

# Frames in Outdoor 802.11 WLANs Provide a Hybrid Binary-Symmetric/Packet-Erasure Channel

Xiaomin Chen, Douglas Leith,  
Hamilton Institute, NUI Maynooth

**Abstract**—Corrupted frames with CRC errors potentially provide a useful channel through which information can be transmitted. Using measurements taken in an outdoor environment, it is demonstrated that for 802.11 wireless links the channel provided by corrupted frames alone (i.e. ignoring frames with PHY errors and frames received correctly) can be accurately modelled as a binary symmetric channel (BSC) provided appropriate pre- and post- processing is carried out. Also, the channel provided by corrupted frames and other frames combined can be accurately modelled as a hybrid binary-symmetric/packet-erasure channel. Importantly, it is found that this hybrid channel offers capacity increases of more than 100% compared to a conventional packet erasure channel over a wide range of RSSIs. This indicates that the potential exists for significant network throughput gains if the information contained in 802.11 corrupted packets is exploited.

## I. INTRODUCTION

Frames sent over an 802.11 wireless link may be received (i) with a *PHY error*, where the PHY header is corrupted by noise/interference and the receiver cannot demodulate the frame, or (ii) with a *CRC error*, where the PHY header is received correctly and the frame is decoded but then fails a CRC check or (iii) *without error*. The resulting information channel is often modelled as a packet erasure channel. That is, frames are received in an “all or nothing” fashion with frames having PHY or CRC errors being discarded and only frames received without error retained. However, the fraction of incorrect bits in frames with CRC errors can be small. For example, Fig. 1 shows measurements of the fraction of corrupted bits in frames with CRC errors on an 802.11 link. It can be seen that even when the frame error rate (FER) is high (91.98% of frames fail the PHY header check or the CRC check), most of the frames received with CRC errors have less than 10% of bits incorrect. Thus, these corrupted frames potentially provide a useful channel through which we can transmit information.

Taking this observation as our starting point, in this paper our aim is to characterise the information channel provided by 802.11 frame transmissions. Using measurements taken in an outdoor environment, we demonstrate that the channel provided by corrupted frames alone (i.e. ignoring frames with PHY errors and frames received correctly) can be accurately modelled as a BSC provided appropriate pre- and post- processing is carried out. Also, the channel provided by corrupted frames and other frames combined can be accurately modelled as a hybrid binary-symmetric/packet-erasure channel. We calculate the capacity of this channel and show




Fig. 1. Measured fraction of incorrect bits vs frame sequence number. Outdoor measurements, 802.11g PHY rate 54Mbps.  $FER = 91.98\%$ .

that the potential exists for significant throughput gains. This complements and extends recent work [6], [7], [8] which study the gains achievable in 802.11 by use of a refined erasure model (partitioning frames into blocks and adding checksums).

## II. PRELIMINARIES

### A. Experimental setup

Experimental data was collected in an open outdoor space with no other interferers present. The FER was adjusted by varying the distance between sender and receiver. Care was taken to ensure repeatability of results – measurements were taken on an open space (a large playing field), sender and receiver were positioned at fixed heights, antenna orientations were held fixed, human operators left the vicinity of the experiment during measurements.

1) *Hardware and software*: An Asus Eee PC 4G Surf equipped with Atheros AR5BXB63 802.11b/g chipsets (AR2425, MAC 14.2, RF5424, PHY 7.0) was used as the access point, running FreeBSD 8.0 with the RELEASE kernel and using the standard FreeBSD ATH driver. A Fujitsu E series Lifebook equipped with a Netgear dual band 802.11a/b/g wireless PC card WAG511 using Atheros AR5212 chipset was used as a client station, running Ubuntu 11.04 and using a modified Linux Madwifi driver. Omni antennas used.

We disabled the Atheros’ Ambient Noise Immunity feature which has been reported to cause unwanted side effects. Transmission power of the laptops was fixed and antenna diversity disabled. In previous work we have taken considerable care to confirm that with this hardware/software setup the wireless stations accurately follow the IEEE 802.11 standard, see [1].

A perl script was used to generate CBR UDP traffic. The content of each UDP packet was a random binary vector. Unless stated otherwise, the UDP payload is 8000 bits, and

the inter-packet interval is 20 ms. Packets were transmitted over the WLAN in the broadcast mode and hence there were no MAC level ACKs or retransmissions. The wireless driver at the receiver was modified to record the receiving status for frames with PHY errors, CRC errors or without errors and to transfer the contents of these from the kernel to user space via the high-speed data relay filesystem (relayfs).

2) *Recovering sequence number for corrupted frames:* Frames received with CRC errors were compared against the corresponding original content in order to determine which specific bits inside the frame were received corrupted. Since the frame header might also be corrupted, the following fitting and pattern matching procedure was used to indirectly recover the frame sequence number. For each corrupted frame we first searched the set of correctly received frames to find the packet received without error closest in time to the corrupted frame. Since the frames are transmitted at a fixed rate, the interval between two consecutive frame timestamps is roughly constant, but to correct for clock skew between transmitter and receiver it is necessary to estimate the relative clock rate and offset. We estimated these using linear least squares fitting of the received timestamps of neighbouring error-free frames. The timestamp of the corrupted frame was then used to determine likely candidates for this frame amongst the transmitted frames. From this set of candidates, the sequence number and payload of each was compared with that of the corrupted frame in order to identify the transmitted frame most likely to correspond to the received corrupted frame.

### B. Runs test


Our statistical analysis makes use of the *runs test* (also called the Wald-Wolfowitz test) [5]. The runs test is a non-parametric test to check the null hypothesis that the elements in a two-valued sequence are independent and identical distributed. Given a 0-1 sequence, a *run* is consecutive entries having the same value *e.g.* in the sequence 1100110111 there is one 0 run, one 00 run, two 11 runs and one 111 run. Under the null hypothesis, the number of runs is a random variable whose conditional distribution is approximately normal with mean  $\mu = \frac{2 \cdot N_1 \cdot N_0}{N} + 1$  and variance  $\sigma^2 = \frac{(\mu-1)(\mu-2)}{N-1}$  where  $N_1$  is the number of 1 values in the sequence,  $N_0$  the number of 0 values and  $N = N_1 + N_0$ . Unless otherwise stated we carry out statistical testing at the 5% significance level.

## III. CHANNEL MODELLING


We proceed by first investigating the channel provided by corrupted frames alone *i.e.* ignoring frames with PHY errors and frames received correctly. We find that, to within statistical error, this can be accurately modelled as a BSC. We then consider the channel provided by corrupted frames and other frames combined. We find that this can be accurately modelled as a hybrid binary-symmetric/packet-erasure channel.

### A. Channel provided by corrupted frames

A binary symmetric channel (BSC) takes binary input  $X \in \{0, 1\}$  and maps this to binary output  $Y \in \{0, 1\}$ . With



(a) Before interleaving



(b) After interleaving

Fig. 2. Per-bit error frequency pattern across a frame before and after interleaving, outdoor,  $FER = 0.5658$ , PHY rate 54Mbps, 5000 frames.

probability  $1 - p$  the channel transmits the input bit correctly and with crossover probability  $p$  the input bit is flipped. That is,  $\mathbb{P}\{Y = 1|X = 1\} = 1 - p = \mathbb{P}\{Y = 0|X = 0\}$  and  $\mathbb{P}\{Y = 0|X = 1\} = p = \mathbb{P}\{Y = 1|X = 0\}$ . Repeated binary channel uses are independent and identically distributed. That is, if a binary vector is transmitted through a BSC, each bit is flipped independently and with identical crossover probability  $p$ . Therefore, to establish that a channel taking binary inputs is a BSC, we need to show

- (1) Repeated binary channel uses are independent and identically distributed (i.i.d.).
- (2) The probability that a 1 is flipped to a 0 after transmission is the same as the probability that a 0 is flipped to a 1 *i.e.* the binary channel is symmetric.

1) *Binary channel uses are i.i.d.:* We begin by presenting raw experimental measurements in Fig. 2(a). This figure plots representative measurements of the bit error frequency for each bit position within a corrupted frame. It can be seen that bit errors are not evenly distributed and the bit error frequency sequence is periodic across a frame. This observation is not new, *e.g.* see [2], [3], [4], and clearly violates the independence requirement of a BSC. Nevertheless, Fig. 2(b) plots the bit error frequency for the same data after interleaving<sup>1</sup>. Interleaving can be readily implemented at the MAC layer – bits in a frame are permuted at the transmitter before a MAC frame goes down to the PHY layer and when the frame is received the inverse permutation is used to recover the original bit order. It can be seen from Fig. 2(b) that after interleaving the periodicity of bit errors appears to be removed, although further analysis is required to confirm this.

To analyse the independence of repeated channel uses within each individual interleaved frame we use the runs test. For each corrupted frame we construct a 0-1 sequence by labelling corrupted bits as 1's and correct bits as 0's. We find that the runs test cannot reject the null hypothesis at the 5% significance level that after interleaving bit errors inside a

<sup>1</sup> That is, we randomly permute the bits within each frame. While the 802.11 PHY already carries out interleaving, this is carried out over small blocks of bits equal in size to an OFDM symbol, *e.g.* 288 bits with 64-QAM, whereas we interleave over a complete frame *i.e.* 8000 bits when the frame size is 1000B. Note that interleaving does not maximise capacity.




Fig. 3. Per-frame bit crossover probability for a sequence of 10000 packets, PHY rate 54Mbps,  $FER = 0.0423$ .

frame are independent. That is, to within statistical error we can conclude that repeated channel uses inside each interleaved frame are independent and Bernoulli distributed with the bit crossover probability  $p = n_c/l$ , where  $n_c$  is the number of corrupted bits within a frame and  $l$  is the frame length.

Next, to analyse the independence of channel uses across multiple frames we concatenate the foregoing binary sequences for successive corrupted frames and apply the runs test. For example, Fig. 3 plots the measured per-frame bit crossover probability for a sequence of 10,000 frames. Consecutive corrupted frames that pass the runs test are labelled using the same marker. It can be seen that the experimental run is partitioned into three segments. Within each segment, corrupted frames form a bit-level channel over which repeated channel uses can be considered to be independent and Bernoulli distributed to within statistical error. The first segment spans 3865 frames (around 77.3s given that the inter-packet interval is 20ms) and the third segment spans 6118 frames (about 122.36s). The second segment has only one frame with an unusually high crossover probability. Neglecting this single bad frame out of 10,000 frames, the measurement results indicate that for long periods the sequence of corrupted frames form a bit-level channel over which repeated channel uses can be considered to be i.i.d. The data in Fig. 3 is for a FER of 0.0423 and PHY rate of 54Mbps. Fig. 4 plots the mean of segment durations over an experimental run of 10000 packets for a range of FERs and PHY rates. It can be seen that the segment duration tends to decrease as the FER increases *i.e.* the duration within which corrupted frames have an i.i.d. bit crossover probability becomes shorter as the FER increases. For FER's less than 20%, at all PHY rates the mean segment duration exceeds 10s. That is, for FER's less than 20% the binary channel provided by corrupted frames is i.i.d. for periods exceeding, on average, 10s. Such time-scales seem sufficient for most channel modelling purposes.

2) *Binary channel is symmetric*: Table I reports the measured bit flip rates for 1 values and 0 values. Results are shown for a range of FER values and PHY rates. It can be seen that the bit flip probabilities for both 1's and 0's are close to each other, and thus, to within experimental error, can be approximately considered as "symmetric".

In summary, after interleaving the bit errors in corrupted frames are, to within statistical error, independent and identically distributed with a symmetric crossover probability, *i.e.* the information channel can be accurately modelled as a BSC.




Fig. 4. Mean segment duration versus FER at different PHY rates.




Fig. 5. Hybrid BSC/packet erasure channel model.

### B. Hybrid binary-symmetric/packet-erasure channel

Our hypothesis is that the channel provided by 802.11 frames can be accurately modelled as a mixed packet erasure/binary symmetric channel. Formally, the channel takes an  $n$ -bit binary vector  $\mathbf{x} \in \{0, 1\}^n$  as input, and outputs received vector  $\mathbf{y} \in \{0, 1\}^n$ . The vector  $\mathbf{y}$  can be received with three possible states: (i)  $\mathbf{0}$  (erased), (ii)  $\mathbf{x}'$  (corrupted) and (iii)  $\mathbf{x}$  (without error). Repeated channel uses are i.i.d. The probability that a received vector is erased is  $\mathbb{P}\{\mathbf{y} = \mathbf{0}\} = r$ . The probability that a non-erased vector is received without error is  $\mathbb{P}\{\mathbf{y} = \mathbf{x} \mid \mathbf{y} \neq \mathbf{0}\} = s$  and so the probability that a vector is correctly received is  $\mathbb{P}\{\mathbf{y} = \mathbf{x}\} = (1-r)s$ . The probability that a vector is corrupted is  $\mathbb{P}\{\mathbf{y} = \mathbf{x}'\} = (1-r)(1-s)$ . In a corrupted vector bits are flipped symmetrically and independently with crossover probability  $p$ . This channel is illustrated schematically in Fig. 5.

We can relate this hybrid model to 802.11 by associating the input vectors with transmitted frames, erasures with PHY errors and corrupted vectors with CRC errors. To establish an equivalence we need to show that (i) frames are received with PHY errors, CRC errors and without error in an i.i.d. fashion, (ii) corrupted frames provide a BSC. We have already established (ii) in Section III-A but it remains to establish (i).

Our hypothesis is that PHY errors, CRC errors and frames received without errors are mutually independent across time. To investigate this hypothesis, we again use the runs test. We construct a 0-1 sequence by labelling frames received with PHY errors as a 1 and other frames as a 0, and then apply the runs test for each individual segment which passes the runs test for bit error independence check. Similarly, we do the same respectively for frames received with CRC errors and without errors. Fig. 6 plots the fraction of time in an experimental run of 10,000 frames within which the runs tests pass for a range of FERs and PHY rates. The runs tests pass over at least 80% of the time, *i.e.* in a run of 10,000 frames over 80% of the time, to within statistical error, frames are received with PHY errors, CRC errors and without error in an i.i.d. fashion.

PHY rate	FER	Flip rate for 1's		Flip rate for 0's		PHY rate	FER	Flip rate for 1's		Flip rate for 0's	
		$\mu_1$	$\sigma_1/\sqrt{N_1}$	$\mu_0$	$\sigma_0/\sqrt{N_0}$			$\mu_1$	$\sigma_1/\sqrt{N_1}$	$\mu_0$	$\sigma_0/\sqrt{N_0}$
54Mbps	0.0835	0.0018	$2.34 \times 10^{-5}$	0.0018	$2.26 \times 10^{-5}$	36Mbps	0.0649	0.0014	$2.72 \times 10^{-5}$	0.0013	$2.48 \times 10^{-5}$
	0.0984	0.0019	$2.40 \times 10^{-5}$	0.0020	$2.34 \times 10^{-5}$		0.1673	0.0019	$1.69 \times 10^{-5}$	0.0020	$1.65 \times 10^{-5}$
	0.1540	0.0023	$1.97 \times 10^{-5}$	0.0023	$1.88 \times 10^{-5}$		0.3314	0.0023	$1.31 \times 10^{-5}$	0.0023	$1.22 \times 10^{-5}$
	0.2932	0.0031	$1.67 \times 10^{-5}$	0.0032	$1.62 \times 10^{-5}$		0.4427	0.0032	$1.34 \times 10^{-5}$	0.0033	$1.30 \times 10^{-5}$
	0.4584	0.0042	$2.19 \times 10^{-5}$	0.0041	$2.04 \times 10^{-5}$		0.5416	0.0037	$1.19 \times 10^{-5}$	0.0035	$1.10 \times 10^{-5}$
	0.5658	0.0067	$2.45 \times 10^{-5}$	0.0065	$2.31 \times 10^{-5}$		0.6697	0.0052	$1.43 \times 10^{-5}$	0.0054	$1.38 \times 10^{-5}$
	0.7492	0.0066	$1.52 \times 10^{-5}$	0.0068	$1.48 \times 10^{-5}$		0.7214	0.0052	$2.27 \times 10^{-5}$	0.0048	$2.08 \times 10^{-5}$
	0.8881	0.0120	$1.92 \times 10^{-5}$	0.0121	$1.83 \times 10^{-5}$		0.9247	0.0122	$3.41 \times 10^{-5}$	0.0115	$3.15 \times 10^{-5}$
48Mbps	0.0547	0.0015	$3.21 \times 10^{-5}$	0.0017	$3.16 \times 10^{-5}$	24Mbps	0.0617	0.0048	$7.66 \times 10^{-5}$	0.0044	$6.99 \times 10^{-5}$
	0.0949	0.0021	$2.91 \times 10^{-5}$	0.0020	$2.71 \times 10^{-5}$		0.3219	0.0041	$1.84 \times 10^{-5}$	0.0039	$1.73 \times 10^{-5}$
	0.1883	0.0021	$1.68 \times 10^{-5}$	0.0020	$1.59 \times 10^{-5}$		0.5220	0.0098	$2.68 \times 10^{-5}$	0.0094	$2.50 \times 10^{-5}$
	0.4319	0.0034	$1.43 \times 10^{-5}$	0.0034	$1.35 \times 10^{-5}$		0.8229	0.0306	$3.19 \times 10^{-5}$	0.0293	$2.97 \times 10^{-5}$
	0.5818	0.0036	$1.81 \times 10^{-5}$	0.0035	$1.70 \times 10^{-5}$	18Mbps	0.0449	0.0045	$1.44 \times 10^{-5}$	0.0046	$1.38 \times 10^{-5}$
	0.6185	0.0044	$1.38 \times 10^{-5}$	0.0046	$1.32 \times 10^{-5}$		0.3481	0.0049	$2.15 \times 10^{-5}$	0.0050	$2.07 \times 10^{-5}$
	0.7586	0.0045	$1.76 \times 10^{-5}$	0.0043	$1.64 \times 10^{-5}$		0.6415	0.0076	$1.99 \times 10^{-5}$	0.0075	$1.89 \times 10^{-5}$
	0.8522	0.0080	$2.30 \times 10^{-5}$	0.0077	$2.15 \times 10^{-5}$		0.8376	0.0091	$1.93 \times 10^{-5}$	0.0091	$1.84 \times 10^{-5}$
	0.9091	0.0146	$2.09 \times 10^{-5}$	0.0145	$1.98 \times 10^{-5}$		0.9776	0.0129	$3.79 \times 10^{-5}$	0.0123	$3.53 \times 10^{-5}$

TABLE I

BIT FLIP RATES FOR 1'S AND 0'S,  $\mu_i$  THE MEAN FLIP RATE OF BIT  $i$ ,  $\sigma_i/\sqrt{N_i}$  THE STANDARD DEVIATION OF FLIP RATE OF BIT  $i$ ,  $N_i$  THE TOTAL NUMBER OF BIT  $i$  IN CORRUPTED FRAMES,  $N = N_0 + N_1$ .




Fig. 6. Fraction of time passing the runs test for PHY errors, CRC errors and good frames.

#### IV. CONCLUSIONS

The capacity of the hybrid channel is

$$C = R(1-r) (s + (1-s)(1-H(p))) \quad (1)$$

where  $r$  is the probability of a PHY error,  $s$  is the probability that a non-erased packet is error free,  $p$  is the crossover probability in corrupted packets,  $H(p) = -p \log_2(p) - (1-p) \log_2(1-p)$  and  $R$  is the PHY bit rate. Using our experimental data, we binned frames according to their RSSI (Received Signal Strength Indication) and calculated the measured crossover probability in corrupted frames and the frame error rate. As we could not measure the RSSI for PHY erasures using our hardware, we do not include these here. Fig. 7 shows the resulting measured capacity vs RSSI. For comparison we also plot the experimentally measured packet erasure channel capacity *i.e.* the capacity when corrupted frames are discarded. It is important to note that these measured curves cannot be directly compared with theoretical calculations since the mapping between RSSI and SNR is not well defined. It can be seen that the capacity of the hybrid channel is strictly greater than the erasure channel capacity, as expected. More interestingly, it can be seen that the hybrid channel offers capacity increases of more than 100% over a wide range of RSSIs. This indicates that the potential exists for significant




Fig. 7. Outdoor experimental channel capacity.

network throughput gains if the information contained in corrupted packets is exploited. Since 802.11n includes the 802.11a/g modulation and coding schemes, our conclusions apply directly. We do, however, expect that indoor links will behave differently from outdoor links due to multipath effects and temporal variations in the environment, and this will be the subject of future work.

#### REFERENCES

- [1] D. Giustiniano, D. Malone, D. J. Leith, and K. Papagiannaki, "Measuring transmission opportunities in 802.11 links", IEEE/ACM Trans Networking, vol. 18(5), pp. 1516 C1529, Oct. 2010.
- [2] B. Han, L. Ji, S. Lee, B. Bhattacharjee, R. R. Miller, "All Bits Are Not Equal - A Study of IEEE 802.11 Communication Bit Errors". Proc. IEEE INFOCOM 2009.
- [3] P. Fuxjager, F. Ricciato, "Collecting Broken Frames: Error Statistics in IEEE 802.11b/g Links". WIOPT 2008.
- [4] A. Willig, M. Kubisch, C. Hoene, A. Wolisz, "Measurements of a Wireless Link in an Industrial Environment Using an IEEE 802.11-Compliant Physical layer". IEEE Trans Ind. Electr., 49(6), 2002.
- [5] A. M. Mood, "The distribution theory of runs", The Annals of Mathematical Statistics, Vol. 11(4), pp. 367-392, 1940.
- [6] K. Jamieson, H. Balakrishnan, "PPR: partial packet recovery for wireless networks", Proc. SIGCOMM, pp409-420, 2007
- [7] K. Ching-Ju, N. Kushman, D. Katabi, "ZipTx: Harnessing Partial Packets in 802.11 Networks", Proc Mobicom, pp351-362, 2008
- [8] B. Han, A. Schulman, F. Gringoli, et al., "Maranello: practical partial packet recovery for 802.11", Proc NSDI, 2010