This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

Measuring Pulsed Interference in 802.11 Links

Brad W. Zarikoff, *Member, IEEE*, and Douglas J. Leith, *Senior Member, IEEE*

***Abstract*—Wireless IEEE 802.11 links operate in unlicensed spectrum and so must accommodate other unlicensed transmitters that 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.**

***Index Terms*—802.11, CSMA/CA, interference.**

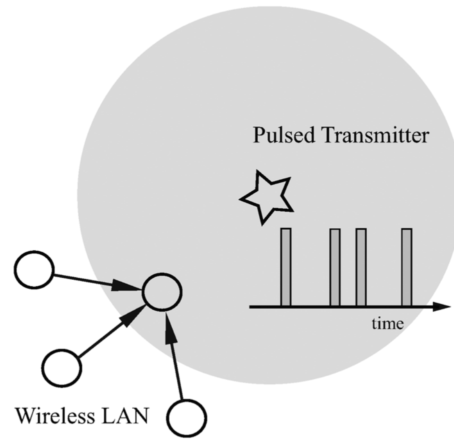

Fig. 1. Illustrating a WLAN with interfering pulsed transmitter (e.g., 802.11 hidden terminal, Bluetooth device, microwave oven, baby monitor, etc.) inducing packet loss.

I. Introduction

Wireless IEEE 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, received signal strength indicator (RSSI)/signal-to-interference-plus-noise ratio (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.

Manuscript received July 27, 2011; revised February 19, 2012 and May 13, 2012; accepted May 24, 2012; approved by IEEE/ACM Transactions on Networking Editor K. Papagiannaki. This work was supported by Science Foundation Ireland under Grants 07/IN.1/I901 and 08/SRC/I1403.

The authors are with the Hamilton Institute, NUI Maynooth, Maynooth, Ireland (e-mail: zarikoff@ieee.org).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TNET.2012.2202686

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. We use the observation that a packet transmission can be thought of as sampling the channel conditions over an interval of time equal to the duration of the packet transmission. By varying the packet transmit duration and observing the corresponding change in packet loss rate, we can infer information about the timing of pulsed interference. This approach is a transmitter-side technique that provides per-link information and is compatible with standard hardware. It significantly extends recent work in [1] and [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.

Detection and measurement of pulsed interference is particularly topical in view of the trend toward 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 Mb/s—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

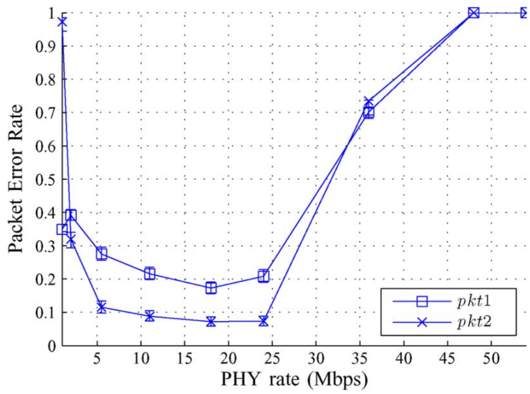

Fig. 2. Experimental measurements of PER versus MCS for an 802.11 network operating on channel 9 and physically located near an operational MWO. See Section IV-B for further details of the experimental setup. Two curves are shown, one for each fragment of a two packet TXOP burst. Observe that the PER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

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 1 Mb/s, 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 are *PHY link-level* approaches using SINR and bit error rate (BER). Second are *MAC approaches* relying on throughput and delay statistics, or frame loss statistics derived from transmitted frames that are not ACKed and/or from signaling messages. Finally, we have *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. References [8] and [9] consider 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. The cognitive radio literature considers PHY-layer techniques for optimizing performance in the presence of interference via joint spectral and temporal analysis [10]. There are some solutions tailored to the ISM band [3], where customized hardware has been devised with the aim of providing a synchronization signal based on periodic interference. However, cognitive radio techniques are largely geared toward interference avoidance and make use of nonstandard hardware.

MAC approaches make up some of the most popular and earliest rate control algorithms. Techniques such as ARF [11], RBAR [12], and RRAA [13] 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 [14], 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., [15] and [16]. However, such approaches can perform poorly in the presence of pulsed interference such as hidden terminals [1].

With regard to combined MAC/PHY approaches, this paper builds upon the packet pair approach proposed in [1] and [2] for estimating the frame error rates due to collisions, noise, and hidden terminals. See also the closely related work in [17]. References [1], [2], and [17] focus on time-invariant channels and do not consider estimation of temporal statistics. Reference [18] 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. Reference [19] modifies the Gilbert–Elliott channel model to model bursty channels. However, it does not consider the MAC layer. There are many examples that use MAC frame error information [20]–[24], 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 [25], which relies on separation between noise errors and collision errors as a means of tuning the packet length and optimizing throughput.

III. Pulsed Interference Temporal Statistics: Nonparametric 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

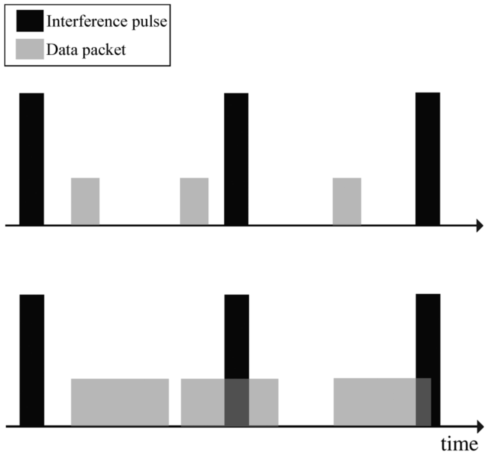

Fig. 3. Schematic illustrating "sampling" of a time-varying channel by data packet transmissions. Since the data transmissions occupy an interval of time, the sampling is of the channel conditions over that interval, rather than at a single point in time. As the duration of the data transmissions increases,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

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 characterized 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.

B. Mathematical Analysis

We now formalize these claims. Consider a sequence of interference pulses indexed by $k = 0, 1, 2, \ldots$, and let T_k denote the start time of the kth interference pulse with $T_0 = 0$, $S_k > 0$ denote the duration of the kth pulse, and $\Delta_k = T_{k+1} - (T_k + S_k) > 0$ be the interval between the end of kth pulse and the start of the $(k+1)$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 Δ_k, $k = 1, 2, \ldots$ are i.i.d. with finite mean. Then, $\Delta_k \stackrel{d}{=} \Delta$, where $\stackrel{d}{=}$ denotes equality in distribution, and let $\mathrm{Prob}[\Delta \le x] = F(x)$. 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 $[t - T_{\mathrm{D}}, t]$. This sampling interval can be thought of as a packet transmission ending at time t. Define indicator function $U_{T_{\mathrm{D}}}(X_t) = 1$ if interval $[t - T_{\mathrm{D}}, t]$ does not overlap with any interference pulse, and $U_{T_{\mathrm{D}}}(X_t) = 0$ otherwise. That is

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

Suppose we transmit a sequence of packets and let $\{t_j\}$ denote the sequence of times when transmissions finish. Assume for the moment that: 1) a packet is lost whenever it overlaps with an interference pulse; and 2) the intervals between packet transmissions are exponentially randomly distributed and are independent of the interference process. We will shortly relax these assumptions. By assumption 1), $U_{T_{\mathrm{D}}}(X_{t_j})$ 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

$$\hat{P}_t(T_{\mathrm{D}}) = 1 - \frac{1}{N(t)} \sum_{j=1}^{N(t)} U_{T_{\mathrm{D}}}(X_{t_j}) \tag{2}$$

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 2) the transmit times $\{t_j\}$ effectively possess the Lack of Anticipation property (the number of packet transmissions in any interval $[t, t+u]$, $u \ge 0$, is independent of $\{X_s\}$, $s \le t$ [26]). When this property holds, by [26, Theorem 1], we almost surely have

$$\lim_{t\to\infty} \hat{P}_t(T_{\mathrm{D}}) = \lim_{t\to\infty} P_t(T_{\mathrm{D}}) =: p(T_D)$$

where

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

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

Assumption 1) 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 2) can be relaxed to any sampling approach that satisfies the Arrivals See Time Averages (ASTA) property; see, for example, [27] and [28].

It remains to show that statistic $p(T_{\mathrm{D}})$ contains useful information about the interference process. We begin by observing that $Y_t = \sup\{k : T_k \le t\}$ 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 $U_{T_{\mathrm{D}}}(X_t) = 1$ for duration $[\Delta_k - T_{\mathrm{D}}]^+$, where $[x]^+$ equals x when $x \ge 0$, and 0 otherwise. The mean value of $U_{T_{\mathrm{D}}}(X_t)$ over a renewal interval is therefore $\int_{T_{\mathrm{D}}}^{\infty}(x - T_D)dF(x)$ and, by the strong law of large numbers

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

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

$$\frac{dp}{dT_{\mathrm{D}}}(T_{\mathrm{D}}) = \frac{1}{\mathbb{E}[S+\Delta]}\int_{T_{\mathrm{D}}}^{\infty} dF(x)$$
$$= \frac{1}{\mathbb{E}[S+\Delta]}\mathrm{Prob}[\Delta > T_{\mathrm{D}}].$$

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

Provided $p(\bullet)$ is differentiable at $T_{\mathrm{D}} = 0$, then

$$\mathbb{E}[S + \Delta] = \frac{1}{\frac{dp(0)}{dT_{\mathrm{D}}}}$$

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

$$\mathrm{Prob}[\Delta > T_{\mathrm{D}}] = \frac{1}{\frac{dp(0)}{dT_{\mathrm{D}}}} \frac{dp}{dT_{\mathrm{D}}}(T_{\mathrm{D}}). \tag{3}$$

Hence, knowledge of statistic $p(T_{\mathrm{D}})$ 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 $F(x) = 1 - \mathrm{Prob}[\Delta > x]$ of the interference pulse interarrival times.

Note that while we can formally differentiate $p(T_{\mathrm{D}})$, its estimate $\hat{p}(T_{\mathrm{D}})$ will be noisy, and so differentiating $\hat{p}(T_{\mathrm{D}})$ is not advisable. The formal differentiation step is merely used to gain insight into the statistical information contained within $p(T_{\mathrm{D}})$, and there is no need to actually differentiate $\hat{p}(T_{\mathrm{D}})$ 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 $p(T_{\mathrm{D}})$ 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 $\mathrm{Prob}(\Delta = T_\Delta) = 1$) and packets are always lost when they overlap with an interference pulse. In this case

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

That is, $p(T_{\mathrm{D}})$ 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 complementary cumulative distribution function (ccdf) $1 - F(T_{\mathrm{D}})$ 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

$$\begin{aligned} p(T_{\mathrm{D}}) &= 1 - \frac{1}{\mathbb{E}[S + \Delta]} \int_{T_{\mathrm{D}}}^{\infty} (x - T_{\mathrm{D}}) dF(x) \\ &= 1 - \lambda_\Delta \int_{T_{\mathrm{D}}}^{\infty} (x - T_{\mathrm{D}}) \lambda_\Delta e^{-\lambda_\Delta x} dx \\ &= 1 - e^{-\lambda_\Delta T_{\mathrm{D}}}. \end{aligned}$$

Fig. 4(c) 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 $p(T_{\mathrm{D}})$ is plotted on a log scale versus T_{D}). The ccdf is also shown in Fig. 4(d) and calculated as $1 - F(T_{\mathrm{D}}) = e^{-\lambda_\Delta T_{\mathrm{D}}}$.

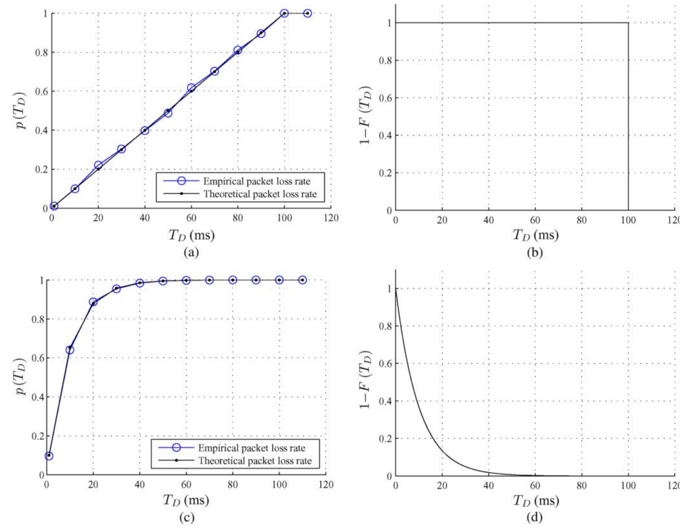

Fig. 4. Theory and simulation for periodic and Poisson interference. Packet transmissions are Poisson with mean rate $\lambda = 0.01$. (a) Periodic interference, period $T_\Delta = 100$ ms. (b) ccdf of Δ for periodic interference. (c) Poisson interference, mean interarrival time $1/\lambda_\Delta = 10$ ms. (d) ccdf of Δ for Poisson interference.

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% [29]), and so it is essential to distinguish between packet loss due to collisions and packet loss due to noise/inteference. 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 802.11 e/n, or the packet fragmentation functionality available in all flavors 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 $T_{\mathrm{D}}/2$, although this can be relaxed. In the remainder of this paper, we will often refer to the first packet in a burst as $\mathrm{pkt1}$, and the second packet as $\mathrm{pkt2}$. It is important to note that the 802.11 MAC only sends $\mathrm{pkt2}$ when an ACK

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

is successfully received for $\mathrm{pkt1}$. 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 $T_{\mathrm{D}}/2$ 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 $\hat{p}_1(T_{\mathrm{D}}/2)$ is then an estimator for

$$p_1\left(\frac{T_{\mathrm{D}}}{2}\right) = 1 - \frac{1-p_{\mathrm{c}}}{\mathbb{E}[S+\Delta]}\int_{\frac{T_{\mathrm{D}}}{2}}^{\infty}\left(x-\frac{T_{\mathrm{D}}}{2}\right)dF(x)$$

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 $p_1(T_{\mathrm{D}}/2)$, 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 $\hat{p}_1(T_{\mathrm{D}}/2)$ is an estimator for

$$p_2\left(\frac{T_{\mathrm{D}}}{2}\right) = 1 - \frac{1}{\left(1-p_1\left(\frac{T_{\mathrm{D}}}{2}\right)\right)\mathbb{E}[S+\Delta]}\int_{T_{\mathrm{D}}}^{\infty}(x-T_{\mathrm{D}})dF(x).$$

Combining the loss statistics $p_1(T_{\mathrm{D}}/2)$ and $p_2(T_{\mathrm{D}}/2)$ for the first and second packets, we can recover our desired loss statistic $p(T_{\mathrm{D}})$ from

$$p(T_{\mathrm{D}}) = 1 - \left(1-p_2\left(\frac{T_{\mathrm{D}}}{2}\right)\right)\left(1-p_1\left(\frac{T_{\mathrm{D}}}{2}\right)\right) \quad (4)$$

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

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

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 1 to

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

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

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

Assuming that the measured packet loss rate approximates $\tilde{p}(T_{\mathrm{D}})$, then given measurements of loss rate for a range of T_{D}-values, we can solve (6) to obtain an estimate for $F(T_{\mathrm{D}})$ 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 $\sum_{i=1}^{K} w_i g_i(T_{\mathrm{D}})$ of orthogonal basis functions $\{g_i(T_{\mathrm{D}})\}$, and select the weights $\{w_i\}$ and $\mathbb{E}[S]$ to minimize the square error between the right-hand side (RHS) of (6) and the measurement of the left-hand side (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 $\tilde{p}(T_{\mathrm{D}})$. Also shown is the loss rate $p(T_{\mathrm{D}})$ when carrier sense is disabled. The bias ϵ between $\tilde{p}(T_{\mathrm{D}})$ and $p(T_{\mathrm{D}})$ is clearly evident. Using this biased data for $\tilde{p}(T_{\mathrm{D}})$ and rectangular basis functions $\{g_i(T_{\mathrm{D}})\}$, solving (6) yields the estimate $\hat{F}(T_{\mathrm{D}})$ shown in Fig. 5(b). It can be seen that $\hat{F}(x)$ accurately estimates the true distribution function $F(T_{\mathrm{D}})$ [also marked in Fig. 5(b)], i.e., that we have successfully 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

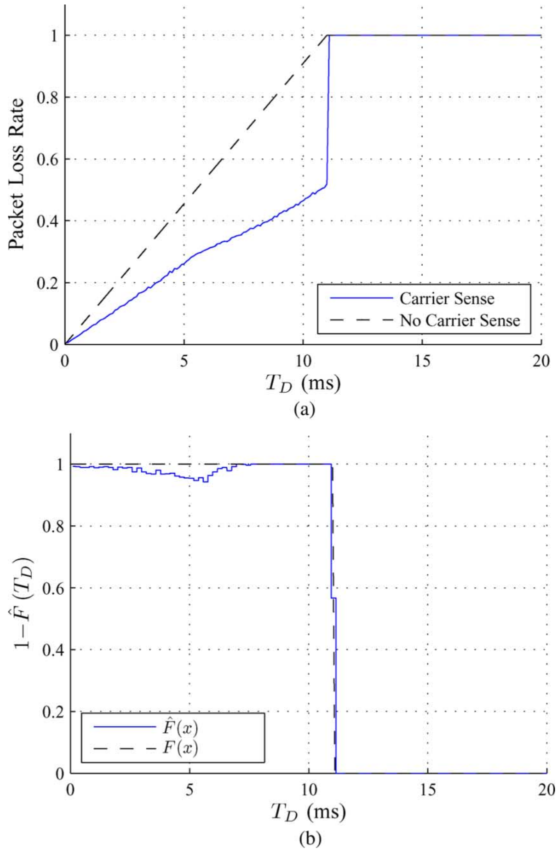

Fig. 5. Simulation example illustrating how the estimation bias introduced by carrier sense can be largely removed using (6). Periodic interference, similar to the microwave oven interference experimentally measured in Section IV-B (period $\Delta = 11$ ms, pulse duration $S = 9$ ms). (a) Packet loss rate versus packet duration $T_{\rm D}$ with and without carrier sense. (b) Estimate $\hat{F}(T_{\rm D})$ of distribution function $F(T_{\rm D})$.

nonparametric 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 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 [30]. Transmission power of the laptops is fixed, and antenna diversity is disabled. Note that these cards do not possess microwave

TABLE I
SPECTRUM ANALYZER 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

over robustness features. 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], [30], and [31] for further details).

A Rohde & Schwarz FSL-6 spectrum analyzer 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 access point (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 analyzer is located such that the energy from each RF source is of similar magnitude.

The MWO operates in the 2.4-GHz ISM band, with significant overlap (>50%) with the WiFi 20-MHz channels 6–13; this was verified using the spectrum analyzer. 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 $T_{\rm D}/2$ is kept to below 1%). The packet duration is adjusted by varying the packet size between 30 and 2110 B (yielding $T_{\rm D}$ from 1.4 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 Mb/s.

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 $T_{\rm D}/2$. Each point is averaged over more than 10^4 observed packets. Using this packet loss data, Fig. 6(b) plots the estimated distribution function $\hat{F}(T_{\rm D})$ for interference pulse interarrival 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 $\hat{F}(T_{\rm D})$ 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

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

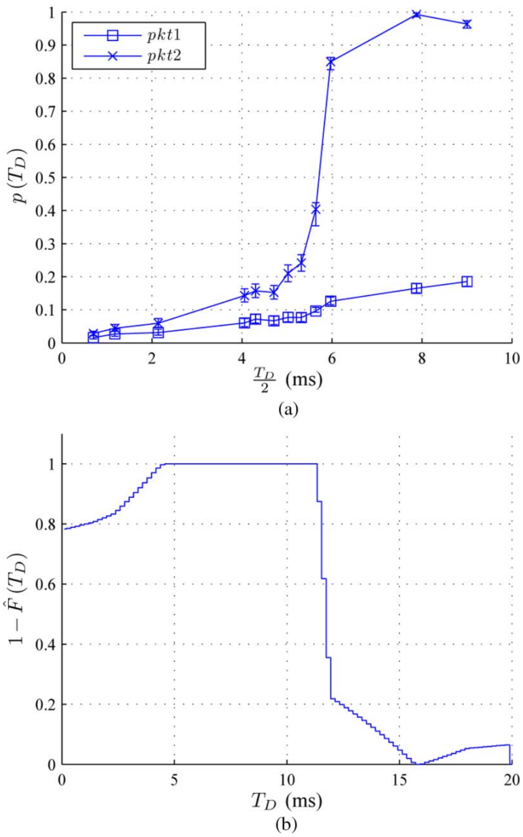

Fig. 6. Experimental measurements with MWO interference. Data frames are transmitted at a PHY rate of 1-Mb/s rate, and the duration $T_{\rm D}$ is varied by adjusting the packet size. Both $\rm pkt1$ and $\rm pkt2$ are equal-length $T_{\rm D}/2$. (a) Measured packet loss rate versus packet duration $T_{\rm D}$. Confidence intervals based on the Clopper–Pearson method are displayed, but are small enough to be partially obscured by the point markers. (b) Interarrival distribution of interference pulses.

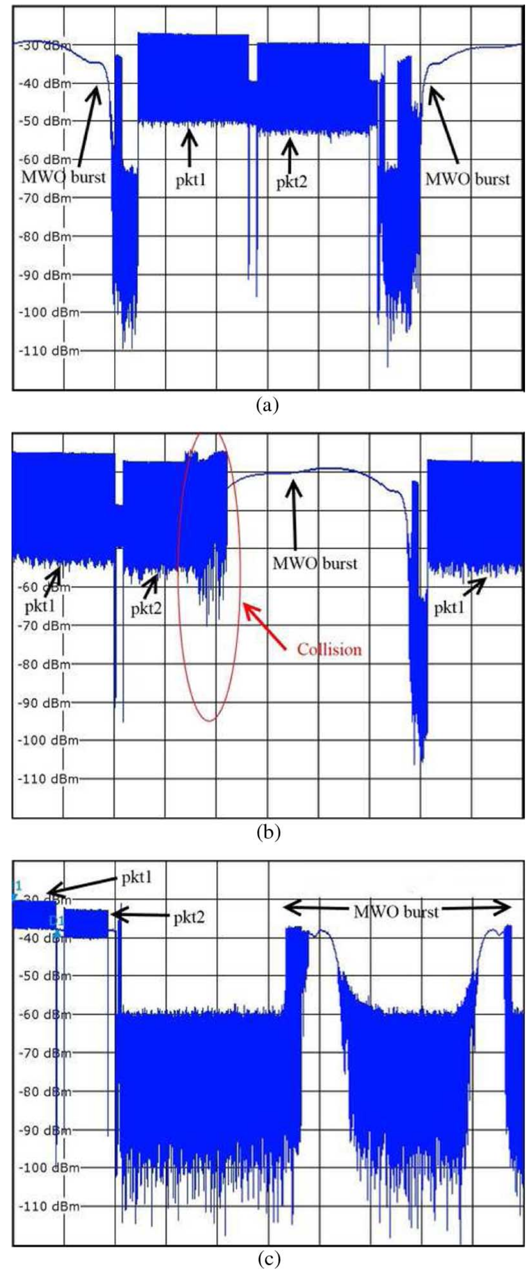

Fig. 7. Spectrum analyzer measurements of MWO interference. (a) Packet pair transmitted between two MWO bursts. The y-axis grid is in 2-ms increments. The packet pair is encoded at the 1-MB/s 802.11 rate, with both packets having duration 4.36 ms. (b) Second packet in a pair suffering a collision with 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.

direct spectrum analyzer 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) to 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 $F(T_{\rm D})$.

3) Validation: Fig. 7(a) presents spectrum analyzer 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 of 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 [32]. Our measurements indicate that the MWO interference consists of pulses with mean interval of 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 testbed consists of a WLAN formed from two client stations and an access point, plus three additional stations configured as hidden nodes (HNs). These HNs are created by modifying the Madwifi driver such that the carrier sense is disabled (using the technique as detailed in [33]) 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 to 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 analyzer). 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.

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 interarrival times estimated using this packet loss data. The data plotted in Fig. 8(b) is the estimate of $1 - F(T_{\mathrm{D}})$ and is displayed using a logarithmic y-axis. Also plotted in Fig. 8(b) is the theory line $1 - F(T_{\mathrm{D}}) = e^{-\lambda T_{\mathrm{D}}}$ 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 nonparametric manner. By making stronger structural assumptions about the interference process, we can alternatively parameterize the distribution function, and our task then becomes one of estimating these model parameters. A fairly direct tradeoff in effort is involved here, which is why it is important to consider both nonparametric and parametric approaches. Namely, we have the bias–variance tradeoff whereby nonparametric 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

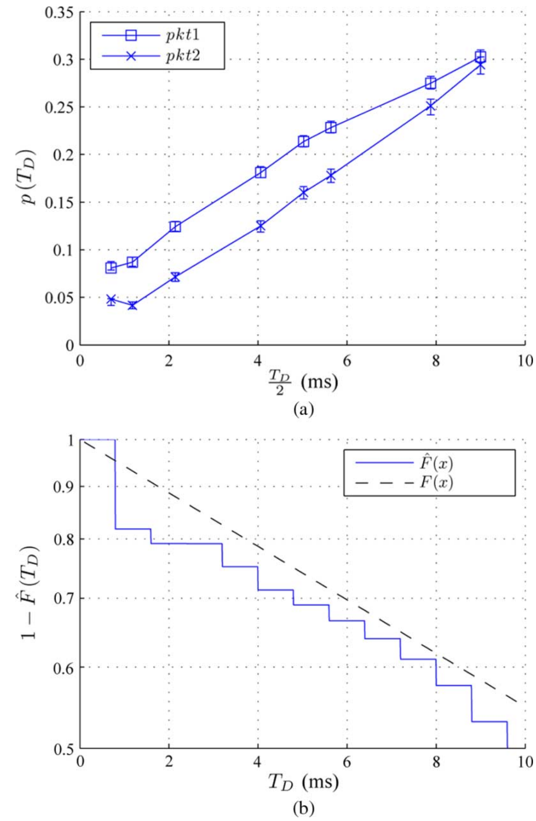

Fig. 8. Experimental measurements; primary network has two nodes transmitting to AP, interference network has three hidden nodes. (a) Measured packet loss rate versus packet duration T_{D}. Confidence intervals based on the Clopper–Pearson method are displayed, but are small enough to be partially obscured by the point markers. (b) Interarrival distribution of interference pulses.

Gilbert–Elliot channel model [34], which is popular for analyzing 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 multistate interference model [35], correlated losses [36], fast fading [37], 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

$$Q = \begin{bmatrix} -\lambda_B & \lambda_B \\ \lambda_G & -\lambda_G \end{bmatrix} \tag{7}$$

and

$$\Pi = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}. \tag{8}$$

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

$$\mathrm{Prob}[Y_{n+1} = j | Y_n = i] = \Pi_{ij}. \tag{9}$$

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

With our the choice of Π, the Y_n flip back and forth between the G and B states so that $\mathcal{Y}$ is of the form $\{\ldots, G, B, G, B, \ldots\}$. Let $\{k\}$ index the subsequence of B states in $\mathcal{Y}$. Let S_k denote the dwell time in the kth 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 $\{V_m, m = 0, 1, 2, \ldots\}$ denote a sequence of random variables that flip back and forth between the Tx and $Idle$ states. The dwell time in the Tx state is a constant T_{D}, and the dwell times in the $Idle$ state are independent exponential random variables with mean $1/\lambda_D$. We index the subsequence 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

$$\begin{aligned} p_{\mathrm{cs}} &= \mathrm{Prob}\left[t_n \in \left[T_{k(n)}, T_{k(n)} + \Delta_{k(n)}\right]\right] \\ &:= \alpha \frac{\lambda_G}{\lambda_G + \lambda_B} \end{aligned}$$

where $0 \leq \alpha \leq 1$ and $\lambda_G/(\lambda_G + \lambda_B)$ is the probability that the interference is in state B. In the following, we consider two limiting situations. First is 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 $\alpha = 1$. Second is 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 $\mathrm{Prob}[\delta_n = 1]$ that packet n is erased depends on $\tilde{S}_n$. Nevertheless, to streamline the presentation, we make the simplifying assumption that $\mathrm{Prob}[\delta_n = 1] = p_B$ whenever $\tilde{S}_n > 0$, and $\mathrm{Prob}[\delta_n = 1] = p_G$ 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 $n \neq m$.

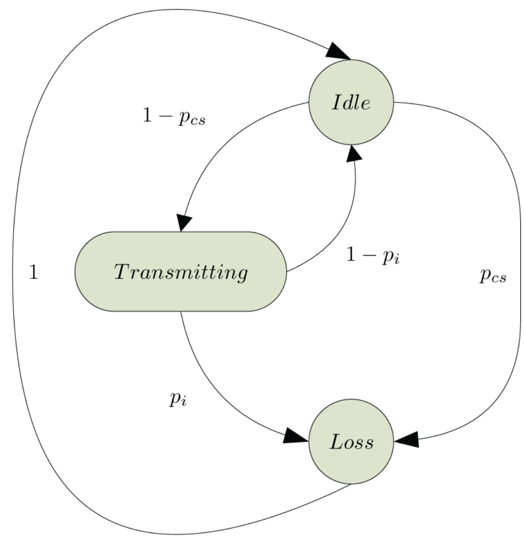

Fig. 9. Slotted-time Markov chain.

5) Packet Error Rate Analysis: To determine the packet error rate as a function of the packet transmit duration, we need to analyze 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

$$\mathbf{P} = \begin{bmatrix} 0 & 1 - p_{\mathrm{cs}} & p_{\mathrm{cs}} \\ 1 - p_{\mathrm{i}}(T_{\mathrm{D}}) & 0 & p_{\mathrm{i}}(T_{\mathrm{D}}) \\ 1 & 0 & 0 \end{bmatrix} \tag{10}$$

where $1 - p_{\mathrm{i}}(T_D) = \exp(-\lambda_B T_{\mathrm{D}})$. The stationary state distribution satisfies $\pi = \pi\mathbf{P}$, where $\pi_1 = \mathrm{Prob}[Idle]$, $\pi_2 = \mathrm{Prob}[Transmitting]$, and $\pi_3 = \mathrm{Prob}[Loss]$. Solving yields

$$\pi^T = \frac{1}{2 + p_{\mathrm{i}}(T_{\mathrm{D}})(1 - p_{\mathrm{cs}})} \begin{bmatrix} 1 \\ 1 - p_{\mathrm{cs}} \\ (1 - p_{\mathrm{cs}})p_{\mathrm{i}}(T_{\mathrm{D}}) + p_{\mathrm{cs}} \end{bmatrix}.$$

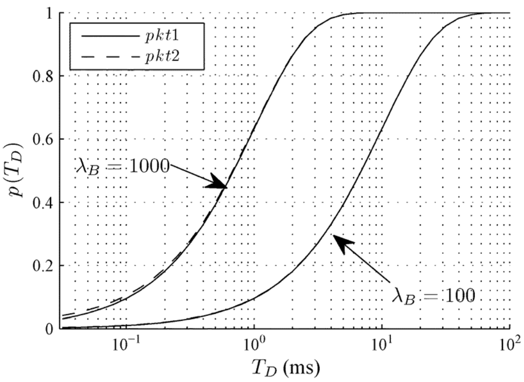

Fig. 10. Packet error rate versus packet duration $T_{\rm D}$; $\lambda_{\rm D} = 30$, variable λ_B, $p_G = 0$, $p_B = 1$, $p_{\rm cs} = 0$, $T_{\rm SIFS} = 10\ \mu$s.

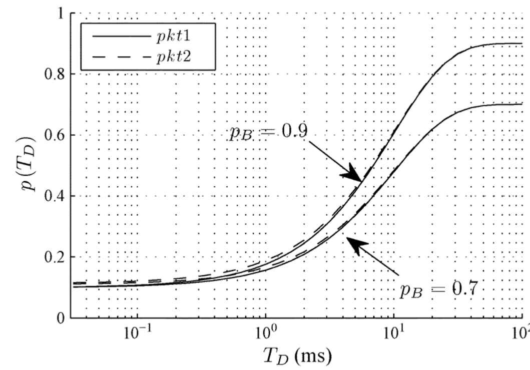

Fig. 11. Packet error rate versus packet duration $T_{\rm D}$; $\lambda_{\rm D} = 30$, $\lambda_B = 100$, $p_G = 0.1$, variable p_B, $p_{\rm cs} = 0$, $T_{\rm SIFS} = 10\ \mu$s.

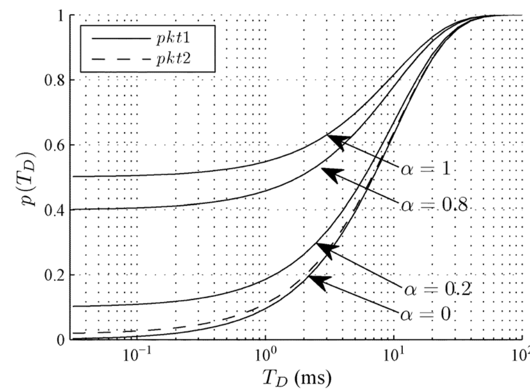

Fig. 12. Packet error rate versus packet duration $T_{\rm D}$; $\lambda_{\rm D} = 30$, $\lambda_B = 100$, $p_G = 0$, $p_B = 0$, variable $p_{\rm cs}$ (by varying α), $T_{\rm SIFS} = 10\ \mu$s.

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

$$\begin{aligned} p_1(T_{\rm D}) &= \frac{(1-p_{\rm i}(T_{\rm D}))\,\pi_2 p_G + p_{\rm i}(T_{\rm D})\pi_2 p_B + p_{\rm cs}\pi_1 p_B}{(1-p_{\rm i}(T_{\rm D}))\,\pi_2 + p_{\rm i}(T_{\rm D})\pi_2 + p_{\rm cs}\pi_1} \\ &= (1-p_{\rm i}(T_{\rm D}))\,(1-p_{\rm cs})p_G \\ &\quad + (p_{\rm i}(T_{\rm D})(1-p_{\rm cs}) + p_{\rm cs})\,p_B \\ &= :G_1(T_{\rm D}, \lambda_B, p_B, p_G, p_{\rm cs}). \end{aligned} \tag{11}$$

The first term in the expression for $p_1(T_{\rm D})$ 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

$$\begin{aligned} p_2(T_{\rm D}) &= (1-p_{\rm i}(T_{\rm D}))\,\frac{\lambda_B}{\lambda_B+\lambda_G} p_G \\ &\quad + \left(1-(1-p_{\rm i}(T_{\rm D}))\,\frac{\lambda_B}{\lambda_B+\lambda_G}\right) p_B \\ &= :G_2(T_{\rm D}, \lambda_B, p_B, p_G, p_{\rm cs}) \end{aligned} \tag{12}$$

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

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 $p_{\rm cs}$. Before proceeding, we briefly illustrate how the model parameters λ_B, p_B, p_G, and $p_{\rm cs}$ affect the observed packet loss versus $T_{\rm D}$ curves. Our aim is to: 1) illustrate the types of loss curves that the model is able to capture; and 2) 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 $p_{\rm cs}$ (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 $p_{\rm cs}$ from measurements of packet loss. The empirical estimators for loss probabilities $p_1(T_{\rm D})$ and $p_2(T_{\rm D})$ are

$$\hat{p}_1(T_{\rm D}) = \frac{1}{N_1}\sum_{n=1}^{N_1}\delta_n^1 \qquad \hat{p}_2(T_{\rm D}) = \frac{1}{N_2}\sum_{n=1}^{N_2}\delta_n^2$$

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 the nth 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 $T_{{\rm D}_1}, T_{{\rm D}_2}, \ldots$ and stacking the corresponding loss probability estimates, we have

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

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

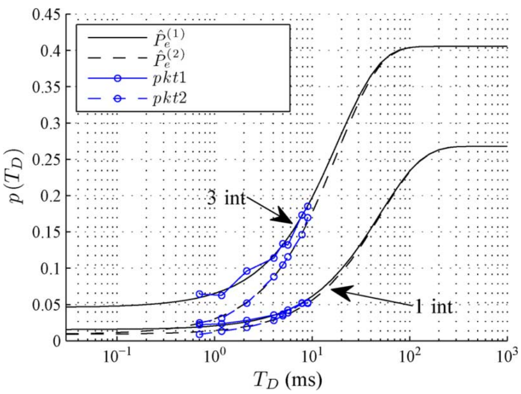

Fig. 13. Experimental measurements and model fit for WLAN with hidden node interference. Data points are for experiments using one and three interferers, with each interferer having a packet transmission rate of $\lambda_B = 20$. Initial values for the parameter estimator are $\hat{\lambda}_B = 20$,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

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

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

for parameters λ_B, p_B, p_G, and $p_{\rm cs}$ are then the values that minimize 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 $T_{\rm D}$. 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 1.4–18 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 corresponding model parameter estimates are given in Table III, obtained using an interior-point solver.

3) Validation: The model parameters that need to be estimated are $p_G, p_B, p_{\rm cs}$, and λ_B. In our experiments, we control the hidden terminal transmitters, and so we know the true value of λ_B. Namely, 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. The value of p_G (the packet loss probability when there are no interference transmissions) is determined by the physical channel properties.

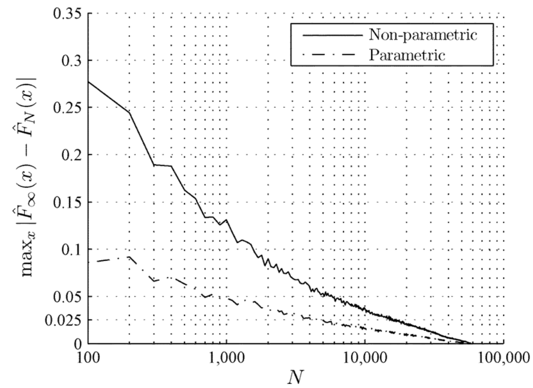

Fig. 14. Convergence of estimates of $F(x)$ versus the number of packets observed. $\hat{F}_N(x)$ denotes the estimate using N packet observations and $\hat{F}_\infty(x)$ denotes the estimate obtained using the full measurement trace. For each N, we take 100 random subsamples of N packets from the full 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 nonparametric estimates. The data set used is from the three interferer experiment; see Fig. 13.

We performed separate measurements without interference and found the packet loss rate to be less than 1%, and it can be seen from Table III that the model estimate for p_G is in good agreement with this (i.e., agrees to within experimental error). The parameter values that we do not fully validate are p_B and $p_{\rm cs}$. However, we note that they are obtained as part of a coupled model, which means that the values obtained are consistent with the accurate estimates obtained for the remaining parameters. Additionally, the estimate of p_B was partially validated using separate experimental tests where we operated the hidden terminals at a high send rate and measured the fraction of packets lost and obtained results consistent with the estimated values of p_B. Observe that p_B increases with the number of hidden terminals—we believe that this increase is genuine and occurs due to the additive nature of the hidden terminal transmissions, i.e., with three transmitters, there is some chance now that two or even three pulses from individual transmitters will overlap/coincide, creating a greater level of packet loss on the measured link.

4) Parametric Versus Nonparametric Estimation: A parametric model makes strong structural assumptions that allow the loss curves to be parameterized 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 nonparametric 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 (11) and (12), and the resulting parameterized 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 nonparametric solution requires

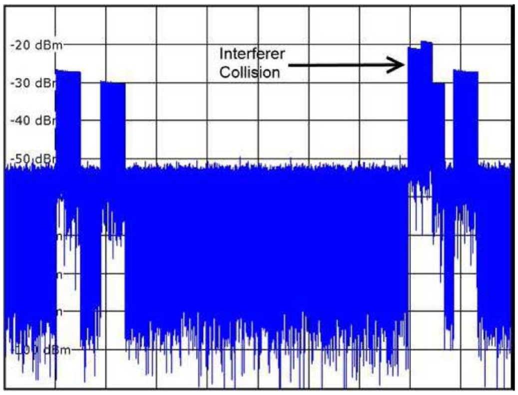

Fig. 15. Spectrum analyzer snapshot of hidden terminal interferers in time. The y-axis grid is in 2-ms increments. Interferer burst durations are fixed at 4.5 ms, with arrivals at 10, 19, 80, 83, and 89 ms. Since each interferer has a different path to the spectrum analyzer antenna, the pulses are at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

$N = 6000$ and $N = 20\,000$ 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 behavior 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 analyzer 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. Conclusion

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] and [2], which establishes a MAC/PHY cross-layer technique capable of classifying lost transmission opportunities into noise-related losses, collision induced losses, and hidden-node losses and of distinguishing these losses from the unfairness caused by exposed nodes and capture effects.

Acknowledgment

Many thanks to K. Duffy and G. 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 Trans. Netw.*, 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 *Proc. 18th IEEE Proc. WiOpt*, 2010, pp. 412–417.

[3] K. Rele, D. Roberson, B. Zhang, L. Li, Y. B. Yap, T. Taher, D. Ucci, and K. Zdunek, "A two-tiered cognitive radio system for interference identification in 2.4 GHz ISM band," in *Proc. 7th IEEE 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. Mobile Comput.*, vol. 1, no. 4, pp. 278–292, Oct.–Dec. 2002.

[5] I. Haratcherev, K. Langendoen, R. Lagendijk, and H. Sips, "Hybrid rate control for ieee 802.11," in *Proc. 2nd ACM MobiWac*, 2004, pp. 10–18.

[6] T. Zhou, X. Wang, and W. Hou, "A fast collision detection algorithm in IEEE 802.11 through physical layer SINR monitoring," in *Proc. 73rd IEEE VTC*, May 2011, pp. 1–5.

[7] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11b mesh network," in *Proc. SIGCOMM*, 2004, pp. 121–132.

[8] S. Rayanchu, A. Mishra, D. Agrawal, S. Saha, and S. Banerjee, "Diagnosing wireless packet losses in 802.11: Separating collision from weak signal," in *Proc. 27th Annu. IEEE INFOCOMM*, Apr. 2008, pp. 735–743.

[9] M. Vutukuru, H. Balakrishnan, and K. Jamieson, "Cross-layer wireless bit rate adaptation," *Comput. Commun. Rev.* vol. 39, pp. 3–14, Oct. 2009.

[10] J. M. , III and G. Q. M. , Jr., "Cognitive radio: Making software radios more personal," *IEEE Pers. Commun.*, vol. 6, no. 4, pp. 13–18, Aug. 1999.

[11] A. Kamerman and L. Monteban, "WaveLAN-II: A high-performance wireless LAN for the unlicensed band," *Bell Labs Tech. J.*, vol. 2, no. 3, pp. 118–133, Autumn, 1997.

[12] G. Holland, N. Vaidya, and P. Bahl, "A rate-adaptive MAC protocol for multi-hop wireless networks," in *Proc. MobiCom*, Jul. 2001, pp. 236–251.

[13] S. H. Y. Wong, H. Yang, S. Lu, and V. Bharghavan, "Robust rate adaptation for 802.11 wireless networks," in *Proc. 12th Annu. ACM MobiCom*, 2006, pp. 146–157.

[14] Q. Pang, V. C. M. Leung, and S. C. Liew, "A rate adaptation algorithm for IEEE 802.11 WLANs based on MAC-layer loss differentiation," in *Proc. 2nd BroadNets*, Oct. 2005, pp. 659–667.

[15] J. Kim, S. Kim, S. Choi, and D. Qiao, "CARA: Collision-aware rate adaptation for IEEE 802.11 WLANs," in *Proc. 25th Annu. IEEE INFOCOMM*, Apr. 2006, pp. 1–11.

[16] D. Malone, P. Clifford, D. Reid, and D. J. Leith, "Experimental implementation of optimal WLAN channel selection without communication," in *Proc. IEEE DySPAN*, 2007, pp. 316–319.

[17] M. Khan and D. Veitch, "Isolating physical PER for smart rate selection in 802.11," in *Proc. 28th Annu. IEEE INFOCOMM*, Apr. 2009, pp. 1080–1088.

[18] M. N. Krishnan, S. Pollin, and A. Zakhor, "Local estimation of probabilities of direct and staggered collisions in 802.11 wlans," in *Proc. IEEE GLOBECOM*, Dec. 2009, pp. 1–8.

[19] E.-I. Kim, 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.*, vol. 51, no. 1, pp. 1–8, 2003.

[20] J. Yin, X. Wang, and D. P. Agrawal, "Optimal packet size in error-prone channel for IEEE 802.11 distributed coordination function," in *Proc. IEEE WCNC*, Mar. 2004, vol. 3, pp. 1654–1659.

[21] X. He, F. Y. Li, and J. Lin, "Link adaptation with combined optimal frame size and rate selection in error-prone 802.11n networks," in *Proc. IEEE ISWCS*, 2008, pp. 733–737.

[22] S. Choi and K. G. Shin, "A class of adaptive hybrid ARQ schemes for wireless links," *IEEE Trans. Veh. Technol.*, vol. 50, no. 2, pp. 777–790, May 2001.

[23] S. Ci and H. Sharif, "An link adaptation scheme for improving throughput in the IEEE 802.11 wireless LAN," in *Proc. 27th Annu. IEEE LCN*, Nov. 2002, pp. 205–208.

[24] P. Lettieri and M. B. Srivastava, "Adaptive frame length control for improving wireless link throughput, range, and energy efficiency," in *Proc. 17th Annu. IEEE INFOCOM*, Mar. 1998, vol. 2, pp. 564–571.

[25] M. Krishnan, E. Haghani, and A. Zakhor, "Packet length adaptation in WLANs with hidden nodes and time-varying channels," in *Proc. IEEE GLOBECOM*, 2011, pp. 1–6.

[26] R. W. Wolff, "Poisson arrivals see time averages," *Oper. Res.*, vol. 30, no. 2, pp. 223–231, 1982.

[27] B. Melamed and W. Whitt, "On arrivals that see time averages," *Oper. Res.*, vol. 38, no. 1, pp. 156–172, 1990.

[28] F. Baccelli, S. Machiraju, D. Veitch, and J. Bolot, "The role of PASTA in network measurement," *IEEE/ACM Trans. Netw.*, vol. 17, no. 4, pp. 1340–1353, Aug. 2009.

[29] D. Malone, K. Duffy, and D. J. Leith, "Modeling the 802.11 distributed coordination function in nonsaturated heterogeneous conditions," *IEEE/ACM Trans. Netw.*, vol. 15, no. 1, pp. 159–172, Feb. 2007.

[30] 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," *Comput. Netw.* vol. 53, no. 2, pp. 141–152, 2009.

[31] K. D. Huang, K. R. Duffy, and D. Malone, "On the validity of IEEE 802.11 MAC modeling hypotheses," *IEEE/ACM Trans. Netw.*, vol. 18, no. 6, pp. 1935–1948, Dec. 2010.

[32] A. Kamerman and N. Erkocevic, "Microwave oven interference on wireless LANs operating in the 2.4 GHz ISM band," in *Proc. IEEE PIMRC*, Sep. 1997, vol. 3, pp. 1221–1227.

[33] E. Anderson, G. Y. , C. Phillips, D. Sicker, and D. Grunwald, "Commodity AR52XX-based wireless adapters as a research platform," Dept. Comput. Sci. Univ. Colorado at Boulder, Tech. Rep. CU-CS-XXXX-08, Apr. 2008.

[34] M. Mushkin and I. Bar-David, "Capacity and coding for the Gilbert-Elliot channels," *IEEE Trans. Inf. Theory*, vol. 35, no. 6, pp. 1277–1290, Nov. 1989.

[35] S. Sivaprakasam and K. S. Shanmugan, "An equivalent Markov model for burst errors in digital channels," *IEEE Trans. Commun.*, vol. 43, no. 2/3/4, pp. 1347–1355, Feb./Mar./Apr. 1995.

[36] M. Yajnik, S. Moon, J. Kurose, and D. Towsley, "Measurement and modeling of the temporal dependence in packet loss," in *Proc. 18th Annu. IEEE INFOCOMM*, Mar. 1999, vol. 1, pp. 345–352.

[37] H. S. Wang and N. Moayeri, "Finite-state Markov channel-a useful model for radio communication channels," *IEEE Trans. Veh. Technol.*, vol. 44, no. 1, pp. 163–171, Feb. 1995.

Brad W. Zarikoff (S'00–M'09) received the B.Eng. degree with distinction in electrical engineering from the University of Victoria, Victoria, BC, Canada, in 2002 and the M.A.Sc. and Ph.D. degrees in engineering science from Simon Fraser University, Burnaby, BC, Canada, in 2004 and 2008, respectively.

He is currently a Research Fellow with the Hamilton Institute, National University of Ireland Maynooth, Maynooth, Ireland. His current research interests include interference mitigation, power-line communication networks, and synchronization for network MIMO systems.

Doug J. Leith (M'02–SM'09) graduated from the University of Glasgow, Glasgow, U.K., in 1986, where he received the Ph.D. degree in engineering in 1989.

In 2001, he moved to the National University of Ireland Maynooth, Maynooth, Ireland, to assume the position of SFI Principal Investigator and to establish the Hamilton Institute (http://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.