

A Non-Parametric Approach to Estimating Ambient Noise Levels in the Presence of Bursty Interference

Brad Zarikoff, Douglas J. Leith

Hamilton Institute, National University of Ireland Maynooth

Abstract—We propose a new transmitter-side approach for estimating the contribution to the packet error rate that is due to background noise, distinct from the contribution due to interference bursts. The technique relies solely on an existing data-ack handshake and the transmitters ability to adjust or monitor the packet size. One immediate application of this information is in rate adaptation. Experimental measurements with microwave oven interference are presented to demonstrate the practical utility of the proposed technique.

Index Terms—Interference, 802.11, CSMA/CA.

I. INTRODUCTION

Emissions in the ISM unlicensed bands are often *pulsed* in nature, and so wireless links using these bands may experience both channel noise and bursty interference. Because of the pulsed nature of the interferers, the SINR measured for one packet may bear little relation to the SINR experienced by other packets. Consequently, rate adaptation algorithms that rely on accurate estimates of the channel error probability via the SINR will perform erratically, converging to values far from optimal. In this paper we present a new approach for estimating the underlying packet loss probability due to noise, which is a function of the channel gain and receiver noise, despite the presence of pulsed interferers.

On wireless links that use ARQ (*e.g.* 802.11), exchanges occur via a data+ACK two-way handshake. As a result, the interference environments at both the receiver and transmitter must be accounted for when measuring link quality, and so receiver-side measurements alone are insufficient. Further, since receive signal strength indicator (RSSI) measurements are based on the estimated noise level within a small section of a frame (*e.g.* the PHY preamble in 802.11a frames on which RSSI estimates are based is 20 μ s long, but the frame may be milliseconds long), RSSI alone is also insufficient as conditions may change during transmission of a frame. We therefore choose to work directly with the data+ACK handshake and measure the proportion of data+ACK handshakes which are not successfully completed. The basic idea is to measure how this proportion changes as the duration of the transmitted data packets is varied. Surprisingly, it turns out that this information is sufficient to allow many characteristics of the pulsed interferers affecting a link to be accurately inferred. Our interest here is in estimating the contribution to packet loss due to noise as opposed to pulsed interference. Since we cannot isolate the link from the interferers, this is




Fig. 1. The interference can be made up of any uncooperative co-channel bursty transmission. For example, a network of similar wireless radios that are operating as hidden nodes, or a microwave oven.

not a straightforward task. This extends recent work in [1] which establishes a packet size sampling technique capable of estimating the distribution of inter-arrival times between interference pulses.

II. RELATED WORK

In a contention-based multi-access channel such as an 802.11 link, a packet error may be caused by three distinct sources: noise, a MAC collision, and an interference collision.¹ Methods for identifying the source of packet errors in 802.11 links are a relatively recent topic of study, prior to which rate control algorithms relied upon estimates from combined channel and MAC error sources. In [2], a receiver-side technique is presented, where collision (lumped MAC and interference) and noise mean statistics are returned via a feedback channel. The channel feedback can be avoided by implicitly relying on first packet ACKs in packet bursts [3] using the TXOP extension in 802.11e/n networks. An extension to [2] that considers packet length adaptation to maximise link throughput is presented in [4]. It relies on a joint packet length optimisation over MAC collisions and channel errors. Our recent work in [1] provides a more detailed view of the interference environment, and uses a parameterised model to separate the three aforementioned error sources, and a non-parameterised model to estimate the distribution of the inter-arrival time between interference bursts. However, this nonparametric approach assumes that losses due to noise can be neglected, and in the current paper we relax this assumption.

III. SYSTEM DEFINITION

We consider a system made up of a communications link d_0 and some unknown number N of interferers d_1 through d_N . The interferers emit RF energy in bursts; we make no distinction if the interferer burst is data modulated or not. A

¹We consider MAC collisions to originate from non-hidden nodes, while interference collisions originate from hidden nodes and heterogeneous technologies.

sequence of packets $p \in \{1, 2, 3, \dots\}$ are transmitted across link d_0 and are subject to corruption by noise and by bursts from the interferers. Our interest is in estimating the error rate for data+ACK handshakes on d_0 due to noise only *i.e.* to estimate the underlying error probability in the absence of interference.

For link d_0 , we lump the multiple interferers into a single equivalent interferer, denoted d_I . Define the duration of the k^{th} transmission from d_I as $S_k > 0$, and the start time of the k^{th} transmission as T_k . The interval between the end of the k^{th} transmission and the beginning of the $(k+1)^{\text{th}}$ transmission is then $\Delta_k = T_{k+1} - (T_k + S_k)$. We assume that the interferers cannot decode or detect transmissions from d_0 , and so the transmission times T_k are independent of transmissions made by d_0 (*i.e.* the interferers act as hidden nodes). We also assume that S_k and Δ_k are both i.i.d. and let $F(\Delta)$ denote the distribution function of the Δ_k 's. Under the further assumptions that

- (i) a data+ACK handshake transmitted on d_0 is lost when its transmission overlaps with an interference pulse,
- (ii) a data+ACK handshake is successful if it does not overlap with an interference pulse, and
- (iii) the interval between each data+ACK transmission on d_0 is exponentially randomly distributed and independent of the interference process,

it has been shown in [1] that the probability that a data+ACK transmission on d_0 overlaps with a burst from d_I is

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

where T_D is the duration of the data+ACK transmission on d_0 and $F_c(\Delta) = 1 - F(\Delta)$ is the ccdf of Δ .

Assumption (i) is mild and involves little loss of generality since it amounts to simply ignoring interference bursts that do not lead to packet loss.² As noted in [1], assumption (iii) can be replaced by the weaker assumption that the sampling approach satisfies the Arrivals See Time Averages property.

However, assumption (ii) is violated in the problem considered here since noise can lead to packet loss even in the absence of interference. We therefore proceed as follows. Let δ_p denote a random variable that has value 1 if data+ACK transmission p transmitted on link d_0 is corrupted due to noise, and value 0 otherwise. Assume that the δ_p , $p \in \{1, 2, 3, \dots\}$ are i.i.d. and independent of the interference bursts. Also, assume that δ_p is independent of the data+ACK transmission duration T_D : experimental results in Section VI show that for small T_D this is a good approximation. Letting $p_G = \text{Prob}[\delta_p = 1]$, we then have

$$p(T_D) = 1 - (1 - p_G)(1 - q(T_D)), \quad (2)$$

where $p(T_D)$ is the probability that a data+ACK transmission on d_0 overlaps with an interference burst and/or is corrupted

due to noise. From here on, we refer to $p(T_D)$ as the packet error probability (PEP) as it includes elements from noise and interference. Note that both $p(T_D)$ and p_G will be a function of the specific modulation and coding scheme (MCS) used for transmissions on d_0 . For now, we assume a fixed MCS. Our objective is to estimate p_G using only measurements of $p(T_D)$ for a range of T_D values.

Useful Mechanisms of the 802.11 MAC: We briefly comment on two features of the 802.11 MAC that will prove important. The first is carrier sense (CS), whereby directly prior to any transmission a station compares a sample of the channel energy to a CS threshold. If the energy is above the threshold, it is presumed that another transmission is currently taking place, and the device will defer transmission. When the energy in an interference burst is above the CS threshold, the timing of station transmissions then becomes coupled to that of the interference bursts. A second feature is TXOP packet bursting. Let the duration of the first data+ACK transmission in a burst be T_{D1} and the duration of the second data+ACK transmission be T_{D2} . The second transmission takes place only if the first ACK is received. By assumption (i), this implies there was no interference burst or noise error event during the first data+ACK transmission. Neglecting the short interframe space (SIFS) between the first and second transmissions, we can therefore assume that the second transmission starts in the absence of an interference burst.

IV. JOINTLY ESTIMATING p_G AND $F(\Delta)$

Our primary focus in this paper is on estimating p_G , but we also briefly comment on jointly estimating p_G and $F(\Delta)$. Fig. 2 provides example curves of loss rate $p(T_D)$ versus transmit duration T_D for periodic interference pulses. Loss rate data is shown for the first data+ACK exchange in a TXOP burst and for the second exchange. Data is also shown with $p_G = 0$ and $p_G = 0.25$, to highlight the impact of p_G on $p(T_D)$. We can make a number of observations. Firstly, the loss rate for the second data+ACK exchange in a burst (dash-dot line) intercepts the y-axis at p_G . We can therefore estimate p_G from the y-axis intercept. Even if measurements for small T_D are not available, smooth extrapolation will still yield an estimate of p_G . For non-periodic interference, similar extrapolations can be done using knowledge of the structure of $p(T_D)$, *e.g.* smoothness of the $p(T_D)$ versus T_D curve. Secondly, comparing Fig. 2(a) and Fig. 2(b) it can be seen that carrier sense has a significant impact on $p(T_D)$. When the energy of interference pulses is above the CS threshold, it can be seen from Fig. 2(a) that the loss rate for first data+ACK exchange in a burst also intercepts the y-axis at p_G and so can be used to estimate p_G . Once we have an estimate of p_G , $F(\Delta)$ can be estimated from this loss rate data using the approach in [1].

V. ESTIMATING p_G

Jointly estimating p_G and $F(\Delta)$ requires measurements of $p(T_D)$ over a range of T_D values. In this section we investigate use of a single sample point to estimate p_G , which potentially allows much faster estimation when p_G alone is

²This is based on the observation that if an interference burst overlaps with a data+ACK transmission on d_0 and prevents a successful transmission, then it will appear to the transmitter as a loss. However, if an interference burst overlaps and a success is declared, it will appear to the transmitter as if there was no interference burst. This amounts to an implicit partition or power threshold, below which low level interference is lumped together with the channel noise.




Fig. 2. Loss rate versus transmit duration T_D for periodic interference pulses with period 11 ms and channel loss rate $p_G = 0.25$. This interference is representative of RF emissions from domestic microwave ovens (MWOs), see later. Loss rate data is shown for both first (dash line) and second (dash-dot line) data+ACK exchanges in a TXOP burst, and when the interference energy is above and below the carrier sense threshold. For comparison, data is also shown when the channel loss rate $p_G = 0$ (solid line).

required (e.g. for rate adaptation). It is evident from (2) that we could directly estimate p_G if there were some T_D where $q(T_D) = 0$, which is true regardless of the shape of $F(T_D)$. Since the distribution function $F(\Delta)$ is monotonically increasing, the smallest $q(T_D)$ we can get is when $T_D \rightarrow 0$. However, $\lim_{T_D \rightarrow 0} q(T_D) = 1 - \mathbb{E}[\Delta]/\mathbb{E}[S + \Delta]$ (as $\lim_{T_D \rightarrow 0} F_c(T_D) = 1$). Since we cannot depend on $\mathbb{E}[S] \ll \mathbb{E}[\Delta]$, nor estimate $\mathbb{E}[S]$ and $\mathbb{E}[\Delta]$, we cannot use this direct approach to accurately estimate p_G . We therefore consider alternative approaches for estimating p_G .

Packet Pair Analysis: Let p_1 denote the loss rate of the first data+ACK exchange in a TXOP burst and p_2 the loss rate of the second exchange. Then $p_1(T_{D1}) = p(T_{D1})$ and $p_2(T_{D1}, T_{D2}) = 1 - (1 - p_G)(1 - p(T_{D1} + T_{D2})) / (1 - p(T_{D1}))$, where we are using the fact that the time between pulses is strictly greater than 0 (else we conjoin the pulses). Then, $\lim_{T_{D2} \rightarrow 0} p_2(T_{D1}, T_{D2}) = p_G$. In this case, we are essentially using the first data+ACK exchange as a carrier sense and the loss rate of second exchange $p_2(T_{D1}, 0)$ is an estimate of p_G . It is also straightforward to include the probability of a collision of first exchange with other synchronised 802.11 links, p_C , by redefining $p_1(T_{D1})$ as $1 - (1 - p(T_{D1}))(1 - p_C)$. However, we will assume that $p_C = 0$ in the following for simplicity.

Carrier Sense Analysis: The second approach we consider is applicable when the carrier sense threshold is set below the received power of the interference bursts. We can include a correction factor in (1), which corrects for pulses that arrive during an interference burst, to get $p_{CS}(T_D) = 1 - \frac{1 - p_G}{\mathbb{E}[S + \Delta]} \left(\mathbb{E}[S]F_c(T_D) + \int_{T_D}^{\infty} F_c(x)dx \right)$. Similarly to the TXOP burst analysis, $\lim_{T_D \rightarrow 0} p_{CS}(T_D) = p_G$. Therefore, we can use $p_{CS}(0)$ as an estimate of p_G .

A. Estimation Bias

The loss probability $p_2(T_D)$ can be estimated using data+ACK measurements as $\hat{p}_2(T_D) = \frac{N_2(T_D)}{N(T_D)}$ where $N(T_D)$ is the number of second exchanges of duration T_D that are attempted and $N_2(T_D)$ is the number of these exchanges for which an ACK is not successfully received. By the law of large numbers this estimator is consistent and unbiased as $N(T_D) \rightarrow \infty$ and the variance converges as $O(\sqrt{N})$. Similarly for estimating $p_{CS}(T_D)$. However, we cannot select T_D or T_{D2} to be zero since practical packet data systems have




Fig. 3. Admissible values of normalised interference rate $\lambda\epsilon$ and channel loss rate p_G for specified levels of maximum estimation bias ρ_{\max} . Interference pulses with exponential inter-arrival times.

lower packet size limits based on protocol overheads (63 μs for ERP-OFDM 802.11, 89.5 μs for DSSS-OFDM, and 126 μs for Bluetooth). Instead we have $T_D \geq \epsilon$, for some $\epsilon > 0$, and this will introduce a bias when estimating p_G . Specifically, we have

$$\hat{p}_G = p_2(\epsilon) = 1 - (1 - p_G)b(\epsilon), \quad (3)$$

where $b(\epsilon)$ is the bias, with $b(0) = 1$. For any $b(\epsilon) < 1$, we can define the error of our estimate \hat{p}_G as $\rho = |p_G - \hat{p}_G|$. Replacing \hat{p}_G with (3), the estimate error is

$$\rho = |(1 - p_G)(1 - b(\epsilon))|. \quad (4)$$

To gain some qualitative insight, we now quantify the magnitude of this bias for two representative examples.

Non-periodic Interference: Suppose that the inter-arrival times Δ between interference pulses are exponentially distributed, with mean inter-arrival time λ . Then for both $p_2(\epsilon)$ and $p_{CS}(\epsilon)$ the bias $b(\epsilon)$ is the same and equal to

$$b(\epsilon) = e^{-\lambda\epsilon}. \quad (5)$$

If we specify a maximum permissible bias error ρ_{\max} , using (5) into (4) we obtain the upper bound $\lambda < -\epsilon^{-1} \log(1 - \rho_{\max}(1 - p_G)^{-1})$. Fig. 3 plots this upper bound in normalised form, $\lambda\epsilon$. It can be seen that as ρ_{\max} and p_G increase, so does the bound on $\lambda\epsilon$. Practically, these curves give an indication of the maximum interference arrival rate that meets a specific level of ρ_{\max} , e.g. we require an error of 1%, and the channel is known to have $p_G > 0.01$, then we need $\lambda\epsilon < 0.01$ (marked with an 'x' in Fig. 3).

Periodic Interference: Suppose that interference pulses are periodic with period $T_\Delta = \mathbb{E}[\Delta]$. Solving (3) for $b(\epsilon)$, we obtain the bias for TXOP bursting to be

$$b_2(\epsilon) = \frac{\max(T_\Delta - T_{D1} - \epsilon, 0)}{\max(T_\Delta - T_{D1}, 0)}, \quad (6)$$

and using carrier-sense

$$b_{CS}(\epsilon) = \frac{\mathbb{E}[S]F_c(\epsilon) + \max(T_\Delta - \epsilon, 0)}{\mathbb{E}[S + \Delta]}. \quad (7)$$

Observe that the interference periodicity penalises the second exchanges in a TXOP burst: for the first exchange, the maximum time until the next interference burst is T_Δ , while for the




Fig. 4. Admissible values of channel noise p_G for specified levels of maximum estimation bias ρ_{\max} . Data is shown for carrier sense based estimation ($b_{CS}(\epsilon)$), $\epsilon = (70, 700, 7000) \mu\text{s}$, and TXOP burst estimation ($b_2(\epsilon)$), $\epsilon = 70 \mu\text{s}$, $T_{D1} = (70, 700, 7000) \mu\text{s}$. Periodic interference with period $T_\Delta = \mathbb{E}[\Delta] = 11\text{ms}$ and pulse duration $\mathbb{E}[S] = 9 \text{ms}$.

second exchange it is $T_\Delta - T_{D1}$. Plots of the admissible p_G for a range of maximum permissible bias errors ρ_{\max} are shown in Fig. 4. For example, in order to achieve an accuracy of 1% when using carrier sense, the minimum transmit duration ϵ must be less than $200 \mu\text{s}$ for $p_G > 0$. For a data rate of 54 Mbps in 802.11a/g, this corresponds to packet payloads of 952 Bytes or less. It is also possible to use the packet pair approach for identifying p_G when carrier sense is enabled. Defining $T_{D\epsilon} = T_{D1} + \epsilon$, we get

$$b_{CS,2}(T_{D\epsilon}) = \frac{\mathbb{E}[S]F_c(T_{D\epsilon}) + \max(T_\Delta - T_{D\epsilon}, 0)}{\mathbb{E}[S + \Delta]}. \quad (8)$$

VI. EXPERIMENTAL MEASUREMENTS

We illustrate the practical utility of the proposed approach using experimental measurements.

Experimental Setup: The experimental setup consists of two stations (one client/one AP) and a 700 W Hinari MX745GLSL microwave oven (MWO). The MWO is operated at maximum power to heat a 1 L bowl of water, and is located approximately 1 m away from the stations. The stations are Asus 700 laptops equipped with Atheros 802.11 a/b/g chipsets (radio 14.2, MAC 8.0, PHY 10.2), running Debian Lenny 2.6.26 and using a modified Linux MadWifi driver based on 10.5.6 HAL and 0.9.4 driver. The beacon period is set to the maximum value of 1 s. Other complicating features such as adaptive noise immunity (ANI), antenna diversity, and rate control are disabled [5], [3]. The MWO operates in the 2.4 GHz ISM band, with significant overlap ($> 50\%$) with the WiFi 20 MHz channels 6 to 13 (2427 MHz to 2482 GHz); this was verified using a spectrum analyser. Our experiments used 802.11 channel 10 and took place in a room that was cleared for additional co-channel interference before, during and after each experiment. The client station initiates data+ACK exchanges with the AP using the standard ping command in a bash script. TXOP bursting is used to generate pairs of exchanges. Both exchanges in a pair are of the same duration T_D and this duration is adjusted by varying the ping size between 30 and 2110 bytes; for a modulation and coding rate of 54 Mbps, we used our spectrum analyser to measure the corresponding T_D values to be in the range $90.2 \mu\text{s}$ to




Fig. 5. Measured packet error rate versus packet duration T_D , for a 54 Mbps 802.11g link using p_{CS} and $p_{CS,2}$. These curves provide an experimental close-up of the carrier sense curves from Fig. 2(a) for some unknown p_G . Results are reported both with and without MWO interference.

$243 \mu\text{s}$. The interval between TXOP bursts is exponentially distributed with rate $\lambda = 30$. Carrier sense is enabled, with threshold -76dBm .

Results: Fig. 5 presents measurements of the packet loss rate versus packet duration T_D for a 54 Mbps link rate. As mentioned, carrier sense is enabled, so curves for p_{CS} and $p_{CS,2}$ are provided. Each point represents an average over 2000 packet transmissions. Note that the RF emissions from the MWO are orders of magnitude above the carrier sense threshold, and are periodic with inter-arrival time 11 ms and on-time approximately 9 ms [1]. Substituting (7) and (8) into (4), and calculating an average $p_G = 0.018$ in the absence of interference, we get $\rho_{CS} \approx 0.004$ and $\rho_{CS,2} \approx 0.009$, respectively, which indicates a bias error for either estimate of under 1%. Comparing the curves with and without MWO interference, it can be seen that the interference adds an additional 1% to 2% to the average $\hat{\rho}(T_D)$. This variation is within the confidence intervals. Convergence of \hat{p}_G is quick, and in our experiments converges to less than 2.5% in approximately 1200 packets.

VII. CONCLUSION

We consider a 802.11 wireless link subject to both channel noise and bursty interference and propose a new approach for estimating the contribution to the packet error rate that is due to channel noise, distinct from the contribution due to interference bursts. This approach is a transmitter-side technique that provides per-link information and is compatible with standard hardware. Experimental measurements with microwave oven interference are presented to demonstrate the practical utility of the proposed technique.

REFERENCES

- [1] B. W. Zarikoff and D. J. Leith, "Measuring Pulsed Interference in 802.11 Links," 2011, <http://arxiv.org/abs/1107.5468>.
- [2] M. N. Krishnan, S. Pollin, and A. Zakhor, "Local estimation of probabilities of direct and staggered collisions in 802.11 w lans," in *Proc. GLOBECOM*, 2009, pp. 1–8.
- [3] D. Giustiniano, D. Malone, D. J. Leith, and K. Papagiannaki, "Measuring Transmission Opportunities in 802.11 Links," *IEEE/ACM Trans Networking*, vol. 18, no. 5, pp. 1516–1529, 2010.
- [4] M. Krishnan and E. Haghani and A. Zakhor, "Packet Length Adaptation in WLANs with Hidden Nodes and Time-Varying Channels," in *submission to IEEE INFOCOM*, 2011.
- [5] K. D. Huang, K. R. Duffy, and D. Malone, "On the validity of IEEE 802.11 MAC modeling hypotheses," *IEEE/ACM Trans Networking*, vol. 18, no. 6, pp. 1935–1948, 2010.