IEEE 802.11 wireless LANs," in *Proceedings of the 26th Annual IEEE Conference on Local Computer Networks (LCN 2001)*, 2001, pp. 348–351.
- [116] G. Bianchi, I. Tinnirello, and L. Scalia, "Understanding 802.11e contention-based prioritization mechanisms and their coexistence with legacy 802.11 stations," *Network, IEEE*, vol. 19, no. 4, pp. 28–34, July-Aug. 2005.
- [117] J. Robinson and T. Randhawa, "Saturation throughput analysis of IEEE 802.11e Enhanced distributed coordination function," *IEEE Journal on Selected Areas in Communications*, vol. 22, no. 5, pp. 917–928, 2004.
- [118] V. Ramaiyan, A. Kumar, and E. Altman, "Fixed Point Analysis of Single Cell IEEE 802.11e WLANs: Uniqueness and Multistability," *Networking, IEEE/ACM Transactions on*, vol. 16, no. 5, pp. 1080– 1093, Oct. 2008.
- [119] A. Banchs, A. Azcorra, C. Garcia, and R. Cuevas, "Applications and challenges of the 802.11e EDCA mechanism: an experimental study," *Network, IEEE*, vol. 19, no. 4, pp. 52–58, July-Aug. 2005.
- [120] D. Gao, J. Cai, and K. N. Ngan, "Admission control in IEEE 802.11e wireless LANs," *Network, IEEE*, vol. 19, no. 4, pp. 6–13, July-Aug. 2005.
- [121] Y. Liu and M. Meng, "Survey of Admission Control Algorithms in IEEE 802.11e Wireless LANs," in *International Conference on Future Computer and Communication. FCC '09.*, June 2009, pp. 230–233.
- [122] P. Serrano, A. Banchs, P. Patras, and A. Azcorra, "Optimal Configuration of 802.11e EDCA for Real-Time and Data Traffic," *Vehicular*

Technology, IEEE Transactions on, vol. 59, no. 5, pp. 2511–2528, Jun 2010.

- [123] P. Serrano, P. Patras, A. Mannocci, V. Mancuso, and A. Banchs, "Control theoretic optimization of 802.11 WLANs: Implementation and experimental evaluation," *Computer Networks*, vol. 57, no. 1, pp. 258 – 272, 2013.
- [124] G. Hiertz, D. Denteneer, L. Stibor, Y. Zang, X. Costa, and B. Walke, "The IEEE 802.11 universe," *Communications Magazine, IEEE*, vol. 48, no. 1, pp. 62–70, January 2010.
- [125] "IEEE Standard for Local and metropolitan area networks: Media Access Control (MAC) Bridges," *IEEE Std 802.1D-2004 (Revision of IEEE Std 802.1D-1998*), 2004.
- [126] T. Higgins, "Does Wi-Fi MultiMedia (WMM) Really Do Anything?" 2009. [Online]. Available: http://www.smallnetbuilder.com/wireless/wireless-features/30833does-wi-fi-multimedia-wmm-really-do-anything-part-1
- [127] G. Bianchi, A. Di Stefano, C. Giaconia, L. Scalia, G. Terrazzino, and I. Tinnirello, "Experimental Assessment of the Backoff Behavior of Commercial IEEE 802.11b Network Cards," in *INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE*, May 2007, pp. 1181–1189.
- [128] R. Bolla, R. Rapuzzi, and M. Repetto, "On the effectiveness of IEEE 802.11e implementations in real hardware," in 6th International Symposium on Wireless Communication Systems (ISWCS 2009)., Sept. 2009, pp. 303–307.
- [129] D. Tang and M. Baker, "Analysis of a local-area wireless network," in *Proceedings of the 6th annual international conference on Mobile computing and networking*, ser. MobiCom '00, 2000, pp. 1–10.
- [130] —, "Analysis of a Metropolitan-Area Wireless Network," Wireless Networks, vol. 8, pp. 107–120, 2002.
- [131] A. Balachandran, G. M. Voelker, P. Bahl, and P. V. Rangan, "Characterizing user behavior and network performance in a public wireless LAN," in *Proceedings of the 2002 ACM SIGMETRICS international conference on Measurement and modeling of computer systems*, ser. SIGMETRICS '02, 2002, pp. 195–205.
- [132] M. Balazinska and P. Castro, "Characterizing mobility and network usage in a corporate wireless local-area network," in *Proceedings of the 1st international conference on Mobile systems, applications and services*, ser. MobiSys '03, 2003, pp. 303–316.
- [133] D. Schwab and R. Bunt, "Characterising the use of a campus wireless network," in *INFOCOM 2004. Twenty-third AnnualJoint Conference* of the IEEE Computer and Communications Societies, vol. 2, march 2004, pp. 862–870.
- [134] T. Henderson, D. Kotz, and I. Abyzov, "The changing usage of a mature campus-wide wireless network," in *Proceedings of the 10th annual international conference on Mobile computing and networking*, ser. MobiCom '04, 2004, pp. 187–201.
- [135] D. Kotz and K. Essien, "Analysis of a campus-wide wireless network," Wirel. Netw., vol. 11, no. 1-2, pp. 115–133, jan. 2005.
- [136] T. Henderson, D. Kotz, and I. Abyzov, "The changing usage of a mature campus-wide wireless network," *Computer Networks*, vol. 52, no. 14, pp. 2690 – 2712, 2008.
- [137] M. Afanasyev, T. Chen, G. M. Voelker, and A. C. Snoeren, "Analysis of a mixed-use urban WiFi network: when metropolitan becomes neapolitan," in *Proceedings of the 8th ACM SIGCOMM conference* on Internet measurement, ser. IMC '08, 2008, pp. 85–98.
- [138] M. Afanasyev, T. Chen, G. Voelker, and A. Snoeren, "Usage Patterns in an Urban WiFi Network," *Networking, IEEE/ACM Transactions on*, vol. 18, no. 5, pp. 1359 –1372, oct. 2010.
- [139] C. Na, J. Chen, and T. Rappaport, "Measured Traffic Statistics and Throughput of IEEE 802.11b Public WLAN Hotspots with Three Different Applications," *Wireless Communications, IEEE Transactions* on, vol. 5, no. 11, pp. 3296 –3305, november 2006.
- [140] E. Halepovic, C. Williamson, and M. Ghaderi, "Wireless data traffic: a decade of change," *Network, IEEE*, vol. 23, no. 2, pp. 20 –26, march 2009.
- [141] J. Yeo, M. Youssef, and A. Agrawala, "A framework for wireless LAN monitoring and its applications," in *Proceedings of the 3rd ACM* workshop on Wireless security, ser. WiSe '04, 2004, pp. 70–79.
- [142] J. Yeo, M. Youssef, T. Henderson, and A. Agrawala, "An accurate technique for measuring the wireless side of wireless networks," in *Papers* presented at the 2005 workshop on Wireless traffic measurements and modeling, ser. WiTMeMo '05, 2005, pp. 13–18.

- [143] M. Rodrig, C. Reis, R. Mahajan, D. Wetherall, and J. Zahorjan, "Measurement-based characterization of 802.11 in a hotspot setting," in *Proceedings of the 2005 ACM SIGCOMM workshop on Experimental* approaches to wireless network design and analysis, ser. E-WIND '05, 2005, pp. 5–10.
- [144] A. Schulman, D. Levin, and N. Spring, "On the Fidelity of 802.11 Packet Traces," in *Passive and Active Network Measurement*, ser. Lecture Notes in Computer Science, 2008, vol. 4979, pp. 132–141.
- [145] R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan, "Analyzing the MAC-level behavior of wireless networks in the wild," in *Proceedings* of the 2006 conference on Applications, technologies, architectures, and protocols for computer communications, ser. SIGCOMM '06, 2006, pp. 75–86.
- [146] Y.-C. Cheng, J. Bellardo, P. Benkö, A. C. Snoeren, G. M. Voelker, and S. Savage, "Jigsaw: solving the puzzle of enterprise 802.11 analysis," in *Proceedings of the 2006 conference on Applications, technologies, architectures, and protocols for computer communications,* ser. SIG-COMM '06, 2006, pp. 39–50.
- [147] K. Lakshminarayanan, S. Seshan, and P. Steenkiste, "Understanding 802.11 performance in heterogeneous environments," in *Proceedings of the 2nd ACM SIGCOMM workshop on Home networks*, ser. HomeNets '11, 2011, pp. 43–48.
- [148] V. Shrivastava, S. Rayanchu, S. Banerjee, and K. Papagiannaki, "PIE in the sky: online passive interference estimation for enterprise WLANs," in *Proceedings of the 8th USENIX conference on Networked systems design and implementation*, ser. NSDI'11, 2011, pp. 25–25.
- [149] D. Dujovne, T. Turletti, and F. Filali, "A Taxonomy of IEEE 802.11 Wireless Parameters and Open Source Measurement Tools," *Communications Surveys Tutorials, IEEE*, vol. 12, no. 2, pp. 249 –262, 2010.
- [150] A. Adya, P. Bahl, R. Chandra, and L. Qiu, "Architecture and techniques for diagnosing faults in IEEE 802.11 infrastructure networks," in *Proceedings of the 10th annual international conference on Mobile computing and networking*, ser. MobiCom '04, 2004, pp. 30–44.
- [151] P. Bahl, R. Chandra, J. Padhye, L. Ravindranath, M. Singh, A. Wolman, and B. Zill, "Enhancing the security of corporate Wi-Fi networks using DAIR," in *Proceedings of the 4th international conference on Mobile* systems, applications and services, ser. MobiSys '06, 2006, pp. 1–14.
- [152] A. Sheth, C. Doerr, D. Grunwald, R. Han, and D. Sicker, "MOJO: a distributed physical layer anomaly detection system for 802.11 WLANs," in *Proceedings of the 4th international conference on Mobile* systems, applications and services, ser. MobiSys '06, 2006, pp. 191– 204.
- [153] Y.-C. Cheng, M. Afanasyev, P. Verkaik, P. Benkö, J. Chiang, A. C. Snoeren, S. Savage, and G. M. Voelker, "Automating cross-layer diagnosis of enterprise wireless networks," in *Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications*, ser. SIGCOMM '07, 2007, pp. 25–36.
- [154] A. Mahanti, C. Williamson, and M. Arlitt, "Remote analysis of a distributed WLAN using passive wireless-side measurement," *Performance Evaluation*, vol. 64, no. 912, pp. 909 – 932, 2007.
- [155] A. P. Jardosh, K. N. Ramachandran, K. C. Almeroth, and E. M. Belding-Royer, "Understanding link-layer behavior in highly congested IEEE 802.11b wireless networks," in *Proceedings of the 2005 ACM* SIGCOMM workshop on Experimental approaches to wireless network design and analysis, ser. E-WIND '05, 2005, pp. 11–16.
- [156] —, "Understanding congestion in IEEE 802.11b wireless networks," in Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement, ser. IMC '05. Berkeley, CA, USA: USENIX Association, 2005, pp. 25–25. [Online]. Available: http://dl.acm.org/citation.cfm?id=1251086.1251111
- [157] G. Bianchi, F. Formisano, and D. Giustiniano, "802.11b/g link level measurements for an outdoor wireless campus network," in World of Wireless, Mobile and Multimedia Networks, 2006. WoWMoM 2006. International Symposium on a, june 2006.
- [158] D. Giustiniano, G. Bianchi, L. Scalia, and I. Tinnirello, "An Explanation for Unexpected 802.11 Outdoor Link-level Measurement Results," in *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*, april 2008, pp. 2432 –2440.
- [159] I. Tinnirello, D. Giustiniano, L. Scalia, and G. Bianchi, "On the sideeffects of proprietary solutions for fading and interference mitigation in IEEE 802.11b/g outdoor links," *Computer Networks*, vol. 53, no. 2, pp. 141 – 152, 2009.
- [160] 3rd Generation Partnership Project; Technical Specification Group Services and System Aspects, "3GPP system to Wireless Local Area

Network (WLAN) interworking; System description (Release 11)," Sept. 2012.

- [161] 3rd Generation Partnership Project; Technical Specification Group GSM/EDGE Radio Access Network, "Generic Access Network (GAN); Stage 2 (Release 11)," Sept. 2012.
- [162] "IEEE Standard for Local and Metropolitan Area Networks- Part 21: Media Independent Handover," *IEEE Std 802.21-2008*, 2009.
- [163] E. Piri and K. Pentikousis, "IEEE 802.21: Media-Independent Handover Services," *The Internet Protocol Journal*, vol. 12, no. 2, June 2009.
- [164] R. Ferrus, O. Sallent, and R. Agusti, "Interworking in heterogeneous wireless networks: Comprehensive framework and future trends," *Wireless Communications, IEEE*, vol. 17, no. 2, pp. 22–31, 2010.
- [165] M. Kassar, B. Kervella, and G. Pujolle, "An overview of vertical handover decision strategies in heterogeneous wireless networks," *Computer Communications*, vol. 31, no. 10, pp. 2607 – 2620, 2008.
- [166] J. Márquez-Barja, C. T. Calafate, J.-C. Cano, and P. Manzoni, "An overview of vertical handover techniques: Algorithms, protocols and tools," *Computer Communications*, vol. 34, no. 8, pp. 985 – 997, 2011.
- [167] T. Ali and M. Saquib, "Analytical Framework for WLAN-Cellular Voice Handover Evaluation," *Mobile Computing, IEEE Transactions* on, vol. 12, no. 3, pp. 447–460, 2013.
- [168] O. Marce, H.-H. Tran, and B. Tuffin, "Double-sided auctions applied to vertical handover for mobility management in wireless networks," *Journal of Network and Systems Management*, pp. 1–24, 2013.
- [169] FCC, "Second Report and Order And Memorandum Opinion and Order, ET Docket No. 08-260," Nov. 2008.
- [170] K. Shin, H. Kim, A. Min, and A. Kumar, "Cognitive radios for dynamic spectrum access: from concept to reality," *Wireless Communications, IEEE*, vol. 17, no. 6, pp. 64–74, 2010.
- [171] "IEEE Standard for Information Technology–Telecommunications and information exchange between systems Wireless Regional Area Networks (WRAN)–Specific requirements Part 22: Cognitive Wireless RAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Policies and Procedures for Operation in the TV Bands," *IEEE Std 802.22-2011*, 2011.
- [172] J. Mitola and J. Maguire, G.Q., "Cognitive radio: making software radios more personal," *Personal Communications, IEEE*, vol. 6, no. 4, pp. 13–18, 1999.
- [173] R. C. Qiu, Z. Hu, H. Li, and M. C. Wicks, *Cognitive Radio Communi*cation and Networking: Principles and Practice. John Wiley & Sons, 2012.
- [174] "PhySimWiFi for NS-3 An IEEE 802.11 OFDM physical layer simulator within NS-3." [Online]. Available: http://dsn.tm.kit.edu/ns3physim.php
- [175] A. Cardenas, S. Radosavac, and J. Baras, "Evaluation of Detection Algorithms for MAC Layer Misbehavior: Theory and Experiments," *Networking, IEEE/ACM Transactions on*, vol. 17, no. 2, pp. 605–617, 2009.