

# E-CSMA: Supporting Enhanced CSMA Performance in Experimental Sensor Networks using Per-neighbor Transmission Probability Thresholds

Shane B. Eisenman

Electrical Engineering, Columbia University  
New York, New York, USA  
shane@ee.columbia.edu

Andrew T. Campbell

Computer Science, Dartmouth College  
Hanover, New Hampshire, USA  
campbell@cs.dartmouth.edu

**Abstract**—A transmitter in a wireless network that uses CSMA, a simple carrier sensing-based MAC protocol, to determine the likelihood of successful packet reception at the intended receiver can easily be misled. At the same time, CSMA variants and hybrid MAC protocols based at least in part on carrier sensing have become the *de facto* standard in wireless sensor networks, underscoring a need to improve its performance. We propose to enhance the *de facto* state of carrier sensing-based MACs in wireless sensor networks by using low cost channel feedback combined with a learning approach to try to better predict the probability of a successful reception, *on a per-receiver basis*. We show results from an experimental wireless sensor network testbed, where our proposal E-CSMA (Enhanced CSMA) provides up to a 55% improvement in network performance.

## I. INTRODUCTION

In wireless networks, the successful use of the carrier sensing approach for determining the likelihood of successful reception requires that a transmitter can accurately determine the state of the channel at the receiver via strictly local sampling. Simple CSMA takes the narrow view that a clear channel at the transmitter implies a clear channel at the receiver, and assumes that the channel will remain clear at the receiver for the duration of the transmission. Clear channel assessment (CCA) via carrier sensing is done by analyzing various signal characteristics (e.g., received signal strength) in the carrier band. In a network where all channel contenders are within radio range of each other, and no external sources of radio interference exist, CCA via carrier sensing works well to determine the likelihood of successful reception at an intended receiver. However, in the general case, hidden and exposed terminals may exist [1], and the time-varying nature of the physical channel due to fading, independent of node interactions, further complicates matters. To illustrate part of the problem, Figure 1 shows the impact of the hidden transmitter problem in a simple topology of Mica2 motes running B-MAC [7]. In Figure 1(a), node T transmits to  $R_1$  and  $R_2$  while node I interferes. Nodes I and T each transmit with a power of  $-10\text{dBm}$  at  $\frac{2}{3}$  of the maximum achievable packet rate. Figure 1(b) shows the packet reception success ratio for  $R_1$  and  $R_2$  with 95% confidence intervals over 5 trials in parentheses. B-MAC is not designed to handle the hidden terminal problem [7], and Figure 1 shows that the hidden terminal problem has a significant impact; all packets to  $R_1$  are successfully received, but nearly 75% of T’s packets to  $R_2$  are corrupted by packets from I, a huge energy waste.

\*This work is supported by the Army Research Office (ARO) under Award W911NF-04-1-0311 on resilient sensor networks.




Fig. 1. (a.) A simple topology to show the hidden terminal problem; (b.) the 75% packet loss from T to  $R_2$  is a huge energy penalty to the network.

Despite the drawbacks of carrier sensing, because CSMA is fully distributed and has low implementation complexity, variants of CSMA and hybrid MAC protocols based in part on carrier sensing are widely deployed in established (e.g., IEEE 802.11 networks [12]) and emerging (e.g., sensor networks [13]) wireless networks. Given the practical concerns that favor a fully distributed, low complexity MAC solution in these networks and the substantial momentum in this direction, there is evident utility in improving the performance of CSMA-style medium access control. In this paper we propose to address the limitations of carrier sensing, without requiring explicit per-packet signaling (e.g., RTS/CTS), based on two motivating observations. First, though a “good” channel as measured at the transmitter does not necessarily indicate a “good” channel at the receiver, there may still be a correlation (positive or negative) between the transmitter and receiver channel states that can be exploited. Second, for each transmitter, the state of the wireless channel at each potential receiver in its radio range is unique. Therefore, we propose to measure the existing correlation empirically using in-band feedback mechanisms to create a probability distribution of successful reception. We further propose that a transmitter maintains such a distribution for each of its potential receivers (i.e., one hop neighbors). In Figure 1, simply by recognizing that a given received signal strength (or signal strength range) corresponds to a high success rate at  $R_1$ , but a low success rate at  $R_2$ , T can avoid transmitting to  $R_2$  under conditions leading to historically low success, while maintaining a high throughput to  $R_1$ .

We discuss tradeoffs concerning the construction and use of the proposed per-receiver probability distributions, and present E-CSMA (Enhanced CSMA), our solution to address the limitations of carrier sensing and enhance the performance of CSMA-style MACs in Section II. To verify the performance improvement of E-CSMA over baseline CSMA, we implement our algorithms in an experimental wireless sensor network, and report results in Section III. Related work is summarized in Section IV before we conclude.

## II. E-CSMA DESIGN

Naive interpretation of carrier sense values, even when averaged over a window of time, can lead to unnecessary collisions and a waste of idle channel bandwidth [9]. Leveraging wireless channel feedback, implicit or explicit, from the receiver to the transmitter, E-CSMA builds and maintains an empirically generated probability distribution of packet reception success *for each receiver*, with respect to observable channel conditions at the transmitter. When a transmitter is faced with a transmit/defer decision, these probability distributions are then referenced, indexed by the currently observable channel conditions. In what follows we describe the design details of E-CSMA, including the construction and maintenance of the reception probability distributions, and how these distributions are utilized when faced with a transmit/defer decision.

### A. Channel State Management

The E-CSMA probability distributions attempt to correlate observable channel conditions at the transmitter with success probability at the receiver. For each transmitted packet the transmitter collects information concerning the local channel conditions under which it is transmitted, and an indication of whether the transmission was successfully received by the intended receiver. In E-CSMA we represent the channel conditions with a finite discrete set  $\mathcal{C}$  having members  $c_i$ , and  $C^t$  represents the observed channel condition sample at the transmitter, and  $S \in \{0, 1\}$  is the reception indication. We aim to empirically track  $P(S = 1 | C^t = c_i)$ , for every  $c_i \in \mathcal{C}$ . This discrete notion of channel conditions maps well to the nature of carrier sensing technology, where at the lowest level the channel state is determined by an in-band sample of the received signal strength (RSS), or a function of the RSS (e.g., difference between the instantaneous RSS value and an exponentially-weighted moving average, hereafter EWMA, noise floor). While the signal strength is analog and thus continuous, in any physical platform design it is available as a quantized digital value and thus is discrete. However, in this case the cardinality of  $\mathcal{C}$  is exponential in the resolution of the RSS ADC, and empirically creating a probability distribution for each possible RSS value may not be practical due to memory constraints or even useful since the learning time across all values of RSS may take too long. Thus, in E-CSMA a transmitter segments the RSS range into bins and collects and maintains average statistics at a bin level granularity only. More specifically, we seek to determine the probability  $P(S = 1 | b_k^l < C^t \leq b_k^u)$ , where  $b_k^l$  and  $b_k^u$  are the lower and upper bin bounds, respectively, for bin  $k$ . There is an inherent tradeoff between bin width and the empiric learning time of the distribution. Even with the reduced storage requirement implied by keeping only bin level statistics, platform memory limitations restrict the number of neighbors about which a node can keep information. As a node's neighborhood changes (e.g., due to mobility or node replacement), a neighbor replacement strategy (e.g., LRU) may be necessary to maintain information about relevant neighbors. (Investigation of neighbor replacement strategy is outside the

scope of this paper.) In practice, the number of radio neighbors with which a given node regularly exchanges data is likely to be relatively low but in general depends on density. For example, even though there may be many potential next hops for a given node, due to link quality filtering likely done by the routing algorithm the number of neighbors about which a node might actively need to maintain state is much lower.

### B. Channel Feedback Mechanism

There are a number of forms channel state feedback may take. With E-CSMA we wish to enhance the link layer behaviour already in use in the network stack. As such we do not force the use of link layer acknowledgments (e.g., Ack, Nak), but will take advantage of these feedback conduits if they are available to give an indication of packet reception success. Otherwise, E-CSMA uses implicit (passive) feedback based on overhearing, with an option to the application to improve performance by adding periodic feedback. In Section III, we implement and test positive and periodic feedback; in the following we provide design details of positive and periodic feedback; a discussion of negative and implicit feedback is omitted due to space constraints.

1) *Positive Feedback*: With positive feedback, a receiver sends back an acknowledgment for each packet received. Positive feedback has a higher overhead cost than the alternatives, but the information fed back is fresh since the success indication is available for each packet shortly after the transmitter finishes its transmission. This is important because the transmitter wants to make the decision about when to transmit the next packet based on knowledge of all the packets it sent before, including the most recent. In E-CSMA, the transmitter keeps a FIFO queue, per receiver, per bin, of the  $M$  most recent feedback results, where a result is 1 if an acknowledgment is received for a given packet and 0 otherwise. The instantaneous estimate of the success probability to a given receiver with channel conditions mapping to a given bin is calculated as a function (e.g., average) of these  $M$  most recent feedback results, where  $M$  controls the granularity of the stored probability.

2) *Periodic Feedback*: With periodic feedback, receivers return the number of packets successfully received (per neighbor, per bin) in the previous interval, and may optionally return the number of packets with failed CRCs over the same interval. Transmitters count the number of packets they send (per neighbor, per bin) over the same interval to compare with the information fed back from the receivers. The feedback interval is defined in terms of number of packets, where the receiver sends back the aggregated acknowledgment after successfully receiving a specified number of packets. E-CSMA uses a two window approach to regulate the probability distributions. The receiver sends feedback (per bin) every interval  $W_R$ , while the transmitter expects feedback (per bin) from the receiver every interval  $W_T$ . If no feedback is received within  $W_T$  then the bin's success probability is multiplicatively decreased (aged) with  $d < 1$ . When feedback is received, the success probability is set to the success ratio calculated from the receiver

acknowledgment and the local counters of the transmitter.  $W_R$  is proportional to the interval that channel access might be too conservative since it controls the feedback frequency, the only way to increase the success probability.  $W_R$  also governs the protocol overhead of the periodic scheme.  $W_T$  is proportional to the interval that channel access might be too aggressive since it controls the probability decay frequency. It is thus also proportional to the number of packets successfully received and number of packets lost.  $W_T$  should be set according to the radio environment stability. To balance throughput and loss,  $d$  should be set proportional to  $W_R/W_T$ .

### C. Transmit/Defer Decision Rule

To ultimately make a decision to transmit or backoff, we must apply a rule to the instantaneous input channel conditions value and the empirically constructed probability distribution. As discussed in Section II-A, we design E-CSMA to discretize the distribution into probability bins, and transmit only if in the past we succeeded at least a fraction  $\alpha$  of the time when the input sample fell in the same bin as the current input sample. This approach does not assume that the success probability is monotonic across input sample values, and provides finer granularity, depending on bin width. To balance the probability of channel loss against the probability of packet drop loss at the transmitter queue, queue occupancy is considered along with  $\alpha$  when a transmit/defer decision must be made. As  $\alpha$  increases the probability of channel loss decreases (since the transmitter is more picky about when to transmit), but delay and the probability of packet drops at the transmitter queue increase; the opposite is true for a small  $\alpha$ . Therefore we define  $\alpha' = \alpha/f(\gamma_t)$ , where  $\gamma_t$  is the queue occupancy percentage at time  $t$ , and use  $\alpha'$  for the decision threshold.

In the discussion to this point we have implicitly assumed a unicast link abstraction. However, in wireless networks a multicast link abstraction is often employed to take advantage of an underlying physical layer that is inherently broadcast. For a multicast link abstraction some method of channel state combining is required. Waiting for a local channel condition that correlates to successful reception at all the intended receivers with high probability maximizes the performance in terms of collision, but the delay is also maximized. In E-CSMA a transmitter waits to transmit until  $n\%$  of the intended receivers have a high probability ( $> p$ ) of successful reception, where  $n$  and  $p$  are set by the application to balance the delay and collision performance, and can change on a per packet basis to adapt to local conditions and application requirements. In the results reported in this paper E-CSMA transmitters only apply their learned probability distributions to transmit decisions involving unicast data packets, leaving the underlying CSMA MAC to handle multicast packets.

When a transmitter decides that it must defer, the simplest approach is to choose a backoff length at random from a fixed backoff window. Alternatives include deferring a time inversely proportional to the success probability for the intended receiver(s), or deferring a time set in accordance to both the local channel conditions at the transmitter and the

local queue occupancy of the transmitter. Backoff is handled by the default MAC for the results reported in this paper.

### D. E-CSMA Medium Access Strategy

Channel access decisions based on empirically constructed probability distributions of reception success provide the greatest performance enhancement when transmitters get timely feedback from receivers, and topology changes slowly. E-CSMA defines a comprehensive medium access strategy designed to shift between the transmit/defer decision given by the underlying CSMA implementation and that given with reference to the transmitter's probability distributions to obtain the best performance. In the following, we summarize the transmit criteria of the decision rule outlined in Section II-C (denoted Condition 1) and of B-MAC [7], a CSMA variant designed for wireless sensor networks (denoted Condition 2).

#### • Condition 1:

Using our binning strategy, the transmit criterion is

$$P(S = 1 | b_k^l < C^t \leq b_k^u) > \alpha'$$

$$\left[ (S_M + \sum_{j=1}^M S_j) / (M + 1) \right]_k = \quad (\text{for positive})$$

$$\left[ R_{last.feedback} \cdot d^i \right]_k = \quad (\text{for periodic})$$

where  $M$  is the FIFO queue size for positive feedback results,  $R_{last.feedback}$  is the ratio calculated and stored in bin  $k$  as a result of the last periodic feedback from the receiver,  $i$  is the number of intervals  $W_T$  that have passed since the last periodic feedback has been received, and other symbols are as previously defined.

#### • Condition 2:

B-MAC uses an outlier detection scheme that compares the minimum of five consecutive channel samples measured at the transmitter to an adaptive noise floor. The transmit criterion is

$$C^t < \text{Threshold}_{B-MAC} = f_{EWMA}(\text{Noise Floor}) - \text{Bias},$$

where  $C^t$  is the minimum of the five samples and  $\text{Bias}$  dictates by how much the sample must differ from the noise floor (the *Outlier Magnitude*), before transmitting.

The E-CSMA channel access strategy is such that for a fraction  $T$  of the packets the transmit/defer decision is made according to Condition 1 and is otherwise  $(1 - T)$  made according to Condition 2. As such,  $T$  controls the duty cycle of Condition 1, and hereafter we use the term *E-CSMA duty cycle* as a synonym for  $T$ .  $T$  evolves in accordance with real-time performance (i.e., successful packet reception) in the network. To track this performance we keep an EWMA of both the E-CSMA performance and the default CSMA performance.  $T$  is increased when E-CSMA performance is relatively better, and decreased in the opposite case. Specifically, the value of  $T$  evolves according to

$$T_{i+1} = T_i + \begin{cases} \tau, & \frac{f_{EWMA}(\text{E-CSMA Succ. Ratio})}{f_{EWMA}(\text{CSMA Succ. Ratio})} \geq 1 + \delta; \\ -\tau, & \frac{f_{EWMA}(\text{E-CSMA Succ. Ratio})}{f_{EWMA}(\text{CSMA Succ. Ratio})} \leq 1 - \delta; \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

Here,  $\delta$  controls the sensitivity of the  $T$  adaptation, and  $\tau$  controls the granularity of the  $T$  adaptation. The adaptation can be packet driven (e.g., updated after every packet transmission) or time driven, where in the latter case  $\Delta t = (t_{i+1} - t_i)$  controls the time scale of the adaptation. Additionally,  $T_{min} \leq T_i \leq T_{max}$ , to allow both CSMA and E-CSMA the chance to recover packet share as the radio environment changes, where  $0 \leq T_{min}, T_{max} \leq 1$  necessarily.

For each transmission, a transmitter generates a random number  $r$  between 0 and 1 inclusive, and compares the value against the current value of  $T$ . If  $r > T$  then Condition 2 must hold in order to transmit, while if  $r \leq T$  Condition 1 must hold. Note that even when packet transmission decisions are made according to Condition 2, RSS/reception success statistics are gathered at the transmitter to keep the E-CSMA probability bins as up to date as possible. Further, the E-CSMA duty cycle  $T$  is updated according to Equation 1 on a per-receiver basis. Note that E-CSMA feedback messages, broadcast messages, and unicast message for which no receiver record currently exists at the transmitter (due to buffer limitation or neighbor replacement strategy) are transmitted according to Condition 2.

### III. EXPERIMENTAL SENSOR NETWORK TESTBED

In this section, we discuss the implementation of E-CSMA on a real sensor network using TinyOS [6] on Mica2 motes [10]. We report testbed results, including the performance impact of using Outlier Magnitude or Raw RSS as the decision input value set, a study of the effect of the decision threshold  $\alpha'$ , and the performance impact of traffic load. For this initial study, we implement and study three feedback mechanisms, ideal, positive and periodic, to demonstrate the potential and feasibility of E-CSMA. Throughout this section we show how use of E-CSMA algorithms can best offer a performance improvement to CSMA, using the default implementation of the CSMA protocol B-MAC [7] as a representative example.

#### A. Mote Testbed Configuration

Our testbed comprises 31 Mica2 motes arranged in a roughly rectangular grid. The motes are connected to MIB600 [10] Ethernet interface boards which are mounted across a suspended ceiling hanging about 4ft from the concrete structural ceiling (six nodes are instead mounted on the dry wall) of a roughly 1600ft<sup>2</sup> room of our research center. The motes hang approximately 10ft above a metal tile raised floor. A line of heavy-duty rigging for electrical conduits and throw switches hangs approximately 2ft from the suspended ceiling between nodes 28-31 and the rest, offering substantial reflective surface area. The Mica2s are powered from the MIB600s, which in turn are powered by a IEEE 802.3af Power over Ethernet hub via standard CAT5E cables. The Ethernet back channel is used for mote programming, experimental parameter configuration and data collection, and as the feedback channel for the ideal feedback scheme. Motes use a frequency of 914.077MHz and are equipped with standard  $\frac{\lambda}{4}$  whip antennas.




Fig. 2. The average cardinality of the neighbor and the null link sets at each tested transmission power. The inner (solid) error bars show the standard deviation and the outer (dashed) error bars show the max and min. The minimum tested power at which all nodes have  $\geq 1$  neighbor is -10dBm.

To determine an appropriate transmit power setting for subsequent experimentation, we measure the cardinality of the neighbor set for each transmitter at a number of radio power levels (*viz.*, -18, -15, -12, -10, -9, -7, -5, -2)dBm. For each power setting each transmitter separately sends fifty 36-byte packets to each other mote in turn. A simple packet delivery ratio is calculated at each mote for each power level, for a total of  $31 \times 30 \times 8 = 7440$  data points. Figure 2 summarizes the results with a plot of average node neighborhood size and null link set size (number of nodes to which the packet delivery ratio is zero) versus power level. In this figure, the ‘Avg. Num. of Null Links’ curve has been shifted to the right by 0.2dBm to allow for better readability. For our experiments we would like the ability to create a multi-neighborhood environment while having a connected topology. For the purposes of this plot, we consider that node  $x$  is a neighbor of node  $y$  if the packet delivery ratio of transmissions from  $x$  to  $y$  is at least 80% (link symmetry is not assumed). From the plot we see the lowest power for which each node has at least one neighbor is -10dBm. Further, at this power, all nodes have at least 7 nodes to which they have a packet delivery ratio of zero (average = 17, max = 26). Thus we consider that a transmit power of -10dBm yields a connected multi-neighborhood environment and we use this power at each transmitter.

#### B. Protocol Parameterization and Experimental Overview

To set the values for bin range  $BR$  and number of bins  $NB$ , we run a simple experiment using two MIB600-mounted Mica2 motes ( $A$  and  $B$ ) on a wooden table top. Motes are equipped with quarter wave antennas. The MAC of mote  $A$  is modified such that it constantly logs and reports RSS readings from the radio (ADC Channel 0 on the Mica2 platform). With mote  $B$  not transmitting, the maximum RSS sample value (minimum received power) observed over a ten minute period is 407. With mote  $B$  transmitting at the maximum power of 5dBm [11], the distance between  $A$  and  $B$  is reduced until the RSS reading is minimum (maximum received power). The minimum RSS sample value observed over a ten minute period is 23. We therefore set  $BR$  for the ‘Raw RSS’ testing to 512 and implement uniform width bins between 0 and

512. Empiric observations across the experimental conditions (i.e., transmitter rate and power) show that only about half of the bins are populated when  $NB = 12$ . However, for our initial investigation we keep  $NB = 12$  to allow for extreme outliers under conditions not yet observed. Following similar observations, we set  $BR$  for the ‘Outlier Magnitude’ testing to 128 and  $NB = 12$ . ‘Raw RSS’ and ‘Outlier Magnitude’ testing is presented in Section III-D

To set the values for the periodic feedback parameters ( $W_R, W_T, d$ ), experimentation (omitted, space constraints) indicates that values of  $W_R, W_T \leq 5$  and  $W_R/W_T$  closer to 1 provide the best network performance. Following this guidance, for the results reported in this paper we use  $W_R = W_T = 4$  and  $d = 0.5$ .

All B-MAC [7] parameters (congestion backoff window size, initial backoff window size, *Bias*) are set to the default values in the TinyOS-1.1.13 [6] implementation.

In the remaining experiments we profile E-CSMA performance for a range of values for the success probability threshold  $\alpha$  and the E-CSMA duty cycle  $T$ . (We use  $f(\gamma) = 1$ ;  $\therefore \alpha' \rightarrow \alpha$ , so we drop the prime notation.) To gain initial insight into system performance under known parameterization, rather than having  $T$  self-regulating according to Equation 1, we initialize  $T$  to specific values, and set  $\delta$  and  $\tau$  to values that effectively disable the performance-based adaptation. We report results on the performance of E-CSMA using positive, periodic and ideal feedback. Ideal feedback is positive feedback sent out-of-band over an Ethernet back channel (for experimental comparison only). While impractical, the results reported for ideal feedback provide a good upper bound on the achievable accuracy of the probability distribution being measured and the resultant network performance improvement possible when there is no in-band signaling. Ultimately, the performance of each is compared against a plain CSMA benchmark, B-MAC.

Unless otherwise noted, each data point in the following plots represents an average of five trials where (regardless of source rate) 200 packets are transmitted by each of the 31 transmitters. In the following, error bars indicate the 95% confidence intervals over the five trials. The reported source rates are average rates for each transmitter. Actual packet transmissions at rate  $r_{src} = 1/t_{src}$  are randomized within the interval  $t_{src}$ . Thus, at every time  $t = n \cdot t_{src}$  the start of the next transmission is set as  $t + rand(0, t_{src} - PKT\_TX\_TIME)$ , where  $rand(min, max)$  returns a number chosen uniformly at random in the supplied interval, and  $PKT\_TX\_TIME$  is the time necessary to transmit a packet at the specified rate  $r_{src}$ .

### C. E-CSMA Performance Metrics

*Link Tax.* The Link Tax is a normalized representation of the energy wasted by dropping packets across a link. Link Tax = (Packets dropped across a link)/(Packets received across a link). Measurements are taken across a time window. Since packet transmission and reception consumes the main portion of energy on a mote minimizing the Link Tax is critical to increasing network lifetime.

*Link Fidelity.* Link Fidelity is simply the number of packets received across a link, i.e., the goodput. This metric provides




Fig. 3. Comparison of ‘Raw RSS’ and ‘Outlier Magnitude’ in terms of average (a) Link Tax and (b) Link Fidelity, as candidates for the E-CSMA decision input value. Outlier Magnitude outperforms Raw RSS by leading to fewer dropped packets per delivered packet network wide.

an indication of how successful the MAC protocol is at propagating offered load within a given window of time.

*Link Power.* Ideally, the MAC should offer high Link Fidelity and low Link Tax to the upper layers. Analogous to the traditional notion of power as throughput/delay, we define Link Power = (Link Fidelity)/(Link Tax) and aim to maximize the Link Power, increasing network lifetime by reducing energy-wasting collisions while offering high goodput.

In the experimental results presented in this section we report the *network average* of Link Tax, Link Fidelity and Link Power, that is, the average Link metric of all neighbor links involved in the particular experiment.

### D. Decision Input Value Selection

In this section we compare the E-CSMA performance achieved when using the Raw RSS value as the input value to the decision process (i.e.,  $C^t$  in Condition 1 in Sec. II-D), versus using the *Outlier Magnitude* (cf. Condition 2 in Sec. II-D), that is, the difference between the Raw RSS value and an adaptive estimate of the noise floor value observed by the transmitter. The current default MAC implementation [7] in TinyOS [6] provides an EWMA of the noise floor, which E-CSMA employs to determine the Outlier Magnitude. If the outlier sign is negative (i.e., if the received noise power is higher than the average noise floor) we always defer. Using this rule, we could be losing throughput by avoiding possibly successful transmission opportunities. However, in sensor networks we believe energy conservation is more important than optimizing for throughput, and thus we are willing to sacrifice some throughput to reduce the collision probability.

To eliminate any biasing effect of a particular in-band feedback mechanism, we compare these two input sets using the ideal feedback scheme. The value of  $T$  is fixed at 1 for this experiment, meaning all packets will be sent using the E-CSMA decision (i.e., by Condition 1). Each transmitter unicasts a packet to a random neighbor (from the neighbor sets identified from the experiment detailed in Section III-A) at a mean rate of 2 packets per second. Qualitatively similar results to those reported here are observed at all other tested rates (viz., 0.25, 0.5, 1, 2, 3, 4, 5 packets/sec) but are suppressed due to space constraints.

In Figures 3(a) and 3(b) we plot the average Link Tax and average Link Fidelity in the network, respectively, versus a

range of  $\alpha$  values. It is clear that using Outlier Magnitude offers superior performance in terms of both number of packets lost (leading to a lower Link Tax) and number of packets successfully received (leading to a higher Link Fidelity). As indicated by Figure 2 of [7], the Raw RSS tends to be noisy and can lead to a false indication of channel state, whereas the Outlier Magnitude is based on an EWMA of recent RSS samples that provides a more stable value stream from which to empirically learn the probability distributions for successful reception. Even though BMAC uses this same outlier detection scheme, the additional information gathered and maintained by E-CSMA can offer performance gains, as seen in Sections III-F and III-G. Figure 3(a) also shows that the delta between the Raw RSS curve and the Outlier Mag curve shrinks as  $\alpha$  increases. This is expected since at low  $\alpha$  the instability of the Raw RSS is more likely to represent a bad channel RSS as belonging to a bin where the empiric success probability is high enough to transmit. As the transmitter becomes more picky about when to transmit (as  $\alpha$  goes up) this is less likely to occur. In both Figure 3(a) and 3(b), the influence of  $\alpha$  is clear: with increasing  $\alpha$  the transmitter is more selective about when to transmit, driving both packet collisions and packet throughput down. The impact of  $\alpha$  is the focus of the next section. Due to its superior performance compared to Raw RSSI, we use Outlier Magnitude as the input value set for the remainder of the experiments.

#### E. Impact of Success Probability Threshold

In this section we evaluate the performance impact of the E-CSMA success probability threshold  $\alpha$ . The value of  $T$  is fixed at 1, so all packets will be sent according to the E-CSMA Condition 1, and nodes transmit at an average rate of 2 packets per second. Qualitatively similar results are observed at all other tested source rates but are suppressed due to space constraints. Figures 4(a) and 4(b) show the Link Tax and Link Fidelity, respectively, for the ideal, positive and periodic feedback schemes. Generally, we see that as  $\alpha$  increases the transmitter becomes more selective about when to transmit and the Link Tax decreases, indicating an energy savings due to fewer packet collisions. However, Link Fidelity also drops off with increasing  $\alpha$ . To disambiguate the situation and point to a value of  $\alpha$  that provides the best blend of energy conservation and throughput, we plot the Link Power, i.e., the ratio of Link Fidelity to Link Tax, in Figure 4(c).

With respect to Link Power, it becomes clear that performance increases monotonically across the tested  $\alpha$  values for the ideal and positive feedback schemes, but this only holds marginally for the periodic feedback scheme, especially between  $\alpha$  of 0.7 and 0.9 where Link Tax, Link Fidelity and Link Power have the same respective values. We note that ideal and positive schemes send an acknowledgment packet for every packet received, providing a relatively smooth feedback signal to the transmitter and incremental updates of the bin probabilities, while the periodic feedback scheme sends only an aggregated acknowledgment for every  $W_R = 4$  packets in our implementation. Therefore, even under relatively good

long term channel conditions, occasional acknowledgment loss has a bigger impact. This is especially true when  $W_R = W_T$ , as it is in our implementation, since the loss of an acknowledgment or even one data packet will cause the bin probability to be multiplicatively reduced by  $d = 0.5$ . Due to the aggressive probability reduction by  $d$ , the rate at which E-CSMA decides to transmit rather than defer for a given receiver/bin is already minimal when  $\alpha$  reaches 0.7. This is reflected in the leveling of the Link Fidelity, Link Tax and Link Power curves for periodic feedback at  $\alpha = 0.7$ . We conjecture that the  $\alpha$  value after which this leveling occurs can be adjusted by changing  $d$ . Figures 4(a), 4(b) and 4(c) provide a guideline to applications that can customize their performance depending on relative importance placed in saving energy by minimizing dropped packets or pushing more packets across the link. For the remaining experiments we use  $\alpha = 0.9$ , since the Link Power is maximal at this value for all three tested feedback mechanisms.

#### F. Impact of E-CSMA Duty Cycle

Here we investigate performance across a range of values for the E-CSMA duty cycle  $T$ , that is, the percentage of time Condition 1 is used to make the transmit decision, versus B-MAC (Condition 2). When  $T = 0$ , all traffic is sent using plain B-MAC (Condition 2); when  $T = 1$ , all traffic is sent using E-CSMA Condition 1. We test several intermediate values to gain insight into the expected performance in a stable radio environment (no node join/death/mobility). While node join/death are likely to occur in a deployed system, these events do not typically take place on MAC time scales, but rather over days, weeks or months. Thus, we expect many static networks to spend most of their time in a relatively stable radio environment. For more dynamic radio environments, if E-CSMA performs poorly  $T$  will automatically be steered towards zero, according to Equation 1 such that in the worst case E-CSMA performance will hover near plain CSMA. In these reported results  $\alpha$  is set to 0.9 and each node transmits at an average rate of 2 packets per second. Qualitatively similar results are observed at all other tested values of  $\alpha$  (*viz.*, 0.1, 0.3, 0.5, 0.7, 0.9) and mote source rate (*viz.*, 0.25, 0.5, 1, 2, 3, 4, 5 packets/sec) but are suppressed due to space constraints.

For Figures 5(a), 5(b) and 5(c), on the ideal feedback curve,  $T = 0$  represents B-MAC without link acknowledgments; on the positive feedback curve,  $T = 0$  represents B-MAC with link acknowledgments; on the periodic feedback curve,  $T = 0$  represents B-MAC with periodic link acknowledgment. While the B-MAC transmit decision is not influenced by the information contained in the feedback, we plot these points to show performance in a network using plain B-MAC when part of the channel bandwidth is consumed by link layer acknowledgments (e.g., when enabled for link layer reliability).

Figure 5(a) shows that for all feedback types, increasing the percentage of transmit decisions made according to Condition 1 (increasing values of  $T$ ) drives the Link Tax down unconditionally, across the tested values. At values of  $T \geq 0.2$ , Link Tax results for all types of feedback are better than B-MAC




Fig. 4. Impact of  $\alpha$  on (a) Link Tax, (b) Link Fidelity, and (c) Link Power for the ideal, positive and periodic feedback schemes.




Fig. 5. Impact of E-CSMA duty cycle  $T$  on (a) Link Tax, (b) Link Fidelity, and (c) Link Power for the ideal, positive and periodic feedback schemes. When  $T = 0$ , all packets are sent using plain B-MAC.

with no acknowledgments (see the dashed line Link Tax = 1.073), despite the bandwidth overhead and increased collision probability the positive and periodic schemes incur because of their in-band feedback.

Figure 5(b) shows that as  $T$  values rise from 0 (plain B-MAC) to 1, the Link Fidelity for all three feedback schemes falls off. While B-MAC can get more packets through on average because of its less discriminating channel access rule, it is at the expense of a marked increase in dropped packets (see Link Tax results from Figure 5(a)). The ideal scheme has the highest Link Fidelity as it uses no in-band feedback packets; periodic feedback is second best as it uses in-band feedback but in an aggregated way to reduce the number of acknowledgment packets. Also the Link Fidelity for the periodic feedback scheme in Figure 5(b) is only 20% away from B-MAC in the worst case; the performance penalty in terms of throughput is small relative to the concomitant 70% reduction in Link Tax. We are further encouraged by the observation that periodic feedback can achieve performance nearly equivalent to that of positive feedback with only  $\frac{1}{4}$  of the packet overhead. This has important implications for energy conservation and network lifetime in a topology of battery-powered motes.

Figure 5(c) shows that Link Power increases with the tested values of  $T$ , suggesting that we can always benefit from increasing the E-CSMA duty cycle, in our testbed. In general (e.g., in a highly dynamic radio environment) it may not always be the case. Thus for actual deployments  $T$  evolves according to Equation 1. Increasing  $T$  clearly provides an advantage when comparing positive-E-CSMA to positive-BMAC ( $T = 0$  on the positive curve), and at all tested values of  $T > 20$ , positive-E-CSMA and periodic-E-CSMA even

outperform BMAC without acknowledgments (see the dashed line Link Power = 87.11), despite the increased collision probability and bandwidth consumption implied by the positive and periodic wireless feedback. At  $T = 1$ , both periodic and positive feedback schemes yield a 35% increase in Link Power over plain B-MAC (difference between periodic and positive at  $T = 1$  and B-MAC's dashed line at 87.11). From Figure 5(c) we see that Link Power does not increase much for periodic and positive feedback schemes as  $T$  increases from 0.8 to 1.0. These results suggest  $T_{max}$  from Equation 1 should be set near 0.8; for fairness  $T_{min}$  should be set near 0.2.

### G. Impact of Traffic Load

In this section, we provide insight into how E-CSMA might perform under a range of realistic offered load conditions. As mentioned in the previous section, many networks are likely to enjoy a fairly stable radio environment, and under these conditions Figure 5(c) shows that Equation 1 would push  $T$  to the  $T_{max}$  bound given the increasing performance advantage over plain CSMA with increasing  $T$ . Based on the reasoning from Section III-F, in this section we take  $T_{max}$  to be 0.8 and evaluate E-CSMA performance at a number of per node source rates, assuming  $T$  is pushed to the bound  $T_{max}$ . The success probability threshold  $\alpha$  remains at 0.9.

Figure 6 shows a plot of Link Power versus the tested source rates, *viz.*, 0.25, 0.5, 1, 2, 3, 4, 5 packets/sec (plots of Link Tax and Link Fidelity versus offered load are omitted due to space constraints). E-CSMA performance results appear for each of the three tested feedback types, and B-MAC results ( $T = 0$ ) for each feedback type are also present for comparison. From the figure, the following trends emerge. Firstly, every curve decreases sharply with increasing rate,




Fig. 6. Link Power versus offered load (per node) for the ideal, positive and periodic feedback schemes. Source rate on the x-axis is for each of the 31 concurrent sources. A zoom of the tail of the bottom five curves is shown as an inset at the same x-axis scale (from 3 to 5 packets/sec).

as channel contention and thus collision probability increases. Even the the margin of improvement given by ideal E-CSMA decreases as the offered load exceeds the channel capacity in a greater percentage of the neighborhoods in the testbed. Secondly, though Periodic B-MAC ( $T = 0$ ) and Positive B-MAC ( $T = 0$ ) converge in the tail (see inset of Figure 6), Positive E-CSMA ( $T = 0.8$ ) maintains a small Link Power advantage over Periodic E-CSMA ( $T = 0.8$ ) across all tested rates. At higher rates fewer packets are successfully received and therefore fewer acknowledgments are sent. Consequently, the acknowledgment packet overhead imposed on B-MAC becomes almost the same for both schemes, nearly eliminating the advantage of the Periodic scheme seen at lower rates. However, Positive E-CSMA continues to throttle transmission attempts more aggressively than Periodic E-CSMA, leading to a lower Link Tax (not shown) and higher Link Power. Most importantly, we see that for positive and periodic feedback, E-CSMA ( $T = 0.8$ ) performance is 15-55% (depending on offered load) better than the corresponding B-MAC ( $T = 0$ ) performance, and in fact all three E-CSMA curves are above all B-MAC curves at all tested rates.

#### H. Importance of Receiver Differentiation

Very recently published work [14] proposes an approach similar to that of E-CSMA. The authors of [14] share the same motivation of mitigating the hidden and exposed terminal problems without explicit signaling, instead letting transmitters build a correlation between observable channel state and reception success probability. However, in [14], probability distributions are not kept per neighbor but per neighborhood; in essence, an average probability distribution is kept by a transmitter across all receivers in its neighborhood. In [14], the authors integrate their scheme with IEEE 802.11 DCF mode and simulate a 25 AP, 2 clients per AP scenario, where the authors state 60% of the total topology suffers from a hidden or exposed terminal. The authors state that their scheme greatly reduces packets drops and increases throughput *for those hidden/exposed terminals*, but average results across all nodes are not included.

In order to evaluate the potential of such a scheme for




Fig. 7. The Link Power Ratio, i.e., the Link Power of baseline E-CSMA over the Link Power of E-CSMA<sub>neighborhood</sub>, shows the advantage of using per-neighbor probability distributions over per-neighborhood. (a) shows the advantage across values of  $T$ , where at  $T = 0$  (B-MAC) the ratio is 1 since the transmit/defer decision does not involve the distributions. (b) shows the advantage across source rate, where at all tested rates baseline E-CSMA performs better (Ratio > 1) for all feedback schemes.

real wireless sensor networks, we implement a variant of E-CSMA, termed E-CSMA<sub>neighborhood</sub> for this discussion, that keeps a single probability distribution at each transmitter rather than one for each receiver. We present results here from our experimental testbed under the same testbed configuration and experimental parameterization as our earlier results. Due to space constraints, we show only comparative results between our baseline E-CSMA proposal and E-CSMA<sub>neighborhood</sub>. Figure 7(a) shows the Link Power Ratio, that is, the Link Power of baseline E-CSMA over the Link Power of E-CSMA<sub>neighborhood</sub>, across a range of  $T$  values; Figure 7(b) shows the Link Power Ratio across a range of transmitter source rates. In each of these figures we see that baseline E-CSMA performs substantially better than E-CSMA<sub>neighborhood</sub>, for all three feedback mechanisms. For example, the periodic feedback scheme performs at least 2x and up to 7x better in terms of Link Power, when  $T = 0.8$  (Figure 7(b)).

When reception statistics are combined across all the neighbors of a transmitter, neighbors with high success in a given bin are averaged out by neighbors with low success. Consequently, transmissions to “good” neighbors are deferred unnecessarily, while transmissions to “bad” neighbors proceed when they should not. The latter event causes more dropped packets in the short term and a commensurate lowering of the empirically measured success probability for that bin. This further throttles transmissions to the “good” neighbors, prohibiting these from bringing the neighborhood success ratio back up. In the end, the distribution settles out at the probabilities for the worst neighbor in each bin. Thus, E-CSMA<sub>neighborhood</sub> is very conservative and has excellent Link Tax performance, but the Link Fidelity is *very* low, leading the Link Power lower. The Link Tax and Link Fidelity components (omitted) of Figures 7(a) and 7(b) bear this out. For example, at 2 packets/sec the periodic feedback scheme has up to a 25% lower Link Tax with E-CSMA<sub>neighborhood</sub>, but also suffers up to a 70% lower Link Fidelity. These results unambiguously show the importance of maintaining per-receiver probability distributions, especially in radio environments (dense nodes, reflections) where all neighbors of a given transmitter are not likely to have the same channel conditions at all times.



#### IV. RELATED WORK

Comments and preliminary results on the limitations of carrier sensing for CC1000 and IEEE 802.11 radio networks are presented in [9]. Noting the strong influence of the capture effect, they empirically study the conditions under which it is appropriate to enable/disable carrier sensing in order to maximize throughput. In contrast, with E-CSMA, we build on the existing CSMA MAC, maintaining success probability profiles for each receiver with a focus on reducing collisions.

Berkeley MAC (B-MAC) [7] uses an adaptive RSS threshold for CCA, whereby each transmitter keeps an EWMA of the estimated channel noise floor by sampling the channel at times when the local transmitter assumes the channel to be idle. However, the CCA is open-loop and the RSS threshold, though adaptive, is common across all receivers.

In [2] the authors propose a class of packet scheduling methods that take wireless channel characteristics into consideration, in the context of a base-station/mobile client W-LAN architecture. Assuming the wireless links to various destinations are statistically independent, the base-station maintains a separate queue for each mobile host; the scheduling policy is based on the current estimated state of the channel between the mobile host and the base-station. The channel state to a given receiver is “bad” when a MAC layer acknowledgment is not received, following a data transmission, and “good” otherwise. Once the channel is marked as “bad”, an estimator for the residual burst period is required, but this is left as an open question. Fragouli, *et al* [3] combine the channel state dependent packet scheduling strategy in [2] with prior work on class-based queuing (Floyd and Jacobson) to provide class-based fairness in a wireless channel. An RTS/CTS exchange is used to predict the channel state.

Primarily targeting mobile wireless networks, Shen, *et al* [4] describe a low complexity channel prediction method using polynomial approximation. The next channel state can be predicted using  $M$  multiplications and  $M$  additions, where  $M$  is the number of the previous channel state samples. This technique relies on channel autocorrelation, and has a relatively short window of accuracy. It may not be a good fit to the often low data event duty cycle of sensor networks. There is a body of channel prediction work on wireless cellular systems, but these target multi-carrier/wideband channels (e.g., [8]) and can not be directly applied to the single narrow band radios of current sensor networks.

Sampath [5] proposes a channel state aware ARQ protocol, including a probabilistic channel prediction scheme for bursty wireless channels based on a variable structure learning automaton (VSLA). The automaton differs from our proposal in that it only supports binary feedback, and relies on high channel autocorrelation for an accurate short-term channel prediction, and operates only on a per neighborhood basis. Further, the work by Sampath, being primarily a mathematical development, neglects several practical considerations (e.g. assumes error-free feedback channel, zero byte feedback, time slotted channel, Gilbert-Elliot channel model).

#### V. CONCLUSION

We have presented E-CSMA, a fully distributed MAC framework targeted at increasing the chance of successful packet reception in wireless sensor networks using carrier sensing. The scheme couples low cost in-band feedback with a local state management strategy to empirically generate success probability distributions per receiver in accordance with locally observable channel conditions at the transmitter. A comprehensive transmit decision process leverages this learned correlation between channel state and reception success probability, balancing success probability with local queue drop probability to improve overall performance. The transmit decision process is adaptive to performance relative to the underlying CSMA MAC on which E-CSMA is built. We have presented experimental results from a 31 node sensor network testbed. Three performance metrics, average link tax, average link fidelity and average link power, are defined that show the ability of E-CSMA to balance packet loss with packet throughput, and quantify the advantage over plain CSMA. Through experimentation we have demonstrated the efficacy of E-CSMA in reducing average link tax with a relatively small decrease in average link fidelity across a range of network conditions. As a side result, we describe a simple experimental methodology for characterizing radio neighborhoods in a wireless sensor network testbed. Such a characterization, based on average numbers of well-connected neighbors and average numbers of non-existent radio links is helpful in understanding how results from a particular testbed might extend to other deployments. Though our focus is on the design space of wireless sensor networks, we believe our algorithms are more generally applicable and can improve performance in other classes of CSMA-based networks.

#### REFERENCES

- [1] V. Bharghavan, *et al.* MACAW: A Media Access Protocol for Wireless LANs. In *Proc. of ACM SIGCOMM*, pp. 212–225, 1994.
- [2] P. Bhagwat, *et al.* Using Channel State Dependent Packet Scheduling to Improve TCP Throughput Over Wireless LANs. In *Proc. of the IEEE Int'l Conf. on Computer Communication*, Jul 1997.
- [3] C. Fragouli, *et al.* Controlled Multimedia Wireless Link Sharing via Enhanced Class-Based Queuing with Channel-State-Dependent Packet Scheduling. In *Proc. of the IEEE Int'l Conf. on Computer Communication*, Jul 1998.
- [4] Z. Shen, *et al.* Short Range Wireless Channel Prediction Using Local Information. In *Proc. of the IEEE Asilomar Conf. on Sigs., Sys., and Comps.*, Nov 2003.
- [5] K. Sampath. A Channel-aware ARQ Protocol for Energy Efficient Wireless Networking. Masters Thesis, Stevens Institute of Technology, May 2004.
- [6] TinyOS homepage. <http://www.tinyos.net>.
- [7] J. Polastre, J. Hill, and D. Culler. Versatile Low Power Media Access for Wireless Sensor Networks. In *Proc. of the 2nd ACM Conf. on Embedded Networked Sensor Systems*, pp. 95–107. Baltimore, Nov 3–5 2004.
- [8] I. Wong, *et al.* Long Range Wireless Channel Prediction for Adaptive OFDM Systems. In *Proc. of the IEEE Asilomar Conf. on Sigs., Sys., and Comps.*, 2004.
- [9] K. Jamieson, *et al.* Understanding the Real-World Performance of Carrier Sense. In *Proc. of ACM SIGCOMM Workshop on Experimental Approaches to Wireless Network Design and Analysis*, Philadelphia, Aug 22 2005.
- [10] MICA2 and MIB600 Data Sheets. <http://www.xbow.com>.
- [11] CC1000 Data Sheet. <http://www.chipcon.com>.
- [12] J. Robinson, *et al.* Experimenting with a Multi-Radio Mesh Networking Testbed. In *1st Workshop on Wireless Network Measurements*, Trento, April 2005.
- [13] A. Arora, *et al.* ExScal: Elements of an Extreme Scale Wireless Sensor Network. In *Proc. of the 11th IEEE Int'l Conf. on Embedded and Real-Time Computing Systems and Applications*, Hong Kong, Aug 17–19 2005.
- [14] C.-c. Chen, *et al.* Self-Learning Collision Avoidance for Wireless Networks. In *Proc. of the 25th IEEE Conf. on Comp. Comms.*, Barcelona, April 23–29 2006.