

Binary Symmetric Channel Based Aggregation with Coding for 802.11n WLANs*

Xiaomin Chen, Vijay G. Subramanian, and Douglas J. Leith

Hamilton Institute, NUI Maynooth, Ireland
{Xiaomin.Chen,Vijay.Subramanian,Doug.Leith}@nuim.ie

Abstract. In this paper we present the first detailed theoretical analysis of the potential performance gains of adopting a BSC paradigm in 802.11 WLANs. Importantly, we also consider the multi-user channel aspect of a WLAN i.e. that transmissions are inherently broadcast in nature. We find that increases in network throughput of more than 100% are possible over a wide range of SNRs. These performance gains are achieved exclusively through software rather than hardware changes.

Keywords: binary symmetric channel, packet aggregation, superposition coding, multi-user channel.

1 Introduction

In this paper we consider the potential benefits of viewing the channel provided by an 802.11 WLAN as a binary symmetric channel (BSC), as opposed to a more conventional packet erasure channel (PEC). That is, rather than simply discarding corrupted frames, we consider viewing a received frame as a binary vector in which an unknown subset of bits have been “flipped”. Although some bits are corrupted/flipped, we can often still extract useful information from corrupted frames thereby increasing the effective transmit rate between wireless stations. This is motivated by a number of observations.

Firstly, for 802.11a/g it has been shown in [1] using experimental measurements that even with a 10 – 30% packet erasure rate typically only a small fraction (usually < 1%) of the bits within corrupted packets are in error. Thus, although noisy, the corrupted packets potentially provide a reasonable channel through which we can transmit information. Modeling the bit error process as a BSC, a simple theoretical analysis (ignoring CSMA/CA features such as collision losses) indicates the potential for significant performance gains. For example, Fig. 1 compares the BSC throughput capacity versus the packet erasure throughput capacity for the set of modulations/rates available in 802.11a/g. See the next section for details on the calculations used to obtain this figure. It can be seen that throughput improvements of 100% or more are indicated over a wide range of SNRs.

* This material is based upon works supported by the Science Foundation Ireland under Grant No. 07/IN.1/I901.

Secondly, recent breakthroughs in efficient capacity-approaching error correction codes such as LDPC codes make these performance gains practically achievable. Since such coding would be introduced above the MAC layer, it is compatible with standard 802.11 hardware i.e. it requires only a software change and so the performance gains essentially come for “free”.




Fig. 1. 802.11a/g BSC and PEC capacities vs. SNR, Rayleigh physical channel. Packet erasure capacities are shown for frame sizes of both 1024 bytes and 8000 bytes.

Thirdly, this BSC paradigm dovetails with the trend towards greater decoupling of the unit of transmission (i.e. frames) used at the MAC/PHY layer from the unit of transmission (i.e. packets) used at the IP layer. For example, to maintain throughput efficiency at high PHY rates, the recent 802.11n standard [2] supports transmission of large frames formed by aggregating multiple packets together. This is because increasing PHY rates leads to faster transmission of the MAC frame payload, but overheads such as PHY headers and MAC contention time typically do not decrease at the same rate and thus begin to dominate the frame transmission time unless amortised across multiple packets, e.g. see [3]. A logical extension is to consider aggregation of packets destined to different receivers into a single MAC frame. Such multi-destination aggregation is currently the subject of much interest because we expect that often there simply may not be enough traffic to an individual destination to always allow large frames to be formed and the network efficiency quickly degrades when small frames are used. This scenario is particularly important for the higher PHY rates proposed for 802.11n [2] and when applications like VoIP [4,5], email and web-browsing are considered. Also, multicast traffic is expected to become increasingly important in WLANs in view of the escalating demand for real-time multimedia applications.

In [6,7] packet erasures are modelled assuming a bit-level BSC model, but the analysis is otherwise based on a packet erasure channel i.e. they do not consider the BSC paradigm of this paper. While [1] demonstrates that packets that fail the CRC check only have a few bits in error, the authors do not consider a BSC modeling paradigm and, in particular, the multi-user BSC paradigm that we are interested in.

Our contributions include:

- The first detailed theoretical analysis of the potential performance gains of adopting a BSC paradigm in 802.11 WLANs. This includes important 802.11 MAC features such as the framing overhead and the overhead of CSMA/CA contention and collisions.
- Consideration of the multi-user channel aspect of a WLAN in the context of the BSC paradigm, i.e. frame transmissions are inherently broadcast in nature and so may be received by multiple stations. We explore this for both unicast and multicast traffic. In the unicast case we consider, in particular, the potential performance gains of using a BSC paradigm for multi-destination aggregation.

2 Channel Models

We begin by reviewing the performance of the PHY layer modulation and FEC available in 802.11 a/g. Since 802.11n uses very similar modulation and FEC (adding a small number of additional OFDM tones and a new 5/6 code rate), our analysis carries over essentially unchanged for single antenna systems. We leave consideration of MIMO 802.11n systems as future work.

2.1 Packet Erasure Rate of Rayleigh Channel

The indoor fading environment is modelled as a Rayleigh channel. Assuming hard-decision Viterbi decoding is used, analytic expressions for the bit error rate (BER) when using each of the various 802.11a/g modulations are described in [8]. The analysis considers Nakagami- m fading channels, and we note that a Rayleigh channel corresponds to a Nakagami channel with $m = 1$. The demodulation BER needs to be adjusted to take account of the error correction provided by convolutional coding. As bit errors in the output of Viterbi decoder are no longer independent, for a packet erasure model, an upper bound on the packet erasure rate (PER) is given in [9] by $p_e = 1 - (1 - p_u)^L$, where L is the length of packet in bits, and p_u is the union bound on the first-event error probability of Viterbi decoding [10]. Fig. 2 shows the resulting PER versus SNR curves for each of the 802.11a/g transmission modes for a packet length of 1024 bytes. The channel capacity is $R \times (1 - p_e)$ at PHY rate R .

2.2 BSC Crossover Probability of Rayleigh Channel

The crossover probability p , i.e. the (symmetric) probability of a 1 being changed into 0 and *vice-versa*, is the probability of a bit error in a BSC model. An upper bound on the bit crossover probability is the sum of the expected number of erroneous bits for all possible incorrectly selected paths while performing Viterbi decoding, as described in [10]. For a binary-tree convolutional code, the expression for this upper bound is given by $\sum_{d=d_{free}}^{\infty} \beta_d p_2(d)$, where $p_2(d)$ is




Fig. 2. Packet erasure rate (PER) vs. SNR, $L = 1024B$, Rayleigh channel

the probability that an incorrect path with Hamming distance d is selected, and β_d is the number of bit errors totaled over all paths with Hamming distance d . Fig. 3 shows the resulting crossover probability vs SNR curves for the 802.11a/g OFDM modulations/rates and a Rayleigh channel. The BSC capacity at PHY rate R is $R \times (1 - H(p))$, where $H(p)$ is the entropy function of a Bernoulli(p) random variable, i.e. $H(p) = -p \log_2(p) - (1 - p) \log_2(1 - p)$.




Fig. 3. Crossover probability vs. SNR, Rayleigh channel

The overall capacity curves shown in Fig. 1 are obtained by selecting the PHY modulation/rate R that maximises the capacity at each SNR. Note that these curves do not include the MAC layer framing overheads, contention time, collision losses etc. In the following sections we extend the analysis to include these overheads for both unicast and multicast traffic. We also extend the analysis to take account of the multi-user nature of the wireless channel, including multi-destination unicast aggregation.

3 Coding in Multi-user Channels

Owing to the broadcast nature of the wireless medium every transmission can be heard by all the receivers within the coverage area. However, the channel quality between the transmitter and every receiver, e.g. the access point (AP) and each client station in a WLAN, is generally different for every receiver, owing to a multitude of reasons such as differences in distance between the transmitter and receivers, obstacles such as walls when operating indoors and differences in the local interference environment. While the 802.11 standard allows transmissions to multiple receivers in one single frame, it constrains the bits transmitted in the same frame to use the same PHY modulation/rate. Hence, the state of the art in 802.11 WLANs is to send multi-user transmissions at the highest PHY rate which the client with the worst channel quality can support, so that all clients can decode the transmission; this is the recommendation for multicast traffic in the current 802.11 standards. Clearly, this is inefficient. One of the outcomes of our analysis is also the quantification of this inefficiency. Since the traditional approach drops packets once the CRC check fails, it is appropriately modelled as a PEC [11].

In contrast to PEC model, the BSC paradigm allows us to transmit information within a frame with different segments encoded with different levels of protection. In this manner we can transmit at different information rates to different destinations while using a single PHY modulation/rate exploiting developments in multi-terminal information theory [11].

The specific multi-user BSC paradigm we consider forms what is known as a physically degraded binary symmetric broadcast channel [11]. For this class of channels superposition coding [12] is known to be capacity-achieving. Superposition coding works by first picking an ordering of the users and constructing the code of every user based upon the codes of all the users before. Decoding starts from the last user for whom the decoder treats every other user's signal as noise. After the last user is decoded, its contribution is subtracted and the remaining users are decoded in a similar nested fashion. For a more concrete illustration of our particular setting, we limit the discussion to two classes of users where for illustration purposes we assume that class 2 experiences no errors and class 1 has a crossover probability $p \in (0, 1/2)$. We have a natural order in this setting where class 2 appears before class 1 in the encoding process. As mentioned earlier instead of partitioning a MAC frame into separate segments for each distinct message, in superposition coding the message (i.e. binary vector) U destined to class 1 and the message V destined to class 2 are summed, modulo 2, to yield the MAC frame body. The binary vector received by the class 1 user can then be viewed as message U corrupted by a bit flips due to a combination of channel noise and the summation with V . This is illustrated in Fig. 4. The message U destined to the class 1 user is first passed through a BSC with crossover probability β determined by the entropy of V . It is then transmitted over the physical BSC channel with crossover probability $p(R)$ at PHY rate R . The channel capacity for the class 1 user in bits per channel use is therefore

$$C_1 = 1 - H(\beta \circ p(R)) \tag{1}$$

where $\beta \circ p(R) = \beta(1 - p(R)) + (1 - \beta)p(R)$.

The channel capacity for the class 2 user is

$$C_2 = H(\beta) \tag{2}$$




Fig. 4. Physically degraded binary symmetric broadcast channel

From the discussion above it is clear that superposition coding can be a complex operation. A simpler but demonstrably sub-optimal choice is a time-sharing based coding scheme [11]. In our setting with two classes of users, time-sharing based coding is such that each MAC frame is partitioned (i.e. time-shared) into one portion intended for class 1 and the other intended for class 2. Using the same setting that is illustrated above for superposition coding, it is clear that the portion intended for class 2 is noise-free and thus not protected. However, the portion intended for class 1 is protected by an error correction code that allows information to be extracted even when bits within the frame are corrupted; the information rate is obviously reduced compared to a noise-free channel. We assume an ideal code with coding rate matched to the BSC capacity. Since the aggregated frame is large, capacity-approaching codes, e.g. LDPC codes [13], exist.

In this paper we present a performance analysis of superposition coding, which is known to be capacity-achieving in multi-user BSC channels, and of the simpler time-sharing coding scheme. We will show in results that there is minimal loss of optimality in using the time-sharing coding scheme. For both schemes our analysis indicates the potential for substantial performance gains over the traditional packet erasure channel paradigm. To our knowledge, this is the first such analysis of multi-user coding in 802.11 WLANs.

4 Modelling Analysis

Consider a wireless network with an AP and two classes of client stations, with n_1 stations in class 1 and n_2 in class 2. The AP has N_{D1} downlink flows destined for n_1 class 1 stations and N_{D2} downlink flows for n_2 class 2 stations, in general

$N_{D_1} \neq n_1$ and $N_{D_2} \neq n_2$. We assume all stations in the same class have the same SNR. Class 1 is located far from the AP with low SNR such that stations in this class are subjected to noisy reception; while class 2 lies within a region where stations have high SNR and thus experience reliable reception at any of the available PHY data rates. The analysis can be readily generalised to encompass situations where each user station has a different SNR, but the two-class case is sufficient to capture many important features of WLAN performance.

To ensure a fair comparison amongst different schemes it is not sufficient to simply compare sum-throughputs. Rather we also need to ensure that schemes provide comparable throughput fairness, since an approach may achieve throughput gains at the cost of increased unfairness. In the following we take a max-min fair approach and impose the fairness constraint that all flows achieve the same throughput. Extension of the analysis to other fairness criteria is, of course, possible.

4.1 Unicast

The vast majority of network traffic is unicast, and contention between multiple stations, with associated collision losses and increased CSMA/CA countdown time, is the norm. For the unicast scenario we set $N_{D_1} = n_1$ and $N_{D_2} = n_2$. The AP aggregates these $N_{D_1} + N_{D_2}$ downlink flows into a single large MAC frame and then transmits it at a single PHY rate. Each client station also has an uplink flow for the AP.

For simplicity we assume that all stations are saturated, although the analysis could be extended to include unsaturated operation using, for example, the approach in [14]. We also assume uplink transmissions by client stations are immediately acknowledged by the AP (rather than, for example, using a block ACK). Similarly, we assume that downlink transmissions are immediately acknowledged by client stations and, to make our analysis concrete, we adopt the approach described in [15] which uses the orthogonality of OFDM subcarriers to allow a group of user stations to transmit feedback signals at the same time, and thereby ACK collisions are avoided. However, these assumptions really just relate to the calculation of the MAC overheads and our analysis could be readily modified to account for alternative acking mechanisms¹.

MAC Model. Transmissions by the AP are subjected to collisions with competing uplink transmissions, while transmissions by client stations are subjected to collisions with the AP and uplink transmissions from other client stations. In the packet erasure setting, transmissions by class 1 stations are also subjected to noise losses, while those in class 2 are assumed to have a high SNR channel and

¹ In particular, for operation on standard 802.11n hardware we might tunnel data packets for multiple destinations via multicast MPDUs aggregated into an A-MPDU (thereby achieving standards-compliant multi-destination aggregation) and generate appropriate acknowledgements either by modifying the receiver NIC driver or at the application layer.

thus not subjected to noise. For class 1 stations we therefore have the probability that a transmission fails (due to collision and/or loss) is

$$p_{f_1} = 1 - (1 - p_{c21})(1 - p_{e21}) \quad (3)$$

where p_{e21} is the probability that the packet is erased due to noise, and p_{c21} is the collision probability of a transmission from a class 1 station,

$$p_{c21} = 1 - (1 - \tau_1)^{n_1-1}(1 - \tau_2)^{n_2}(1 - \tau_0) \quad (4)$$

with τ_0 , τ_1 and τ_2 being the attempt probability of AP, class 1 stations and class 2 stations, respectively. As the AP and class 2 stations have the same back-off probability (the probability that a transmission fails), according to Bianchi model [16], $\tau_0 = \tau_2$. The backoff probability for the AP or a class 2 station is

$$p_{f_2} = 1 - (1 - \tau_1)^{n_1}(1 - \tau_2)^{n_2} \quad (5)$$

The usual Bianchi [16] expression gives the relation between the attempt probability τ and the probability p_f that a transmission fails. However, we make use of expression (6) in [17] that builds upon the Bianchi expression taking into account a finite number of retransmission attempts and losses due to decoding errors.

Packet Erasure Channel. Similarly to the approach used in 802.11n A-MPDUs [2], we consider a situation where messages addressed to distinct destinations are aggregated together to form a single large MAC frame. We do not present results here without aggregation since the throughputs are strictly lower than when aggregation is used [3]. To keep our discussion concrete, we assume the frame format is as shown in Fig. 5. However, it is important to stress that this really just relates to the calculation of the MAC overheads and our analysis could be readily modified to account for alternative frame formats. In Fig. 5 a sub-header is prefixed to each IP packet to indicate its receiver address, source address and packet sequence information. An FCS checksum is used to detect corrupted packets. Since the sub-header already contains the receiver address, source address and sequence control, the MAC header removes these three fields, but keeps other fields unchanged from the standard 802.11 MAC header. For simplicity, we assume that the MAC header is transmitted at the same PHY rate as the PLCP header and thus can be assumed error-free in the following analysis, although we will relax this assumption in future work. The frame format is known to all user stations, so that each station can correctly locate its packet even if some bits in the frame are corrupted.

In the PEC scenario, the downlink transmission rate is determined by the worst client which has the lowest SNR, hence we have the downlink PHY rate R_1 used by the AP equal to the uplink PHY rate of class 1 stations R_{21} in order to meet the max-min fairness objective. For a given SNR and R_1 , the union bound on the first-event error probability of Viterbi decoding is $p_u(R_1)$. The packet erasure rate of an uplink packet from class 1 stations $p_{e21}(R_1)$ is then

$$p_{e21}(R_1) = 1 - (1 - p_u(R_1))^{L_{21}(R_1)} \quad (6)$$




Fig. 5. Erasure channel frame format [2]

where $L_{21}(R_1)$ is the frame size in bits, given by

$$L_{21}(R_1) = DBPS(R_1) \times \left\lceil \frac{(x_{21}(R_1) + L_{machdr} + L_{FCS}) \times 8 + 6 + 16}{DBPS(R_1)} \right\rceil \quad (7)$$

where $DBPS(R_1)$ represents data bits per symbol at PHY rate R_1 , $x_{21}(R_1)$ is the payload size in bytes of an uplink packet from a class 1 station, and L_{machdr} and L_{FCS} are the length in bytes of the standard MAC header and the FCS field. The expected payload delivered by an uplink packet of a class 1 station is

$$E_{pld_{21}}(R_1) = x_{21}(R_1)(1 - p_u(R_1))^{L_{21}(R_1)} \quad (8)$$

Let x_{22} denote the payload size in bytes of class 2 uplink transmissions. By assumption, uplink transmissions by class 2 stations are assumed to be loss free at all supported rates and so we may take the transmission rate $R_{22} = 54\text{Mbps}$. The expected payload of an uplink packet from a class 2 station at R_1 is

$$E_{pld_{22}}(R_1) = x_{22}(R_1) \quad (9)$$

Turning now to the AP, similar to the approach used in 802.11n, the aggregated MAC frame then consists of $N_{D1} + N_{D2}$ unicast packets. Let $x_{11}(R_1)$ and $x_{12}(R_1)$ denote the payload size in bytes to class 1 and class 2 stations at R_1 respectively. The length of a MAC frame L is thus

$$L = N_{D1}x_{11}(R_1) + N_{D2}x_{12}(R_1) + (N_{D1} + N_{D2})(L_{subhdr} + L_{FCS}) \quad (10)$$

where L_{subhdr} is the sub-header length. The expected payload delivered to a class 1 station in a downlink packet is

$$E_{pld_{11}}(R_1) = x_{11}(R_1) \times (1 - p_u(R_1))^{(x_{11}(R_1) + L_{subhdr} + L_{FCS}) \times 8} \quad (11)$$

The expected payload delivered to a class 2 station in a downlink packet is

$$E_{pld_{12}}(R_1) = x_{12}(R_1) \quad (12)$$

To equalize the throughput of each flow, we require

$$x_{12}(R_1) = x_{11}(R_1)(1 - p_u(R_1))^{(x_{11}(R_1) + L_{subhdr} + L_{FCS}) \times 8} \quad (13)$$

$$x_{22}(R_1) = x_{12}(R_1) \tag{14}$$

$$\tau_1(1 - \tau_2)(1 - p_u(R_1))^{L_{21}(R_1)}x_{21}(R_1) = \tau_2(1 - \tau_1)x_{22}(R_1) \tag{15}$$

For a given R_1 and a fixed L we can solve equations (10) and (13) to obtain x_{11} and x_{12} . As τ_0 , τ_1 and τ_2 also depend on R_1 , combining the expressions for τ_0 , τ_1 and τ_2 from the MAC model with expression (15), we can solve to obtain $x_{21}(R_1)$.

To obtain the throughputs it remains to derive expressions for the expected duration of a MAC slot. There are four possible types of MAC slot. We consider each in turn.

- Type 1 - AP transmits: Observe that the duration of AP frames is larger than that of the client stations (due to aggregation). Hence, if the AP transmits during the slot, then regardless of whether it suffers from a collision the duration of the slot in μs is

$$T_{AP} = T_{phyhdr1} + T_{DN}(R_1) + T_{sifs} + T_{phyhdr} + T_{ack} + T_{difs} \tag{16}$$

where $T_{phyhdr1}$ is the PHY/MAC header duration for an aggregated frame; T_{phyhdr} is the standard PHY header duration; T_{sifs} and T_{difs} are respectively DIFS and SIFS durations. DIFS is the 802.11 distributed coordination function (DCF) interframe space, and SIFS is the short interframe space [2]. T_{ack} is the transmission duration of an ACK frame; $T_{DN}(R_1)$ is the transmission duration of a downlink MAC frame,

$$T_{DN}(R_1) = \lceil (L \times 8 + 6 + 16) / DBPS(R_1) \rceil \times 4 \tag{17}$$

- Type 2 - class 1 transmits: Observe that the duration of class 1 frames is larger than that of class 2 stations. Hence, if a class 1 station wins the transmission opportunity, and its transmission does not collide with a downlink transmission, but might collide with other uplink transmissions, the duration in μs is

$$T_{21} = T_{phyhdr} + T_{UP1}(R_1) + T_{sifs} + T_{phyhdr} + T_{ack} + T_{difs} \tag{18}$$

where $T_{UP1}(R_1)$ is the transmission duration of an uplink MAC frame from a class 1 station, given by

$$T_{UP1}(R_1) = 4 \times \left\lceil \frac{(x_{21}(R_1) + L_{machdr} + L_{FCS}) \times 8 + 22}{DBPS(R_{21})} \right\rceil \tag{19}$$

The probability that a collision occurs among uplink packets (not involving downlink packet) is

$$p_{CUP} = (1 - \tau_0)(1 - (1 - \tau_1)^{n_1}(1 - \tau_2)^{n_2} - n_1\tau_1(1 - \tau_1)^{n_1-1}(1 - \tau_2)^{n_2} - n_2\tau_2(1 - \tau_2)^{n_2-1}(1 - \tau_1)^{n_1}) \quad (20)$$

The probability that a collision occurs only among uplink packets from class 2 stations is

$$p_{CUP2} = (1 - \tau_0)(1 - \tau_1)^{n_1}(1 - (1 - \tau_2)^{n_2} - n_2\tau_2(1 - \tau_2)^{n_2-1}) \quad (21)$$

Hence, the probability that the duration is T_{21} is

$$p_{T_{21}} = n_1\tau_1(1 - \tau_1)^{n_1-1}(1 - \tau_2)^{n_2}(1 - \tau_0) + p_{CUP} - p_{CUP2} \quad (22)$$

- Type 3 - only class 2 transmits: If a class 2 station makes a transmission which does not collide with an AP or class 1 station transmissions, but might collide with uplink packets from other class 2 stations, the duration within a slot time in μs is

$$T_{22} = T_{phyhdr} + T_{UP2}(R_1) + T_{sifs} + T_{phyhdr} + T_{ack} + T_{difs} \quad (23)$$

where $T_{UP2}(R_1)$ is the transmission duration in μs of an uplink MAC frame from a class 2 station,

$$T_{UP2}(R_1) = 4 \times \left\lceil \frac{(x_{22}(R_1) + L_{machdr} + L_{FCS}) \times 8 + 22}{54 \times 4} \right\rceil \quad (24)$$

The probability that the duration is T_{22} is given by

$$p_{T_{22}} = n_2\tau_2(1 - \tau_2)^{n_2-1}(1 - \tau_0)(1 - \tau_1)^{n_1} + p_{CUP2} \quad (25)$$

- Type 4 - idle slot: If no transmission occurs, the duration is a PHY slot σ . This event occurs with probability

$$p_{Idle} = (1 - \tau_1)^{n_1}(1 - \tau_2)^{n_2}(1 - \tau_0) \quad (26)$$

Combining these yields the expected MAC slot duration,

$$E_T = p_{Idle}\sigma + \tau_0 T_{AP} + p_{T_{21}} T_{21} + p_{T_{22}} T_{22} \quad (27)$$

The network throughput is then given by

$$S(R_{21}) = \frac{N_{D1}X_1 + N_{D2}X_2}{E_T} \quad (28)$$

where

$$X_1 = \tau_0(1 - \tau_1)^{n_1}(1 - \tau_2)^{n_2} E_{pld_{11}} + \tau_1(1 - \tau_1)^{n_1-1}(1 - \tau_2)^{n_2}(1 - \tau_0) E_{pld_{21}} \quad (29)$$

$$X_2 = \tau_0(1 - \tau_1)^{n_1}(1 - \tau_2)^{n_2} E_{pld_{12}} + \tau_2(1 - \tau_1)^{n_1}(1 - \tau_2)^{n_2-1}(1 - \tau_0) E_{pld_{22}} \quad (30)$$

We select the downlink PHY rate R_1^* (equal to the uplink PHY rate of class 1 stations) from the set \mathcal{R} of supported 802.11a/g rates so as to maximise this throughput given the channel SNR.

BSC Time-Sharing Coding. For the BSC paradigm we start by considering the simpler time-sharing coding scheme. As in the erasure channel case, MAC frames are constructed by aggregating two portions intended for the different classes of stations with the portion meant for class 1 stations coded (based on BSC crossover probability) and the remainder (for class 2 stations) being uncoded. Note that each portion is in itself composed of sub-frames meant for the different users. We also apply similar coding to protect uplink transmissions of class 1 stations to allow information to be recovered from corrupted uplink frames.

Let x_{11} denote the information payload size for a class 1 station and x_{12} for a class 2 station. Suppose a downlink PHY rate R_1 is chosen. The crossover probability for class 1 stations is $p(R_1)$ and the number of coded bits for x_{11} is $x_{11}/(1 - H(p(R_1)))$. To equalize the downlink throughputs of stations in both classes, we therefore require

$$x_{11}(R_1) = x_{12}(R_1) \tag{31}$$

Given a frame size L , we have that

$$L = N_{D1} \cdot \frac{x_{11}(R_1) + L_{subhdr} + L_{FCS}}{1 - H(p(R_1))} + N_{D2} \cdot (x_{12}(R_1) + L_{subhdr} + L_{FCS}) \tag{32}$$

and hence $x_{11}(R_1)$ and $x_{12}(R_1)$ can be solved. As downlink packets to class 1 stations are erasure-free in the BSC paradigm, the expected payload delivered to a class 1 station is $E_{pld_{11}} = x_{11}(R_1)$.

To equalize the uplink and downlink throughput from/to class 2 stations we require $x_{22}(R_1) = x_{12}(R_1)$. Since erasure-free frames are delivered, $p_{e_{21}} = 0$ and $\tau_0 = \tau_1 = \tau_2$. Thus, to equalize the uplink and downlink throughput from/to class 1 stations we require $x_{21}(R_1) = x_{11}(R_1)$.

The uplink PHY rate of a class 1 station R_{21}^* is selected in terms of maximising its BSC capacity, i.e.

$$R_{21}^* = \arg \max_{r \in \mathcal{R}} r \cdot \left(1 - H(p(r)) \right) \tag{33}$$

The expected duration in a slot time E_T and the network throughput $S(R_1)$ are derived in a similar manner to the erasure case. We select the downlink rate R_1^* so as to maximise the network throughput given the channel SNR.

BSC Superposition Coding. The MAC frames in this setting are constructed in two steps. Once a value of β has been determined, the V and U bit vectors are generated from the aggregated IP packets of each class. These are then added to generate the MAC frame. Despite the scheme being more complicated, the analysis with superposition coding is similar to the BSC time-sharing case. The main difference lies in the calculation of the downlink payload size.

Suppose the downlink PHY rate used by the AP is R_1 , the downlink BSC capacity in bits per channel use between the AP and a class 1 station is

$1 - H(\beta \circ p(R_1))$, and that between the AP and a class 2 station is $H(\beta)$. The MAC frame body is formed by superimposing N_{D2} downlink unicast packets destined to class 2 stations to N_{D1} downlink unicast packets destined to class 1 stations. Let x_{11} denote the information payload size for a class 1 station, and x_{12} the information payload size for a class 2 station. Given the MAC frame size L , we have

$$L = \frac{N_{D1}(x_{11}(R_1) + L_{subhdr} + L_{FCS})}{1 - H(\beta \circ p(R_1))} = \frac{N_{D2}(x_{12}(R_1) + L_{subhdr} + L_{FCS})}{H(\beta)} \quad (34)$$

To equalize the downlink throughputs of stations in both classes, we require

$$x_{11}(R_1) = x_{12}(R_1) \quad (35)$$

Thus we find the relationship $N_{D1}H(\beta) = N_{D2}(1 - H(\beta \circ p(R_1)))$. The ratio N_{D1}/N_{D2} fixes the value of β . With the value of β determined, the downlink unicast payload size for each client station (class 1 or class 2) is then given by

$$\begin{aligned} x(R_1) &= \max\left(\frac{LH(\beta)}{N_{D2}} - L_{subhdr} - L_{FCS}, 0\right) \\ &= \max\left(\frac{L(1 - H(\beta \circ p(R_1)))}{N_{D1}} - L_{subhdr} - L_{FCS}, 0\right) \end{aligned} \quad (36)$$

The uplink PHY rate of class 1 stations R_{21}^* is selected in the same manner as BSC time-sharing case. We also protect uplink transmissions using the ideal code. Again, the PHY rate R_1^* is chosen to maximise system throughput.

4.2 Multicast

In the multicast scenario that we consider, the AP multicasts only two downlink flows, which are aggregated into a large MAC frame. Flow 1 is communicated to n_1 class 1 stations and flow 2 is communicated to n_2 class 2 stations, respectively. There are no competing uplink flows. Therefore, we can compute the throughput using the analysis in Section 4.1 by setting the following parameter values: $N_{D1} = N_{D2} = 1$; $x_{21} = x_{22} = 0$; $p_{e21} = p_{c21} = 0$; $\tau_1 = \tau_2 = 0$; $\tau_0 = 2/(W_0 + 1)$, where W_0 is the minimum contention window size. The expected payload and MAC slot duration can be calculated using the same way as the unicast analysis, but for a multicast network, we consider per-station multicast saturation throughput,

$$S(R_1) = \frac{\tau_0 E_{pld11}(R_1)}{E_T(R_1)} = \frac{\tau_0 E_{pld12}(R_1)}{E_T(R_1)} \quad (37)$$

The optimal PHY rate R_1^* is selected to maximise the per-station throughput.

5 Performance

5.1 Unicast

We compare the throughput performance of the erasure channel and BSC schemes for the protocol parameters detailed in Table 1. Fig. 6 shows the sum-throughputs achieved by the different approaches for a network consisting of 20 client stations, 10 in class 1 and 10 in class 2. This is quite a large number of saturated stations for an 802.11 WLAN and suffers from a high level of collision losses. Comparing with Fig. 1, it can be seen that the throughput is significantly reduced due to the various protocol overheads and collisions that have now been taken into account. Nevertheless, the relative throughput gain of the BSC based approaches compared to the erasure channel approach continues to exceed 50% for a wide range of SNRs. BSC time-sharing coding achieves very similar performance to the more sophisticated superposition coding. Fig. 7 shows the corresponding result for a smaller number of client stations, 5 in class 1 and 5 in class 2. The overall throughput is higher than that with 20 stations because of the lower chance of collisions, and the gain offered by BSC approaches is even higher i.e. more than 75% over a wide range of SNRs.

Table 1. Protocol parameters used in simulations

T_{sifs} (μs)	16	L_{subhdr} (bytes)	16
T_{difs} (μs)	34	L_{FCS} (bytes)	4
Idle slot duration σ (μs)	9	L_{machdr} (bytes)	24
T_{ack} (μs)	24	CW_{min}	16
T_{phyhdr} (μs)	20	CW_{max}	1024
$T_{phyhdr1}$ (μs)	36	Retry limit	7




Fig. 6. Unicast maximum network throughput vs. SNR of class 1 stations, Rayleigh channel, $L = 8000$ bytes, with $n_1 = n_2 = 10$ stations




Fig. 7. Unicast maximum network throughput vs. SNR of class 1 stations, Rayleigh channel, $L = 8000$ bytes, with $n_1 = n_2 = 5$ stations

Fig. 8 illustrates how the number of stations affects these results. The decrease in network throughput with increasing number of stations is evident, as is the significant performance gain offered by the BSC schemes. For smaller numbers of stations (which is perhaps more realistic), the throughput gain offered by the BSC approaches is larger e.g. nearly up to 70% for 2 stations and falling to around 50% with 20 stations. The proportion of class 1 and class 2 stations can be




Fig. 8. Unicast maximum network throughput vs. varying total number of stations for a fixed proportion of class 2 stations $n_1 = n_2$, $SNR = 22$ dB, $L = 8000$ bytes, Rayleigh channel

expected to affect the relative performance of the erasure and BSC schemes. This is because we now have multiple transmitting stations, and each station defers its contention window countdown on detecting transmissions by other stations. Since class 1 transmissions are of longer duration than class 2 transmissions, we expect that the network throughput will rise as the number of class 1 stations falls and indeed we find that this is the case. See, for example, Fig. 9 which plots

the network throughput versus the varying proportion of class 2 stations while maintaining the total number of client stations constant as $n_1 + n_2 = 10$.




Fig. 9. Unicast maximum network throughput vs. varying proportion of class 2 stations for a fixed total number of stations $n_1 + n_2 = 10$, $SNR = 22\text{dB}$, $L = 8000$ bytes, Rayleigh channel




Fig. 10. Multicast per-station maximum throughput vs SNR of class 1 stations, $L = 8000$ bytes, Rayleigh channel, with $n_1 = n_2 = 10$ stations

5.2 Multicast

Fig. 10 shows the per-station multicast throughput for a network with $n_1 = 10$ class 1 stations and $n_2 = 10$ class 2 stations. The throughput is much higher than the unicast case as shown in Fig. 6 because of the absence of collisions among different stations. Nevertheless, both of the BSC schemes (time-sharing and superposition coding) continue to offer substantial performance gains over the erasure channel approach, increasing throughput by almost 100% over a wide range of SNRs. The superposition coding scheme performs slightly better than the time-sharing scheme, but the difference is minor. Fig. 11 shows the corresponding results with a larger MAC frame size of 65536 bytes, which is the maximum frame size allowed in the 802.11n standard. The performance gain




Fig. 11. Multicast per-station maximum throughput vs. SNR of class 1 stations, $L = 65536$ bytes, Rayleigh channel, with $n_1 = n_2 = 10$ stations

offered by the BSC approaches increases as the frame size is increased. Since the per-station multicast throughput is independent of the number of stations, we only show results for one value of n_1 and n_2 .

6 Conclusions

In this paper we consider the potential benefits of viewing the channel provided by an 802.11 WLAN as a binary symmetric channel (BSC), as opposed to a more conventional packet erasure channel. That is, rather than simply discarding corrupted frames we consider viewing a received frame as a binary vector in which an unknown subset of bits have been “flipped”. We present analysis results for both multicast and unicast traffic, taking account of important MAC layer overheads such as collision losses. Importantly, we also consider the multi-user channel aspect of a WLAN i.e. that transmissions are inherently broadcast in nature. We find that increases in network throughput of more than 100% are possible over a wide range of SNRs and that the much simpler time-sharing scheme yields most of these gains. To our knowledge, this is the first detailed analysis of multi-user BSC coding in 802.11 WLANs. We note that these performance gains involve software rather than hardware changes, and so essentially come for “free”.

References

1. Lin, K., Kushman, N., Katabi, D.: ZipTx: Harnessing Partial Packets in 802.11 Networks. In: Proc. of ACM MobiCom 2008, San Francisco, pp. 351–362 (2008)
2. IEEE 802.11n-2009 - Amendment 5: Enhancements for Higher Throughput. IEEE-SA (2009)
3. Li, T., Ni, Q., Malone, D., Leith, D., Xiao, Y., Turletti, T.: Aggregation with Fragment Retransmission for Very High-speed WLANs. IEEE/ACM Trans. on Networking 17(2), 591–604 (2009)

4. Verkaik, P., Agarwal, Y., Gupta, R., Gupta, A.C.: SoftSpeak: Making VoIP Play Well in Existing 802.11 Deployment. In: Proc. of NSDI 2009, Boston, MA (2009)
5. Wang, W., Liew, S.C., Pang, Q., Li, V.O.K.: A Multiplexing-multicast Scheme that Improves System Capacity of Voice-over-IP on Wireless Lan by 100%. In: Proc. of ISCC 2004, Washington, DC, USA, vol. 2, pp. 472–477 (2004)
6. Zheng, Y., Lu, K., Wu, D., Fang, Y.: Performance Analysis of IEEE 802.11 DCF in Binary Symmetric Channels. In: Proc. of GLOBECOM 2005 (2005)
7. Zheng, Y., Lu, K., Wu, D., Fang, Y.: Performance Analysis of IEEE 802.11 DCF in Imperfect Channels. *IEEE Trans. on Vehicular Technology* 55(5), 1648–1656 (2006)
8. Simon, M.K., Alouini, M.S.: *Digital Communication over Fading Channels*, 2nd edn. Wiley-IEEE Press (2005)
9. Pursley, M.B., Taipale, D.J.: Error Propabilities for Spread-spectrum Packet Radio with Convolutional Codes and Viterbi Decoding. *IEEE Trans. on Communications* 35(1), 1–12 (1987)
10. Viterbi, A.J.: Convolutional Codes and Their Performances in Communication Systems. *IEEE Trans. on Communications Technology* 19(5), 751–772 (1971)
11. Cover, T.M., Thomas, J.A.: *Elements of Information Theory*, 2nd edn. John Wiley & Sons, Inc. (2006)
12. Bergmans, P.P., Cover, T.M.: Cooperative Broadcasting. *IEEE Trans. on Information Theory* IT-20(3), 317–324 (1974)
13. MacKay, D.J.C.: *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, New York (2003)
14. Malone, D., Duffy, K., Leith, D.: Modeling the 802.11 Distributed Coordination Function in Nonsaturated Heterogeneous Conditions. *IEEE/ACM Trans. on Networking* 15(1), 159–172 (2007)
15. Dutta, A., Saha, D., Grunwald, D., Sicker, D.: SMACK - A Smart ACKnowledgment Scheme for Broadcast Messages in Wireless Networks. In: Proc. of SIGCOMM 2008 (2008)
16. Bianchi, G.: Performance Analysis of the IEEE 802.11 Distributed Coordination Function. *IEEE Journal on Selected Areas in Communications* 18(3), 535–548 (2000)
17. Ni, Q., Li, T., Turletti, T., Xiao, Y.: Saturation Throughput Analysis of Error-prone 802.11 Wireless Networks. *Wiley Journal of Wireless Communications and Mobile Computing* (2005)