

Measuring Pulsed Interference in 802.11 Links

Brad W. Zarikoff, Douglas J. Leith
Hamilton Institute, NUI Maynooth

Abstract—Wireless 802.11 links operate in unlicensed spectrum and so must accommodate other unlicensed transmitters which generate pulsed interference. We propose a new approach for detecting the presence of pulsed interference affecting 802.11 links, and for estimating temporal statistics of this interference. This approach builds on recent work on distinguishing collision losses from noise losses in 802.11 links. When the intervals between interference pulses are i.i.d., the approach is not confined to estimating the mean and variance of these intervals but can recover the complete probability distribution. The approach is a transmitter-side technique that provides per-link information and is compatible with standard hardware. We demonstrate the effectiveness of the proposed approach using extensive experimental measurements. In addition to applications to monitoring, management and diagnostics, the fundamental information provided by our approach can potentially be used to adapt the frame durations used in a network so as to increase capacity in the presence of pulsed interference.

I. INTRODUCTION

Wireless 802.11 links operate in unlicensed spectrum and so must accommodate other unlicensed transmitters. These transmitters include not only other 802.11 WLANs but also Bluetooth devices, Zigbee devices, domestic appliances *etc.* Importantly, the resulting interference is often pulsed in nature. That is, the interference that consists of a sequence of “on” periods (or pulses) during which the interference power is high, interspersed by “off” periods where the interference power is lower, illustrated schematically in Fig. 1. The former might be thought of as corresponding to a packet transmission by a hidden terminal and the latter as the idle times between these transmissions. For this type of interferer, RSSI/SINR measurements are of limited assistance since the SINR measured for one packet may bear little relation to the SINR experienced by other packets. A further complicating factor is that in 802.11 links frame loss due to collisions is a feature of normal operation in 802.11 WLANs, and thus we need to be careful to distinguish losses due to collisions and losses due to channel impairment.

In this paper we propose a new approach for detecting the presence of pulsed interference affecting 802.11 links and for estimating temporal statistics of this interference under mild assumptions. Our approach is a transmitter-side technique that provides per-link information and is compatible with standard hardware. This significantly extends recent work in [1], [2] which establishes a MAC/PHY cross-layer technique capable of classifying lost transmission opportunities into noise-related losses, collision induced losses, hidden-node losses and of distinguishing these losses from the unfairness caused by exposed nodes and capture effects.

Supported by Science Foundation Ireland grants 07/IN.1/1901 and 08/SR-C/11403.




Fig. 1. Illustrating a WLAN with interfering pulsed transmitter (e.g. 802.11 hidden terminal, Bluetooth device, microwave oven, baby monitor, *etc.*) inducing packet loss.

Detection and measurement of pulsed interference is particularly topical in view of the trend towards increasingly dense wireless deployments. In addition to being of interest in their own right for network monitoring, management and diagnostics, our temporal statistic measurements can be used to adapt network parameters so as to significantly increase network capacity in the presence of pulsed interference. This is illustrated in Fig. 2, which shows experimental measurements of packet error rate (PER) versus modulation and coding scheme (MCS) for an 802.11 network in the presence of a pulsed microwave oven (MWO) interferer. Two curves are shown, one for each fragment of a two packet TXOP burst (below we discuss in more detail our interest in using packet pairs). Observe that the PER is lowest at a PHY rate of 18-24 Mbps – importantly, the PER rises not only for higher PHY rates, as is to be expected due to the lower resilience to noise at higher rates, but also rises for *lower* PHY rates. The increase in PER at lower PHY rates is due to the pulsed nature of the interference – since the frame size in our experiment is fixed, the time taken to transmit a frame increases as the PHY rate is lowered, increasing the likelihood that a frame “collides” with an interference burst. At a PHY rate of 1Mbps, the frame duration is longer than the maximum interval between interference pulses and, as a result, the PER is close to 100%. We discuss this example in more detail in Section IV-B, but it is clear the appropriate choice of PHY rate can lead to significant throughput gains in such situations. We briefly note that this type of MAC layer adaptation complements proposed PHY layer interference avoidance techniques such as cognitive radio [3].

II. RELATED WORK

Previous work on estimating 802.11 channel conditions can be classified into three categories. First, *PHY link-level* approaches using SINR and bit-error rate (BER). Second, *MAC approaches* relying on throughput and delay statistics, or frame loss statistics derived from transmitted frames which are not ACKed and/or from signaling messages. Finally *cross-layer MAC/PHY approaches* that combine information at both MAC and PHY layers.

Most work on PHY layer approaches is based on SINR measurements, e.g. [4]–[6]. The basic idea is to *a priori* map SINR measures into link quality estimates. However, it is well known that the correlation between SINR and actual packet delivery rate can be weak due to time-varying channel conditions [7], pulsed interference being one such example of a time-varying channel. [8] considers loss diagnosis by examining the error pattern within a physical-layer symbol, with the aim of exposing statistical differences between collision and weak signal based losses, but does not consider pulsed interference. The cognitive radio literature considers PHY layer techniques for optimising performance in the presence of interference via joint spectral and temporal analysis [9]. There are some solutions tailored to the ISM band [3], where customised hardware has been devised with the aim of providing a synchronisation signal based on periodic interference. However, cognitive radio techniques are largely geared towards interference avoidance and make use of non-standard hardware.

MAC approaches make up some of the most popular and earliest rate control algorithms. Techniques such as ARF [10], RBAR [11] and RRAA [12] attempt to use frame transmission successes and failures as a means to indirectly measure channel conditions. However, these techniques cannot distinguish between noise, collision, or hidden noise sources of error. In [13], rate control via loss differentiation is suggested via a modified ARF algorithm; it was shown to greatly improve performance via the inclusion of a NAK signal, but this requires a modification to the 802.11 MAC. Use of RTS/CTS signals has been proposed for distinguishing collisions from channel noise losses, e.g. [14], [15]. However, such approaches can perform poorly in the presence of pulsed interference such as hidden terminals [1].

With regard to combined MAC/PHY approaches, the present paper builds upon the packet pair approach proposed in [1], [2] for estimating the frame error rates due to collisions, noise and hidden terminals. See also the closely related work in [16]. [1], [2], [16] focus on time-invariant channels and do not consider estimation of temporal statistics. [17] considers a similar problem to [1], but uses channel busy/idle time information.

Some work has been done on packet length adaptation as a means of exploiting a time-varying channel. [18] modifies the Gilbert-Elliott channel model to model bursty channels; however, they do not consider the MAC layer. There are many examples that use MAC frame error information [19]–[23], but they lack the ability to distinguish between noise and collisions. There has been some recent interesting work on a cross-layer model for packet length adaptation in [24], which




Fig. 2. Experimental measurements of packet error rate (PER) versus modulation and coding scheme (MCS) for an 802.11 network operating on channel 9 and physically located near an operational microwave oven (MWO). See Section IV-B for further details of the experimental setup. Two curves are shown, one for each fragment of a two packet TXOP burst. Observe that the PER is minimised around 18–24 Mbps and rises at both lower and higher MCS rates due to the pulsed nature of the interference.

relies on separation between noise errors and collision errors as a means of tuning the packet length and optimising throughput.

III. PULSED INTERFERENCE TEMPORAL STATISTICS: NON-PARAMETRIC ESTIMATION

A. Basic Idea

We start with the observation that packet transmissions over a time-varying wireless link can be thought of as sampling the channel conditions. Each sample covers an extended interval of time, equal to the duration T_D of the packet transmission, see Fig. 3. On a channel with pulsed interference, the frequency with which packet transmissions overlap with interference pulses (and so the level of packet loss) depends on the duration of the packet transmissions relative to the intervals between pulses, and on the durations of the pulses. For example, it is easy to see that when the packet duration T_D is larger than the maximum time between interference pulses, then every packet transmission overlaps with at least one interference pulse and we can expect to observe a high rate of packet loss. Conversely, when the packet duration T_D is much smaller than the time between interference pulses, most of the packet transmissions will not encounter an interference pulse and we can expect a much lower rate of packet loss. Hence, by varying the packet transmit duration and observing the corresponding change in packet loss rate, we can hope to infer information about the timing of the interference pulses. We can make this intuitive insight more precise as follows. Assume that the intervals between pulses are i.i.d. so that they are characterised by a probability distribution function. Then, we will shortly show that the information contained in such packet loss information is sufficient to fully reconstruct this distribution function. This, somewhat surprising, result has important practical implications. Namely, that even when the interference pulses are not directly observable (which we expect to usually be the case), we are nevertheless still able to reconstruct key temporal statistics of the interference process from easily measured packet loss statistics.




Fig. 3. Schematic illustrating “sampling” of a time-varying channel by data packet transmissions. Since the data transmissions occupy an interval of time, the sampling is of the channel conditions over that interval, rather than at a single point in time. As the duration of the data transmissions increases, the chance that a data transmission overlaps with an interference pulse also tends to increase.

B. Mathematical Analysis

We now formalise these claims. Consider a sequence of interference pulses indexed by $k = 0, 1, 2, \dots$ and let T_k denote the start time of the k^{th} interference pulse with $T_0 = 0$, $S_k > 0$ denote the duration of the k^{th} pulse and $\Delta_k = T_{k+1} - (T_k + S_k) > 0$ be the interval between the end of k^{th} pulse and the start of the $(k + 1)^{\text{th}}$ pulse. Defining state vector $X_t := (t, T_{k(t)}, S_{k(t)}, \Delta_{k(t)})$, $t \in \mathbb{R}^+$, the sequence $\{X_t\}$ forms a stochastic process with $T_{k+1} = T_k + S_k + \Delta_k$, $T_0 = 0$, $k(t) = \sup\{k : T_k < t\}$. We assume that the random variables $\Delta_k, k = 1, 2, \dots$ are i.i.d. with finite mean. Then $\Delta_k \stackrel{d}{=} \Delta$, where $\stackrel{d}{=}$ denotes equality in distribution, and let $\text{Prob}[\Delta \leq x] = F(x)$. Similarly, we assume that the pulse durations $\{S_k\}$ are i.i.d. with finite mean and $S_k \stackrel{d}{=} S$.

Pick a sampling interval $[t - T_D, t]$. This sampling interval can be thought of as a packet transmission ending at time t . Define indicator function $U_{T_D}(X_t) = 1$ if interval $[t - T_D, t]$ does not overlap with any interference pulse, and $U_{T_D}(X_t) = 0$ otherwise. That is,

$$U_{T_D}(X_t) = \begin{cases} 1 & t \in [T_k + S_k + T_D, T_{k+1}) \text{ for some } k \\ 0 & \text{otherwise} \end{cases}. \quad (1)$$

Suppose we transmit a sequence of packets and let $\{t_j\}$ denote the sequence of times when transmissions finish. Assume for the moment that (i) a packet is lost whenever it overlaps with an interference pulse and (ii) the intervals between packet transmissions are exponentially randomly distributed and are independent of the interference process. We will shortly relax these assumptions. By assumption (i), $U_{T_D}(X_{t_j})$ equals 1 if the packet transmitted at time t_j is received successfully and 0 otherwise. Hence, the empirical estimate of the packet loss rate is

$$\hat{P}_t(T_D) = 1 - \frac{1}{N(t)} \sum_{j=1}^{N(t)} U_{T_D}(X_{t_j}), \quad (2)$$

where $N(t)$ is the number of packets transmitted in interval $[0, t]$. Provided the packet duration T_D is sufficiently small relative to the mean time between packets, by assumption (ii) the transmit times $\{t_j\}$ effectively possess the Lack of Anticipation property (the number of packet transmissions in any interval $[t, t + u]$, $u \geq 0$, is independent of $\{X_s\}$, $s \leq t$ [25]). When this property holds, by [25, Theorem 1] we almost surely have

$$\lim_{t \rightarrow \infty} \hat{P}_t(T_D) = \lim_{t \rightarrow \infty} P_t(T_D) =: p(T_D)$$

where

$$P_t(T_D) = 1 - \frac{1}{t} \int_0^t U_{T_D}(X_s) ds.$$

That is, the packet loss rate estimator (2) provides an asymptotically unbiased estimate of the mean value of U_{T_D} .

Assumption (i) can be replaced by the weaker requirement that the packet loss rate is higher when a packet transmission overlaps with an interference pulse than when it does not. We consider this in more detail later, in Section V. Assumption (ii) can be relaxed to any sampling approach that satisfies the Arrivals See Time Averages (ASTA) property, see for example [26], [27].

It remains to show that statistic $p(T_D)$ contains useful information about the interference process. We begin by observing that $Y_t = \sup\{k : T_k \leq t\}$ is a renewal process – since the Δ_k and S_k are i.i.d., the start times $\{T_k\}$ of the interference pulses are renewal times. The mean time between renewals is $\mathbb{E}[S + \Delta]$. On each renewal interval $t \in [T_k, T_{k+1})$ we have that $U_{T_D}(X_t) = 1$ for duration $[\Delta_k - T_D]^+$, where $[x]^+$ equals x when $x \geq 0$ and 0 otherwise. The mean value of $U_{T_D}(X_t)$ over a renewal interval is therefore $\int_{T_D}^{\infty} (x - T_D) dF(x)$ and, by the strong law of large numbers,

$$p(T_D) = 1 - \frac{1}{\mathbb{E}[S + \Delta]} \int_{T_D}^{\infty} (x - T_D) dF(x).$$

Since $F(\bullet)$ is a distribution function it is differentiable almost everywhere, and thus so is $p(\bullet)$. At every point T_D where $p(\bullet)$ is differentiable we have

$$\begin{aligned} \frac{dp}{dT_D}(T_D) &= \frac{1}{\mathbb{E}[S + \Delta]} \int_{T_D}^{\infty} dF(x) \\ &= \frac{1}{\mathbb{E}[S + \Delta]} \text{Prob}[\Delta > T_D]. \end{aligned}$$

Provided $p(\bullet)$ is differentiable at $T_D = 0$, then

$$\mathbb{E}[S + \Delta] = \frac{1}{dp(0)/dT_D}$$

since $\text{Prob}[\Delta > 0] = 1$, and so

$$\text{Prob}[\Delta > T_D] = \frac{1}{dp(0)/dT_D} \frac{dp}{dT_D}(T_D). \quad (3)$$

Hence, knowledge of statistic $p(T_D)$ as a function of T_D is sufficient to allow us to calculate not only the mean time between interference pulses $\mathbb{E}[S + \Delta]$, but also the entire distribution function $F(x) = 1 - \text{Prob}[\Delta > x]$ of the interference pulse inter-arrival times.

Note that while we can formally differentiate $p(T_D)$, its estimate $\hat{p}(T_D)$ will be noisy and so differentiating $\hat{p}(T_D)$ is




Fig. 4. Theory and simulation for periodic and Poisson interference. Packet transmissions are Poisson with mean rate $\lambda = 0.01$.

not advisable. The formal differentiation step is merely used to gain insight into the statistical information contained within $p(T_D)$ and there is no need to actually differentiate $\hat{p}(T_D)$ in order to infer characteristics of the interference process (*e.g.* see the examples in the next section).

C. Two Simple Examples

We present two simple examples illustrating the use of statistic $p(T_D)$ and for which explicit calculations are straightforward.

1) *Periodic impulses*: The first example is where the interference consists of periodic impulses with period T_Δ (so $\text{Prob}(\Delta = T_\Delta) = 1$) and packets are always lost when they overlap with an interference pulse. In this case,

$$p(T_D) = 1 - \frac{1}{\mathbb{E}[S + \Delta]} \int_{T_D}^{\infty} (x - T_D) dF(x) \\ = \begin{cases} \frac{T_D}{T_\Delta} & T_D \leq T_\Delta \\ 1 & T_D > T_\Delta \end{cases}.$$

That is, $p(T_D)$ is a truncated line with slope T_Δ . Fig. 4(a) plots this theory line, along with the measured packet loss rate obtained from simulations. The interference period T_Δ can be directly estimated from the slope of the measured line of packet loss versus T_D . The ccdf $1 - F(T_D)$ shown in Fig. 4(b) can be calculated using (3) or deduced based on the interference period.

2) *Poisson interference*: The second simple example is where the interference pulses are Poisson impulses, with rate λ_Δ . In this case,

$$p(T_D) = 1 - \frac{1}{\mathbb{E}[S + \Delta]} \int_{T_D}^{\infty} (x - T_D) dF(x) \\ = 1 - \lambda_\Delta \int_{T_D}^{\infty} (x - T_D) \lambda_\Delta e^{-\lambda_\Delta x} dx \\ = 1 - e^{-\lambda_\Delta T_D}.$$

Fig. 4(b) shows the corresponding measured packet loss rate obtained from simulations. Once again, the rate parameter λ_Δ can be directly estimated from the measured curve of packet loss versus T_D (namely from the slope when $p(T_D)$ is plotted on a log scale versus T_D). The ccdf is also shown in Fig. 4(d), and calculated as $1 - F(T_D) = e^{-\lambda_\Delta T_D}$.

D. Distinguishing Collision and Interference Losses in 802.11

The foregoing analysis focuses on packet loss due to interference and ignores other sources of packet loss. As already noted, packet loss due to collisions is part of the proper operation of the 802.11 MAC. In even quite small wireless LANs, the loss rate due to collisions can be significant (*e.g.* in a system with only two users, the collision probability can approach 5% [28]) and so it is essential to distinguish between packet loss due to collisions and packet loss due to noise/interference. To achieve this we borrow the packet-pair bursting idea first proposed in [1]. We make use of the following properties of the 802.11 MAC:

- 1) Time is slotted, with well-defined boundaries at which frame transmissions by a station are permitted.
- 2) The standard data-ACK handshake means that a sender-side analysis can reveal any frame loss.
- 3) Transmissions occurring before a DIFS are protected from collisions. This is used, for example, to protect ACK transmissions, which are transmitted after a SIFS interval.

Using property 3, when two frames are sent in a burst with a SIFS between them, the first frame is subject to both collision and noise losses but the second frame is protected from collisions and only suffers from noise/interference losses. Such packet-pair bursts can be generated in a number of ways (*e.g.* using the TXOP functionality in 802.11e/n, or the packet fragmentation functionality available in all flavours of 802.11).

For 802.11 links, we therefore consider sampling the channel using packet pair bursts rather than using single packets. For simplicity we will assume that the duration of both packets is the same and equal to $T_D/2$, although this can be relaxed. In the remainder of this paper we will often refer to the first packet in a burst as *pkt1*, and the second packet as *pkt2*. It is important to note that the 802.11 MAC only sends *pkt2* when an ACK is successfully received for *pkt1*. To retain the Lack of Anticipation property, when no ACK is received for the first packet we introduce a virtual transmission of the second packet *i.e.* no actual packet is transmitted but the sender still pauses for the time that it would have taken to send the second packet. In practice this is straightforward to implement by simply adding $T_D/2$ to the interval between packet pairs when an ACK for the first packet is not received. With this procedure, when the intervals between the completion of one packet pair and the start of the next packet pair form a Poisson process, the packet loss statistics will satisfy the ASTA property. Assuming that packet collisions occur independently of interference pulses, the packet loss rate for the first packet in the pair $\hat{p}_1(T_D/2)$ is then an estimator for

$$p_1(T_D/2) = 1 - \frac{1 - p_c}{\mathbb{E}[S + \Delta]} \int_{T_D/2}^{\infty} (x - \frac{T_D}{2}) dF(x),$$

where p_c is the packet collision probability. Note that it is difficult to separate out the contribution p_c due to collisions from measurements of $p_1(T_D/2)$, as already discussed. The second packet in a pair is only transmitted if the first packet was received successfully (per the standard 802.11 TXOP and fragmentation semantics) and so the second packet measurement data is censored. We therefore have that the packet loss rate for the second packet in the pair $\hat{p}_1(T_D/2)$ is an estimator for

$$p_2(T_D/2) = 1 - \frac{1}{(1-p_1(T_D/2))\mathbb{E}[S+\Delta]} \int_{T_D}^{\infty} (x - T_D) dF(x).$$

Combining the loss statistics $p_1(T_D/2)$ and $p_2(T_D/2)$ for the first and second packets, we can recover our desired loss statistic $p(T_D)$ from

$$p(T_D) = 1 - (1 - p_2(T_D/2))(1 - p_1(T_D/2)), \quad (4)$$

and in this way separate out the contribution to packet loss from interference from the contribution due to collisions.

E. Carrier Sense


The 802.11 MAC uses carrier sense to distinguish between busy and idle slots on the wireless medium. If the energy on the channel is sensed above the carrier-sense threshold, then the PHY_CCA.indicate(BUSY) signal will be issued by the PHY to indicate to the MAC layer that the channel is busy. Consequently, when an interference pulse is above the carrier-sense threshold at the transmitter, packet transmissions will not start. Instead, a packet waiting to be transmitted will be queued until the channel is sensed idle (PHY_CCA.indicate(IDLE)), and then transmitted. This means that the packet transmission times are no longer independent of the interference process and the ASTA property is generally lost. In particular, the packet loss rate is biased and tends to be underestimated since packet transmissions that should have started during an interference pulse (and so likely to have led to a packet loss) are deferred until after the pulse finished (and so much less likely to be lost since the time to the next interference pulse is then maximal).

When the duration of the interference pulses is short relative to the time between pulses, then the magnitude of this bias can be expected to be small. When the interference pulse duration is larger, an approximate compensation for the bias can be carried out as follows. Consider the indicator function


$$\tilde{U}_{T_D}(X_t) = \begin{cases} 1 & t \in [T_k + T_D, T_{k+1}) \text{ for some } k \\ 0 & \text{otherwise} \end{cases}.$$

This modifies (1) by lumping the time when the interference pulse is active into the *good* window, roughly capturing the fact that packet transmissions scheduled during a pulse will be deferred until the pulse finishes. When the interference pulse on and off times are i.i.d., this modified loss statistic is equal with probability one to

$$\begin{aligned} \tilde{p}(T_D) &= 1 - \frac{1}{\mathbb{E}[S+\Delta]} \int_0^{\infty} dG(y) \int_{T_D}^{\infty} (y+x-T_D) dF(x) \\ &= 1 - \frac{1}{\mathbb{E}[S+\Delta]} \left(\mathbb{E}[S] F_c(T_D) - \int_{T_D}^{\infty} (x-T_D) dF_c(x) \right) \\ &= p(T_D) - \epsilon, \end{aligned} \quad (5)$$



(a) Packet loss rate versus packet duration T_D with and without carrier sense



(b) Estimate $\hat{F}(T_D)$ of distribution function $F(T_D)$

Fig. 5. Simulation example illustrating how the estimation bias introduced by carrier sense can be largely removed using (6). Periodic interference, similar to the microwave oven interference experimentally measured in Section IV-B (period $\Delta = 11$ ms, pulse duration $S = 9$ ms).

where $F_c(x) = 1 - F(x)$ is the ccdf, $G(y) = \text{Prob}[S > y]$ and $\epsilon = \frac{\mathbb{E}[S]}{\mathbb{E}[S+\Delta]} F(T_D)$ is an approximation to the estimation bias. Using integration by parts and that $\tilde{p}(0) = 1 - \frac{\mathbb{E}[S]+\mathbb{E}[\Delta]}{\mathbb{E}[S+\Delta]}$, (5) can be rewritten as

$$\tilde{p}(T_D) = 1 - \frac{\mathbb{E}[S]}{\mathbb{E}[S] + \mathbb{E}[\Delta]} F_c(T_D) + \frac{1}{\mathbb{E}[S] + \mathbb{E}[\Delta]} \int_{T_D}^{\infty} F_c(x) dx. \quad (6)$$

Assuming that the measured packet loss rate approximates $\tilde{p}(T_D)$, then given measurements of loss rate for a range of T_D values we can solve equation (6) to obtain an estimate for $F(T_D)$ and $\mathbb{E}[S]$. This can be carried out in a number of ways – one simple approach is to write $F(x)$ as a weighted sum of $\sum_{i=1}^K w_i g_i(T_D)$ of orthogonal basis functions $\{g_i(T_D)\}$, and select the weights $\{w_i\}$ and $\mathbb{E}[S]$ to minimise the square error between the RHS of (6) and the measurement of the LHS. We illustrate use of this approach in Fig. 5, which presents data generated using a simulation with carrier sense and periodic interference. The on-time of the interference pulses is $S = 9$ ms and the time between pulses is $\Delta = 11$ ms. Fig. 5(a) plots the measured packet loss rate versus T_D which is assumed to approximate $\tilde{p}(T_D)$. Also shown is the loss rate $p(T_D)$ when carrier sense is disabled. The bias ϵ between $\tilde{p}(T_D)$ and $p(T_D)$ is clearly evident. Using this biased data for $\tilde{p}(T_D)$ and rectangular basis functions $\{g_i(T_D)\}$, solving (6) yields the estimate $\hat{F}(T_D)$ shown in Fig. 5(b). It can be seen that $\hat{F}(x)$ accurately estimates the true distribution function $F(T_D)$ (also marked in Fig. 5(b)) *i.e.* that we have successfully

TABLE I
SPECTRUM ANALYSER DETAILS AND SETUP FOR ZERO SPAN
MEASUREMENTS.

Model	Rohde & Schwarz FSL6 with optional pre-amp
Video BW	10 MHz
Resolution BW	10 & 20 MHz
Sweep time	20 ms
Antenna	LM Technologies LM254 2.4 GHz dipole

compensated for the carrier sense bias. In particular, the sharp transition at 11 ms is accurately estimated.

IV. EXPERIMENTAL MEASUREMENTS

In this section, we present experimental measurements demonstrating the power and practical utility of the proposed non-parametric estimation approach. We collected data in two separate measurement campaigns. The first consists of measurements on an 802.11 link affected by interference from a domestic microwave oven (MWO). Such interference is common, and so of considerable practical importance. The second shows measurements from an 802.11 lab testbed, with two transmitting nodes and a number of hidden nodes acting as the pulsed interference source.


A. Hardware and Software

Asus 700 laptops equipped with Atheros 802.11 a/b/g chipsets (radio 14.2, MAC 8.0, PHY 10.2) were used as client stations, running Debian Lenny 2.6.26 and using a modified Linux Madwifi driver based on 10.5.6 HAL and 0.9.4 driver. A Fujitsu Lifebook P7010 equipped with a Belkin Wireless G card using an Atheros 802.11 a/b/g chipset (AR2417, MAC 15.0, PHY 7.0) was used as the access point, running FreeBSD 8.0 with the RELEASE kernel and using the standard FreeBSD ATH driver. The beacon period is set to the maximum value of 1 s. We disabled the Atheros' Ambient Noise Immunity feature which has been reported to cause unwanted side effects [29]. Transmission power of the laptops is fixed and antenna diversity is 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 and the packet pair measurement approach is correctly implemented (see [1], [29], [30] for further details).


A Rohde & Schwarz FSL-6 spectrum analyser is used to verify that the test channels are unoccupied and also to capture the time-domain traces (see Table I for details).

B. Microwave Oven Interference

1) *Experimental Setup*: The experimental setup consisted of one client station, the AP and a 700 W microwave oven. During the experiments, the MWO is operated at maximum power to heat a 2 L bowl of water, and is located approximately 1 m away from the client station and AP; the exact geometry of the setup is not important since the MWO is close enough to the laptops to disrupt communications. The antenna connected to the spectrum analyser is located such that the energy from each RF source is of similar magnitude.



(a) Measured packet loss rate versus packet duration T_D . Confidence intervals based on the Clopper-Pearson method are displayed, but are small enough to be partially obscured by the point markers.



(b) Inter-arrival distribution of interference pulses

Fig. 6. Experimental measurements with microwave oven (MWO) interference. Data frames are transmitted at a PHY rate of 1 Mbps rate and the duration T_D is varied by adjusting the packet size. Both $pkt1$ and $pkt2$ are equal length $\frac{T_D}{2}$.

The MWO operates in the 2.4 GHz ISM band, with significant overlap ($> 50\%$) with the WiFi 20 MHz channels 6 to 13; this was verified using the spectrum analyser. Our 802.11 experiments used channels 7 and 9 and took place in a room that was cleared for co-channel interference before, during and after each experiment.

The client station transmits packets to the AP with the MTU, FRAG and packet size set to values that ensure that both $pkt1$ and $pkt2$ are of nearly identical duration (the deviation of $T_D/2$ is kept to below 1%). The packet duration is adjusted by varying the packet size between 30 and 2110 bytes (yielding T_D from 1.4 ms to 18 ms). These packets are generated using the standard `ping` command in a bash script. The interval between each set of packet pairs is exponentially distributed with rate $\lambda = 30$ packets per second, and the modulation and coding rate is fixed at 1 Mbps.

2) *Inferring Interference Statistics From Packet Loss Measurements*: Fig. 6(a) presents the measured packet loss rate between the client station and the AP versus the packet duration $T_D/2$. Each point is averaged over more than 10^4 observed packets. Using this packet loss data, Fig. 6(b) plots the estimated distribution function $\hat{F}(T_D)$ for interference pulse inter-arrival times. We use the approach described in Section III-E to compensate for the bias introduced by carrier sense at the client station. It can be seen that $\hat{F}(T_D)$ exhibits a sharp


transition around 11 ms, along with some residual probability mass between 11 and 15 ms. This indicates that the MWO interference is estimated to be approximately periodic with period $\Delta = 11$ ms. We confirm the accuracy of this inference independently using direct spectrum analyser measurements of the MWO interference in the next section, see Fig. 7.

Before proceeding however, it is worth comparing the experimentally measured 802.11 loss data in Fig. 6(a) with the simulation data in Fig. 4(a). This comparison highlights the additional complexity introduced by carrier sense and the censoring of second packet loss data. Nevertheless, our approach is able to successfully disentangle these effects in a principled way and thereby estimate $F(T_D)$.


3) *Validation*: Fig. 7(a) presents spectrum analyser data showing two interference pulses generated by the MWO. A packet pair transmission by the client station can also be seen, lying between the interference bursts (this particular packet pair transmission is successfully received by the AP, verified by noting the presence of MAC ACKs at the end of each packet). From this and other data, we find that the MWO interference is approximately periodic, with period $T = 1/f = 20$ ms *i.e.* a frequency of 50 Hz, as expected due to the AC circuitry that is driving the MWO. The profile of the interference bursts is, however, not uniform. Fig. 7(b) shows a measured interference burst where the interference power is roughly constant over the duration (approximately 9 ms) of the pulse. Fig. 7(c) shows an interference pulse where the interference power dips during the middle of the pulse, so as to effectively create two narrower pulses spaced approximately 4 ms apart. This variation in burst energy profile is attributed to frequency instability of the MWO cavity magnetron, a known effect in MWOs [31]. Our measurements indicate that the MWO interference consists of pulses with mean interval 11 ms between pulses, with some deviation (Fig. 6(b)). These direct measurements are therefore in good agreement with the estimated distribution function, which was derived indirectly using packet loss measurements.

C. 802.11 Network With Hidden Nodes


1) *Experimental Setup*: This test bed consists of a WLAN formed from two client stations and an access point, plus three additional stations configured as hidden nodes. These hidden nodes (HNs) are created by modifying the Madwifi driver such that the carrier sense is disabled (using the technique as detailed in [32]) and setting the NAV to zero for all packets – this effectively makes the HNs unresponsive to any packets that they decode from the client, or energy that may trigger a physical carrier sense. A script generates ping traffic on the hidden nodes having exponentially distributed intervals between packet transmissions, with a mean interval of 50 ms. The ping packets sent are of duration 4.5 ms (verified via the spectrum analyser). Since the transmissions by each HN are Poisson with intensity $\lambda = 20$ packets/s, the aggregate interference is also Poisson and with intensity $\lambda = 60$ packets/s. The experiments used channel 13 of the ISM band, and took place in a room that was cleared for co-channel interference before, during and after the experiments.



(a) Packet pair transmitted between two MWO bursts. The y-axis grid is in 2 ms increments. The packet pair is encoded at the 1 MBps 802.11 rate, with both packets having duration 4.36 ms.




(b) Second packet in a pair suffering a collision with a MWO burst; after the MWO burst has finished and carrier sense indicates the channel is idle, the packet is retransmitted. The y-axis grid is in 2 ms increments.




(c) Packet pair and a MWO burst. The y-axis grid is in 2 ms increments. The resolution bandwidth is set to 20 MHz, and thus captures about 99% of the WLAN signal. The MWO burst has a dip in the middle, which is attributed to frequency instability in the MWO cavity magnetron.

Fig. 7. Spectrum analyser measurements of microwave oven (MWO) interference.



(a) Measured packet loss rate versus packet duration T_D . Confidence intervals based on the Clopper-Pearson method are displayed, but are small enough to be partially obscured by the point markers.



(b) Inter-arrival distribution of interference pulses.

Fig. 8. Experimental measurements; primary network has two nodes transmitting to AP, interference network has three hidden nodes.

2) *Inferring Interference Statistics From Packet Loss Measurements*: Fig. 8(a) plots the measured packet loss rate in the WLAN versus the packet duration. Note that this loss rate includes a contribution due to collisions between the two client stations in the WLAN and a contribution due to interference from the hidden nodes. Nevertheless, using our packet pair approach we are able to disentangle these two sources of packet loss. Fig. 8(b) plots the resulting distribution of interference pulse inter-arrival times estimated using this packet loss data. The data plotted in Fig. 8(b) is the estimate of $1 - F(T_D)$, and is displayed using a logarithmic y -axis. Also plotted in Fig. 8(b) is the theory line $1 - F(T_D) = e^{-\lambda T_D}$ corresponding to Poisson distributed interference with rate $\lambda = 60$ packets/s. It can be seen that the estimated data is approximately linear on this log scale, as expected for a Poisson distribution, and that the slope is close to the expected value of $\lambda = 60$. The offset between the Poisson theory line and the estimated line is explained by the presence of a baseline packet loss rate of approximately 5% in our experimental setup – this baseline loss rate is confirmed by separate measurements (not shown here).

V. PULSED INTERFERENCE TEMPORAL STATISTICS: PARAMETRIC ESTIMATION

Thus far we have considered estimating the interference distribution function in a non-parametric manner. By making stronger, structural assumptions about the interference process,

we can alternatively parameterise the distribution function and our task then becomes one of estimating these model parameters. A fairly direct trade-off in effort is involved here, which is why it is important to consider both non-parametric and parametric approaches. Namely, we have the bias-variance trade-off whereby non-parametric approaches make only weak assumptions about the interference process, but require more measurement data, whereas parametric approaches make strong assumptions, but require less measurement data for the same estimation accuracy (assuming that the model structure is accurate).

In this section we present a parametric estimation approach for one class of model. The model is related to the two-state Gilbert-Elliot channel model [33], which is popular for analysing communication channels with bursty losses, extended to incorporate carrier sensing and the packet transmission process. Although simple, this model is useful and we demonstrate its effectiveness for estimating hidden terminal interference. A number of extensions are possible, including to a multi-state interference model [34], correlated losses [35], fast fading [36] and so on, but we leave consideration of these extensions to future work.

A. Parametric Packet Loss Model

1) *Interference*: We model pulsed interference as switching randomly between two states, “good” (G) and “bad” (B), with exponentially distributed dwell times in each state. Formally, let $\mathcal{S} = \{G, B\}$ denote the set of interference states,

$$Q = \begin{bmatrix} -\lambda_B & \lambda_B \\ \lambda_G & -\lambda_G \end{bmatrix}, \quad (7)$$

and

$$\Pi = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}. \quad (8)$$

Let $\mathcal{Y} = \{Y_n, n = 0, 1, 2, \dots\}$ be a sequence of random variables taking values in \mathcal{S} and representing the evolving state, with

$$\text{Prob}[Y_{n+1} = j | Y_n = i] = \Pi_{ij}. \quad (9)$$

With our choice of Π , the Y_n flip back and forth between the G and B states so that \mathcal{Y} is of the form $\{\dots, G, B, G, B, \dots\}$. Let $\{k\}$ index the sub-sequence of B states in \mathcal{Y} . Let S_k denote the dwell time in the k^{th} B state and Δ_k the dwell time in the following G state. The dwell times S_k and Δ_k are independent exponential random variables having, respectively, mean $1/\lambda_B$ and $1/\lambda_G$. The sequence $T_{k+1} = T_k + S_k + \Delta_k$ is the sequence of jump times at which the interference enters state B .

2) *Packet Transmissions*: The wireless station performing measurements transmits a sequence of packets to a destination station, with exponentially distributed pauses between transmissions. Similar to the foregoing interference model, we let $\{Tx, Idle\}$ be the two transmitter states, where Tx corresponds to transmission of a packet. Let $\{V_m, m = 0, 1, 2, \dots\}$ denote a sequence of random variables which flip back and forth between the Tx and $Idle$ states. The dwell time in the Tx state is a constant T_D , the dwell times in the $Idle$ state are

independent exponential random variables with mean $1/\lambda_D$. We index the sub-sequence of Tx states by packet numbers in $\{n\}$, and let t_n denote the time when transmission of packet n starts.

3) *Carrier Sense*: The interference state at the packet transmit time t_n is $Y_{k(n)}$ where $k(n) = \sup\{k : T_k \leq t_n\}$. Let

$$p_{cs} = \text{Prob}[t_n \in [T_{k(n)}, T_{k(n)} + \Delta_{k(n)}]] \\ := \alpha \frac{\lambda_G}{\lambda_G + \lambda_B},$$

where $0 \leq \alpha \leq 1$ and $\frac{\lambda_G}{\lambda_G + \lambda_B}$ is the probability that the interference is in state B . In the following, we consider two limiting situations. Firstly, where the carrier sense threshold lies above the noise level in both interference states, in which case the packet transmission times are decoupled from the interference state and $\alpha = 1$. Secondly, where the carrier sense threshold lies above the noise level in interference state G but below the noise level in state B , in which case $\alpha = 0$.

4) *Packet Loss*: Packets are discarded when they fail a checksum test at the receiver. Hence, we treat the channel as an erasure channel. Let δ_n denote a random variable that takes value 1 when packet n is erased and value 0 otherwise. Let \tilde{S}_n denote the time that the channel spends in state B during the transmission of packet n . In general, we expect that the probability $\text{Prob}[\delta_n = 1]$ that packet n is erased depends on \tilde{S}_n . Nevertheless, to streamline the presentation we make the simplifying assumption that $\text{Prob}[\delta_n = 1] = p_B$ whenever $\tilde{S}_n > 0$ and $\text{Prob}[\delta_n = 1] = p_G$ otherwise, where p_B and p_G are channel packet loss rate parameters in the B and G states respectively. We also assume that packet erasures occur independently, i.e. the random variables δ_n, δ_m are independent for $n \neq m$.

5) *Packet Error Rate Analysis*: To determine the packet error rate as a function of the packet transmit duration, we need to analyse two coupled stochastic processes, namely the channel and transmission processes. The joint process takes state values in $\{G, B\} \times \{Idle, Tx\}$. Since our interest is in counting the frequency of packet losses, observe that we can lump the $(Idle, G)$ and $(Idle, B)$ states together, since we know that no packet loss can occur in these $(Idle, \bullet)$ states. Also, when the system first enters state (Tx, B) , then a packet loss occurs and we do not need to keep count of the number of subsequent transitions between (Tx, G) and (Tx, B) . We can therefore partition time into slots, with each slot being of three possible types: *Idle* (corresponding to the lumped $(Idle, \bullet)$ states), *Loss* (corresponding to lumping of states (Tx, G) and (Tx, B) after the first transition from (Tx, G) to (Tx, B)) and *Transmitting* (corresponding to a dwell time in state (Tx, G)). The transitions between these slots are as shown in Fig. 9 and Table II.

The transition matrix P of this slotted time Markov chain is:

$$\mathbf{P} = \begin{bmatrix} 0 & 1 - p_{cs} & p_{cs} \\ 1 - p_i(T_D) & 0 & p_i(T_D) \\ 1 & 0 & 0 \end{bmatrix}, \quad (10)$$

where $1 - p_i(T_D) = \exp(-\lambda_B T_D)$. The stationary state distribution satisfies $\pi = \pi \mathbf{P}$, where $\pi_1 = \text{Prob}[Idle]$,




Fig. 9. Slotted time Markov chain.

$\pi_2 = \text{Prob}[Transmitting]$, and $\pi_3 = \text{Prob}[Loss]$. Solving yields,

$$\pi^T = \frac{1}{2 + p_i(T_D)(1 - p_{cs})} \begin{bmatrix} 1 \\ 1 - p_{cs} \\ (1 - p_{cs})p_i(T_D) + p_{cs} \end{bmatrix}.$$

The packet error probability for the first packet in a pair is

$$p_1(T_D) = \frac{(1 - p_i(T_D))\pi_2 p_G + p_i(T_D)\pi_2 p_B + p_{cs}\pi_1 p_B}{(1 - p_i(T_D))\pi_2 + p_i(T_D)\pi_2 + p_{cs}\pi_1} \\ = (1 - p_i(T_D))(1 - p_{cs})p_G + (p_i(T_D)(1 - p_{cs}) + p_{cs})p_B \\ =: G_1(T_D, \lambda_B, p_B, p_G, p_{cs}). \quad (11)$$

The first term in the expression for $p_1(T_D)$ corresponds to the event where the interference stays in state G throughout a packet transmission and a packet loss occurs. The second term corresponds to the event that a packet transmission starts with the interference in state G , but the interference changes to state B during the course of the transmission and a packet loss occurs. The third term corresponds to the event that a packet transmission starts with the interference in state B and a packet loss occurs.

Conditioned on the first packet transmission being successful, the packet error probability for the second packet in a pair is

$$p_2(T_D) = (1 - p_i(T_D)) \frac{\lambda_B}{\lambda_B + \lambda_G} p_G \\ + \left(1 - (1 - p_i(T_D)) \frac{\lambda_B}{\lambda_B + \lambda_G} \right) p_B \\ =: G_2(T_D, \lambda_B, p_B, p_G, p_{cs}), \quad (12)$$

where the $\frac{\lambda_B}{\lambda_B + \lambda_G}$ factor accounts for the event that the interference is in the B state upon starting transmission of *pkt2*.

B. Model parameters

Equations (11) and (12) together form a parametric model of the packet pair loss process, which is described by parameters λ_B, p_B, p_G and p_{cs} .

Before proceeding, we briefly illustrate how the model parameters λ_B, p_B, p_G and p_{cs} affect the observed packet

TABLE II
MARKOV MODEL STATE TRANSITIONS.

Idle \rightarrow Transmitting	(start Tx, interference in state G):	$1 - p_{cs}$
Idle \rightarrow Loss	(start Tx, interference in state B):	p_{cs}
Transmitting \rightarrow Idle	(interference in state G throughout Tx):	$1 - p_i = \exp(-\lambda_B T_D)$
Transmitting \rightarrow Loss	(interference enters state B during Tx):	$p_i = 1 - \exp(-\lambda_B T_D)$
Loss \rightarrow Idle	(Tx of damaged packet ends):	1




Fig. 10. Packet error rate versus packet duration T_D ; $\lambda_D = 30$, variable λ_B , $p_G = 0$, $p_B = 1$, $p_{cs} = 0$, $T_{SIFS} = 10 \mu s$.




Fig. 11. Packet error rate versus packet duration T_D ; $\lambda_D = 30$, $\lambda_B = 100$, $p_G = 0.1$, variable p_B , $p_{cs} = 0$, $T_{SIFS} = 10 \mu s$.

loss versus T_D curves. Our aim is to (i) illustrate the types of loss curves that the model is able to capture and (ii) gain some intuitive insight into the role of the various model parameters. Fig. 10 shows the impact of λ_B , which produces a horizontal shift in the loss curves. Fig. 11 shows the impact of p_B , which determines the right-hand asymptote of the loss curves. Fig. 12 shows the impact of the carrier sense parameter p_{cs} (by varying α), which produces a vertical shift in the left-hand asymptote. Although not shown, the impact of p_G also produces a vertical shift in the left-hand asymptote.

C. Maximum Likelihood Parameter Estimation

Our objective is to estimate the model parameters λ_B , p_B , p_G and p_{cs} from measurements of packet loss. The empirical estimators for loss probabilities $p_1(T_D)$ and $p_2(T_D)$ are

$$\hat{p}_1(T_D) = \frac{1}{N_1} \sum_{n=1}^{N_1} \delta_n^1 \quad \hat{p}_2(T_D) = \frac{1}{N_2} \sum_{n=1}^{N_2} \delta_n^2,$$

where N_1 is the number of first packets, N_2 the number of second packets, δ_n^1 is the indicator function that equals 1 when




Fig. 12. Packet error rate versus packet duration T_D ; $\lambda_D = 30$, $\lambda_B = 100$, $p_G = 0$, $p_B = 0$, variable p_{cs} (by varying α), $T_{SIFS} = 10 \mu s$.

the n^{th} first packet is lost and 0 otherwise, and similarly δ_n^2 for second packets. Collecting packet loss measurements for a sequence of packet durations T_{D_1}, T_{D_2}, \dots and stacking the corresponding loss probability estimates we have

$$\begin{bmatrix} \hat{p}_1(T_{D_1}) \\ \hat{p}_2(T_{D_1}) \\ \hat{p}_1(T_{D_2}) \\ \hat{p}_2(T_{D_2}) \\ \vdots \end{bmatrix} = \begin{bmatrix} G_1(T_{D_1}, \lambda_B, p_B, p_G, p_{cs}) \\ G_2(T_{D_1}, \lambda_B, p_B, p_G, p_{cs}) \\ G_1(T_{D_2}, \lambda_B, p_B, p_G, p_{cs}) \\ G_2(T_{D_2}, \lambda_B, p_B, p_G, p_{cs}) \\ \vdots \end{bmatrix} + \eta, \quad (13)$$

where η denotes the estimation error in the packet loss estimates. For N_1, N_2 sufficiently large, the estimation error η is close to being Gaussian distributed. The maximum likelihood estimates for parameters λ_B , p_B , p_G and p_{cs} are then the values that minimise the square error between the LHS and RHS in (13).

D. Experimental Measurements

1) *Experimental Setup*: We revisit the WLAN experimental setup discussed in Section IV-C, but now change the setup slightly so that only a single wireless client (rather than two clients) transmits in the WLAN. This change is introduced because, for simplicity, we have not included packet collisions in our parametric model.

2) *Packet Loss Measurements*: Fig. 13 shows the measured packet loss rate versus the packet duration T_D . Note that the range of packet durations that we can use is constrained by the maximum 802.11 frame size of 2272 B to lie in the interval 1.4 ms to 18 ms. Two sets of results are shown, for one and for three hidden nodes active. Each experimental point is calculated as the average of more than 6×10^5 packet transmissions. Also shown are the maximum likelihood fits to this data using parametric model (11) and (12); the




Fig. 13. Experimental measurements and model fit for WLAN with hidden node interference. Data points are for experiments using 1 and 3 interferers, with each interferer having a packet transmission rate of $\lambda_B = 20$. Initial values for the parameter estimator are $\hat{\lambda}_B = 20$, $\hat{p}_{cs} = 0$, $p_G = 0$, and $p_B = 0.5$. Model parameters are given in Table III.

Number of interferers	$\hat{\lambda}_B$	\hat{p}_{cs}	\hat{p}_G	\hat{p}_B
1	19.9932	0.0286	0.0080	0.2678
3	54.7173	0.1011	0.0055	0.4055

TABLE III

DETAILS OF THE MAXIMUM LIKELIHOOD PARAMETER ESTIMATES FOR MEASUREMENT DATA IN FIG. 13

corresponding model parameter estimates are given in Table III, obtained using an interior-point solver.

3) *Validation*: The hidden node interferers each make transmissions with exponentially distributed idle time between packets so that the mean transmit rate is 20 packets/s. When one interferer is active, we expect $\lambda_B = 20$ and when three interferers are active we expect $\lambda_B = 60$. It can be seen from Table III that the model estimates are close to these predictions. Measurements taken with no hidden node interferers active indicate that the baseline packet loss rate is less than 1% and it can be seen from Table III that the model estimate for p_G is in good agreement with this. While it is difficult to similarly validate the estimates for parameters p_{cs} and p_B , we note that the estimated values are very reasonable.

4) *Parametric vs Nonparametric Estimation*: A parametric model makes strong structural assumptions that allow the loss curves to be parameterised using a small number of parameters. Since there are fewer parameters, we expect to be able to estimate their values with less data, but at the cost of introducing a bias if the structural assumptions turn out to be incorrect. Fig. 14 plots $\max_x |\hat{F}_\infty(x) - \hat{F}_N(x)|$ versus the number of observed packets N for both the parametric and non-parametric approaches, where $\hat{F}_N(x)$ is the estimate of $F(x)$ obtained using N observations and $\hat{F}_\infty(x)$ is the estimate using all 6×10^5 observations. For the parametric model, the parameter estimates are fed back into the model equations (11) and (12), and the resulting parameterised P_e curves are used to calculate $\hat{F}_N(x)$. This provides a rough indication of how estimates converge as the amount of data is increased. It can be seen that the parametric solution converges to within 5% of the asymptotic estimate after $N = 900$ packets and to within 2.5% after $N = 4000$ packets, while the non-parametric solution requires $N = 6000$ and $N = 20000$




Fig. 14. Convergence of estimates of $F(x)$ versus the number of packets observed. $\hat{F}_N(x)$ denotes the estimate using N packet observations and $\hat{F}_\infty(x)$ denotes the estimate obtained using the full measurement trace. For each N , we take 100 random subsamples of N packets from the full measurement trace, calculate $\max_x |\hat{F}_\infty(x) - \hat{F}_N(x)|$ for each subsample, and average this value over the 100 subsamples to obtain the curves shown. Data is shown for both parametric and non-parametric estimates. The data set used is from the three interferer experiment, see Fig. 13.




Fig. 15. Spectrum analyser snapshot of hidden terminal interferers in time. The y-axis grid is in 2 ms increments. Interferer burst durations are fixed at 4.5 ms, with arrivals at 10, 19, 80, 83 and 89 ms. Since each interferer has a different path to the spectrum analyser antenna, the pulses are at different power levels. The third and fourth pulses collide, resulting in a stepped feature.

packets, respectively, to achieve the same level of estimation accuracy.

5) *Discussion*: It is interesting to note that, despite its simplicity, the parametric model used here is remarkably effective at capturing the behaviour in a complex physical environment. For example, the model ignores the fact that the interference power will depend on the *number* of hidden node transmissions taking place at the same time. This effect can be seen in the spectrum analyser measurements in Fig. 15, where overlapping transmissions by interferers leads to a stepped interference pulse profile. The model also assumes that the duration of interference pulses is exponentially distributed, but this will not be the case in our experimental setup. More complex parametric models are also possible, and in particular can leverage the wealth of research on bursty communications channels, but we leave this to future work.

VI. CONCLUSIONS

In this paper we propose a new approach for detecting the presence of pulsed interference affecting 802.11 links, and for estimating temporal statistics of this interference. Our

approach is a transmitter-side technique that provides per-link information and is compatible with standard hardware. This significantly extends recent work in [1], [2] which establishes a MAC/PHY cross-layer technique capable of classifying lost transmission opportunities into noise-related losses, collision induced losses, hidden-node losses and of distinguishing these losses from the unfairness caused by exposed nodes and capture effects.

VII. ACKNOWLEDGEMENTS

Many thanks to Ken Duffy and Giuseppe Bianchi for their helpful comments.

REFERENCES

- [1] D. Giustiniano, D. Malone, D. J. Leith, and K. Papagiannaki, "Measuring Transmission Opportunities in 802.11 Links," *IEEE/ACM Transactions on Networking*, vol. 18, no. 5, pp. 1516–1529, Oct. 2010.
- [2] D. J. Leith and D. Malone, "Field Measurements of 802.11 Collision, Noise and Hidden-Node Loss Rates," in *IEEE Proc. 8th International Symp. on Modeling and Opt. in Mobile, Ad Hoc and Wireless Networks (WiOpt)*, 2010, pp. 412–417.
- [3] K. Rele and D. Roberson and B. Zhang and L. Li and Y. B. Yap and T. Taher and D. Ucci and K. Zdunek, "A Two-Tiered Cognitive Radio System for Interference Identification in 2.4 GHz ISM Band," in *7th IEEE Consumer Commun. and Networking Conf. (CCNC)*, Jan. 2010, pp. 1–5.
- [4] D. Qiao, S. Choi, and K. Shin, "Goodput Analysis and Link Adaptation for IEEE 802.11a Wireless LANs," *IEEE Trans. on Mobile Computing*, vol. 1, no. 4, pp. 278–292, Oct-Dec. 2002.
- [5] I. Haratcherev, K. Langendoen, R. Legendijk, and H. Sips, "Hybrid rate control for IEEE 802.11," in *Proceedings of the 2nd Int. Workshop on Mobility Management & Wireless Access Protocols (MobiWac)*, ser. MobiWac '04. New York, NY, USA: ACM, 2004, pp. 10–18. [Online]. Available: <http://doi.acm.org/10.1145/1023783.1023787>
- [6] T. Zhou and X. Wang and W. Hou, "A Fast Collision Detection Algorithm in IEEE 802.11 through Physical Layer SINR Monitoring," in *Proc. of the 73rd IEEE Vehic. Tech. Conf. (VTC)*, May 2011.
- [7] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11b mesh network," in *Proc. of the 2004 Conf. on Applications, Tech., Architectures, and Protocols for Computer Commun. (SIGCOMM)*, 2004, pp. 121–132.
- [8] S. Rayanchu and A. Mishra and D. Agrawal and S. Saha and S. Banerjee, "Diagnosing Wireless Packet Losses in 802.11: Separating Collision from Weak Signal," in *27th Annual IEEE Int. Conf. on Computer Commun. (INFOCOMM)*, Apr. 2008, pp. 735–743.
- [9] J. M. III and G. Q. M. Jr., "Cognitive radio: making software radios more personal," *IEEE Personal Communications*, vol. 6, no. 4, pp. 13–18, Aug. 1999.
- [10] A. Kamerman and L. Monteban, "Goodput analysis and link adaptation for IEEE 802.11a wireless lans," *Bell Labs Tech. Jour.*, pp. 118–133, Summer 1997.
- [11] G. Holland and N. Vaidya and P. Bahl, "A Rate-Adaptive MAC Protocol for Multi-Hop Wireless Networks," in *Proc. of the Int. Conf. on Mobile Computing and Networking (MobiCom)*, July 2001, pp. 236–251.
- [12] S. H. Y. Wong, H. Yang, S. Lu, and V. Bharghavan, "Robust rate adaptation for 802.11 wireless networks," in *Proc. of the 12th Annual Int. Conf. on Mobile Computing and Networking (MobiCom)*. New York, NY, USA: ACM, 2006, pp. 146–157. [Online]. Available: <http://doi.acm.org/10.1145/1161089.1161107>
- [13] Q. Pang, V. C. M. Leung, and S. C. Liew, "A rate adaptation algorithm for IEEE 802.11 wlan based on mac-layer loss differentiation," in *2nd Int. Conf. on Broadband Networks (BroadNets)*, Oct. 2005, pp. 659 – 667.
- [14] J. Kim and S. Kim and S. Choi and D. Qiao, "CARA: Collision-aware Rate Adaptation for IEEE 802.11 WLANs," in *25th Annual IEEE Int. Conf. on Computer Commun. (INFOCOMM)*, Apr. 2006.
- [15] D. Malone, P. Clifford, D. Reid, and D. J. Leith, "Experimental Implementation of Optimal WLAN Channel Selection Without Communication," in *Proc. IEEE Dynamic Spectrum Access Netw. (DySPAN)*, 2007, pp. 316–319.
- [16] M. Khan and D. Veitch, "Isolating Physical PER for Smart Rate Selection in 802.11," in *28th Annual IEEE Int. Conf. on Computer Commun. (INFOCOMM)*, Apr. 2009, pp. 1080–1088.
- [17] M. N. Krishnan, S. Pollin, and A. Zakhor, "Local estimation of probabilities of direct and staggered collisions in 802.11 wlan," in *IEEE Global Telecommunications Conference, 2009. (GLOBECOM)*, Dec. 2009, pp. 1–8.
- [18] E.-I. Kim and J.-R. Lee and D.-H. Cho, "Throughput Analysis of Data Link Protocol with Adaptive Frame Length in Wireless Networks," *AEU Int. J. Electron. Commun.*, pp. 1–8, 2003.
- [19] J. Yin, X. Wang, and D. P. Agrawal, "Optimal packet size in error-prone channel for IEEE 802.11 distributed coordination function," in *IEEE Wireless Communications and Networking Conference (WCNC)*, vol. 3, March 2004, pp. 1654–1659.
- [20] X. He and F. Y. Li and J. Lin, "Link Adaptation with Combined Optimal Frame Size and Rate Selection in Error-prone 802.11n Networks," in *Proc. of IEEE Int. Symp. on Wireless Commun. Systems (ISWCS)*, 2008, pp. 733–737.
- [21] S. Choi and K. G. Shin, "A Class of Adaptive Hybrid ARQ Schemes for Wireless Links," *IEEE Trans. on Vehic. Tech.*, pp. 777–790, 2001.
- [22] S. Ci and H. Sharif, "An Link Adaptation Scheme for Improving Throughput in the IEEE 802.11 Wireless LAN," in *Proc. of 27th Annual IEEE Conf. on Local Computer Networks (LCN)*, Nov. 2002, pp. 205–208.
- [23] P. Lettieri and M. B. Srivastava, "Adaptive frame length control for improving wireless link throughput, range, and energy efficiency," in *Proc. of the IEEE 17th Annual Joint Conf. of the IEEE Computer and Commun. Soc. (INFOCOM)*, vol. 2, Mar 1998, pp. 564–571.
- [24] M. Krishnan and E. Haghani and A. Zakhor, "Packet Length Adaptation in WLANs with Hidden Nodes and Time-Varying Channels," in *submitted to IEEE INFOCOMM*, 2011.
- [25] R. W. Wolff, "Poisson Arrivals See Time Averages," *Operations Research*, vol. 30, no. 2, pp. 223–231, 1982.
- [26] B. Melamed and W. Whitt, "On Arrivals That See Time Averages," *Operations Research*, vol. 38, no. 1, pp. 156–172, 1990.
- [27] F. Baccelli, S. Machiraju, D. Veitch, and J. Bolot, "The Role of PASTA in Network Measurement," *IEEE/ACM Trans. on Networking*, vol. 17, no. 4, pp. 1340–1353, 2009.
- [28] D. Malone and K. Duffy and D.J. Leith, "Modeling the 802.11 Distributed Coordination Function in Nonsaturated Heterogeneous Conditions," *IEEE/ACM Trans. on Networking*, vol. 15, no. 1, pp. 159 –172, Feb. 2007.
- [29] I. Tinnirello, D. Giustiniano, L. Scalia, and G. Bianchi, "On the side-effects 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, qoS Aspects in Next-Generation Networks. [Online]. Available: <http://www.sciencedirect.com/science/article/B6VRG-4TTMK0V-1/2/71fc79de16b0855ac93813d9babf0dc1>
- [30] K. D. Huang, K. R. Duffy, and D. Malone, "On the validity of IEEE 802.11 MAC modeling hypotheses," *IEEE/ACM Transactions on Networking*, vol. 18, no. 6, pp. 1935 – 1948, 2010.
- [31] A. Kamerman and N. Erkocevic, "Microwave Oven Interference on Wireless LANs Operating in the 2.4 GHz ISM Band," in *IEEE Int. Symp. on Personal, Indoor and Mobile Radio Commun. (PIMRC)*, vol. 3, September 1997, pp. 1221–1227.
- [32] E. Anderson and G. Y. and C. Phillips and D. Sicker and D. Grunwald, "Commodity AR52XX-Based Wireless Adapters as a Research Platform," University of Colorado at Boulder, Department of Computer Science, Campus Box 430, Technical Report CU-CS-XXXX-08, April 2008.
- [33] M. Mushkin and I. Bar-David, "Capacity and coding for the Gilbert-Elliot channels," *IEEE Transactions on Information Theory*, vol. 35, no. 6, pp. 1277–1290, Nov. 1989.
- [34] S. Sivaprakasam and K. S. Shanmugan, "An equivalent Markov model for burst errors in digital channels," *IEEE Transactions on Communications*, vol. 43, no. 234, pp. 1347–1355, Feb/Mar/Apr 1995.
- [35] M. Yajnik and S. Moon and J. Kurose and D. Towsley, "Measurement and Modeling of the Temporal Dependence in Packet Loss," in *18th Annual IEEE Int. Conf. on Computer Commun. (INFOCOMM)*, vol. 1, Mar. 1999, pp. 345 – 352.
- [36] H. S. Wang and N. Moayeri, "Finite-state Markov channel-a useful model for radio communication channels," *IEEE Transactions on Vehicular Technology*, vol. 44, no. 1, pp. 163–171, Feb. 1995.

PLACE
PHOTO
HERE

Brad Zarikoff received the B.Eng. degree with distinction in electrical engineering from the University of Victoria, Victoria, Canada, in 2002 and the M.A.Sc. and Ph.D. degrees from Simon Fraser University, Burnaby, Canada, in 2004 and 2008, respectively. He is currently a research fellow at the Hamilton Institute, National University of Ireland Maynooth. His current research interests include interference mitigation, power line communication networks, and synchronisation for network MIMO systems.

PLACE
PHOTO
HERE

Doug Leith graduated from the University of Glasgow in 1986 and was awarded his PhD, also from the University of Glasgow, in 1989. In 2001, Prof. Leith moved to the National University of Ireland, Maynooth to assume the position of SFI Principal Investigator and to establish the Hamilton Institute (www.hamilton.ie) of which he is Director. His current research interests include the analysis and design of network congestion control and resource allocation in wireless networks.