

Measurement-Based Models of Delivery and Interference in Static Wireless Networks

Charles Reis
University of
Washington

Ratul Mahajan
Microsoft Research

Maya Rodrig
University of
Washington

David Wetherall
University of
Washington

John Zahorjan
University of
Washington

ABSTRACT

We present practical models for the physical layer behaviors of packet reception and carrier sense with interference in static wireless networks. These models use measurements of a real network rather than abstract RF propagation models as the basis for accuracy in complex environments. Seeding our models requires N trials in an N node network, in which each sender transmits in turn and receivers measure RSSI values and packet counts, both of which are easily obtainable. The models then predict packet delivery and throughput in the same network for different sets of transmitters with the same node placements. We evaluate our models for the base case of two senders that broadcast packets simultaneously. We find that they are effective at predicting when there will be significant interference effects. Across many predictions, we obtain an RMS error for 802.11a and 802.11b of a half and a third, respectively, of a measurement-based model that ignores interference.

Categories and Subject Descriptors

C.4 [Performance of systems]: Modeling techniques

General Terms

Measurement, performance

Keywords

Modeling, interference, RSSI

1. INTRODUCTION

Wireless networks such as 802.11 have enjoyed an unprecedented adoption rate in recent years, and their deployed base continues to grow. Originally envisioned to support mobile devices, wireless has also proved popular in more static settings that involve PCs and laptops in homes and offices because it removes the need for wires [2, 15]. A fundamental issue in these networks is interference, in which transmissions from one sender-receiver pair affect those of other pairs. Interference defines the spatial boundaries for spectrum reuse, and it directly impacts the assignment of senders to channels [18], network capacity [10], and routing choices [8].

It is thus startling that packet delivery under interference is poorly understood for real networks. Common protocols such as 802.11 make conservative scheduling decisions, serializing the transmissions of senders who can hear each other in case there is harmful interference. Most explorations of protocols with respect to inter-

ference use simple abstract models that may assume signal propagation is a simple function of distance, that coverage of radios is circular, that interference range is twice the transmission range, and so forth. Unfortunately, empirical data from experimental wireless networks has shown that all of these models are largely inaccurate [14, 17]. RF propagation in realistic environments is sufficiently complex that the only existing feasible method for estimating packet delivery between two nodes is to measure it.

As a response, there has been a shift towards experimental wireless networks in which packet delivery and higher level metrics have been measured for real radios working in particular network and protocol designs [3, 6, 11]. This approach is valuable because it mitigates the problem of unrealistic RF models; it has improved the understanding of wireless behaviors. But it trades one peril for another. Measuring experimental networks lacks a crucial benefit of analysis and simulation: the ability to explore a large space of configurations with reproducible results. It is too time-consuming to run experimental networks in a wide variety of settings, and results in one setting do not necessarily predict results in a different setting or even at a later time. This undermines the value of experiments by making it difficult to meaningfully compare results.

In our work, we ask whether it is possible to combine the strengths of both methods: can we use simple measurements on a wireless network to capture its RF characteristics and then predict how it will perform when running under different settings? This paper is a first step in that direction in which we derive a practical model of packet delivery under interference.

We use measurements on a running network to seed our model because predicting RF propagation in a complex environment is a hopelessly challenging task. To ensure that our model is applicable to real networks, we rely on only *received signal strength indicator* (RSSI) values and pair-wise delivery counts, since both are easily obtained using commodity wireless cards. We record this information when there is a single-sender as observed at all receivers, which requires N trials to obtain N^2 parameters for an N node network. We then formulate low-level models for packet reception and carrier-sense by relating the traditional notion of SINR (signal to interference plus noise ratio) to our measurements. We investigate 802.11 characteristics, both in a controlled setting with attenuators and on a building network, to provide a foundation for the models. These models are in turn fed into a higher-layer system model that predicts packet delivery and interference for the same node placements but different sets of transmitters. We view this as a foundation for exploring other higher-level design choices, such as RTS/CTS exchanges, routing and channel assignments.

We have evaluated the base case of our models, in which two senders compete to transmit fixed-size broadcast packets at a set bit-rate, on our in-building 802.11 testbed. We find that they are effective at identifying the situations in which there will be significant interference and predicting the magnitude of the effect. Across many random trials the RMS error of our throughput predictions is

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGCOMM'06, September 11–15, 2006, Pisa, Italy.

Copyright 2006 ACM 1-59593-308-5/06/0009 ...\$5.00.




Figure 1: Our wireless testbed, consisting of fifteen 802.11 a/b/g nodes. The width of the building is 184 ft.

11% of the channel bitrate for 802.11a and 9% for 802.11b. This is comparable to the temporal variability we observe in the wireless medium. In contrast, the RMS error of a naive model that ignores interference is two or three times higher, 24% for 802.11a and 31% for 802.11b, with predictions that are often poor when there is significant interference. To further demonstrate the utility of our models, we show how they can be used to predict the conflict graph of a network. We view these results as promising, and are hopeful that future work will extend them to cover a larger fraction of the many transmission options: more than two senders, mixtures of packet sizes and rates, unicast traffic with acknowledgements and retransmissions, and so on.

The rest of this paper is organized as follows. The network environments we experiment with are described in Section 2. We study wireless characteristics to support our models in Section 3. We develop our models in Section 4. In Section 5, we evaluate our models. We then present related work and conclude.

2. EXPERIMENTAL PLATFORMS

The work described in this paper uses measurements for three purposes: (i) to gain an understanding of the characteristics of wireless networks (ii) as inputs to the modeling process we propose, and (iii) as the basis of an experimental evaluation of the accuracy of our model. We now describe the environments we measured for these purposes.

2.1 Testbed Experiments

The primary vehicle for our empirical studies is an indoor testbed of fifteen stationary PCs shown in Figure 1. The testbed is located on the third floor of our building, mimicking a deployment of wireless nodes in an office scenario. Each testbed node has an 802.11 a/b/g Atheros card that we operate using the “stripped” MadWiFi driver [4].

We depend on *received signal strength indicator* (RSSI) values that are reported by all commodity wireless cards. RSSIs are estimates of the signal energy at the receiver during packet reception, measured during the PLCP headers of arriving packets and reported on proprietary (and different) scales. Our Atheros cards, for example, report RSSI as $10\log_{10}(\frac{S+I}{n})$, where S is the strength of the incoming signal, I is the interfering energy in the same band, and n is a constant (-95 dBm) that represents the “noise floor” inside the

radio. Atheros RSSI is thus dB relative to the noise floor. To give results that are independent of card vendors, we transform RSSI values to *received signal strength* (RSS) values that give absolute energy levels. That is, RSS is defined to be $S + I$. Additionally, we abuse notation slightly by reporting specific values in (log scale) dBm units, as is common practice, while writing formulas such as the two above to manipulate (non-log scale) quantities such as mW.

Because our focus is raw packet delivery probabilities, we operate in broadcast mode, which suppresses MAC-level features such as ACKs and RTS/CTS exchanges.

The testbed operates in a noisy environment, with many active people and energy sources, including a building-wide 802.11b/g production network. We use Channel 3 for 802.11b experiments, which is separate but non-orthogonal to Channels 1 and 6 of the official network (which also uses Channel 11). This network acts as a realistic source of external interference for our testbed, given current dense urban deployments of wireless networks. It complicates our efforts to predict wireless performance, but we find predictability despite it.

2.2 Attenuator Experiments

To identify the causes of variability in live wireless networks, we also conduct controlled experiments with the same Atheros cards. This is valuable for verifying key characteristics of the hardware, while limiting the impact of competing energy sources.

Specifically, we disconnect the antennas from two cards and attach shielded SMA cables, along with both fixed and variable attenuators to control the signal strength at the receiver. We surround the receiver with RF shielding foam to minimize influences from the local environment, since these cards have some ability to receive packets even without an attached antenna. Our configuration follows sensitivity experiment guidelines provided by Intersil [1].

3. WIRELESS CHARACTERISTICS

In this section, we characterize wireless delivery in our testbed and via attenuator experiments to identify the key effects for our models. In successive subsections, we study the feasibility of predicting packet delivery using practical measurements from commodity hardware, the nature and impact of external interference, and temporal stability of a wireless environment. We perform experiments for 802.11a and 802.11b but present results only for the latter due to space limitations.

3.1 Packet Delivery and RSSs

The SINR ratio is widely used in the literature to model packet delivery probabilities: packets are successfully received if $\frac{S}{I+n}$ is above a certain threshold, and otherwise are not. Applying this model in practice, though, presents several problems. For one, the SINR model itself is only approximate, as it ignores factors such as multipath [3]. Moreover, commodity wireless cards do not report the information required to use it. For instance, our cards report only their version of RSS, the minimum feedback allowed by the 802.11 standard. Some other cards also report an estimate of I by measuring energy in the air when no packets are being sent, but this estimate may be inaccurate during packet delivery, and we will show that it is not necessary in any case.

Because our ultimate goal is to construct a *practical* model, we only employ measures that are widely available. We therefore turn our attention to RSS and its use in predicting packet delivery probabilities. We begin with experiments performed in a controlled setting, followed by those taken in an uncontrolled environment.

Controlled setting When the interference I is negligible, correctly measured RSS should perfectly reflect packet delivery. We




Figure 2: (a) *Delivery probability for Atheros cards as a function of mean RSS, in attenuator experiments.* (b) *The CDF of observed RSS values for a particular attenuation, as circled on the left graph.*

use our attenuator setup to test this hypothesis. We repeat five minute floods of large broadcast packets (1084 bytes, including headers) with a set of nine different signal attenuations, from minimal attenuation to a level that prevents any packets from being received. The entire process is repeated for three separate rounds.

Figure 2(a) shows the probability of successful delivery as a function of measured RSS. Each point represents mean values over a given five minute round. There is a narrow transition range of less than 5 dBm, below which no packets are received, and above which packets are received with near certainty. This is an encouraging outcome, as RSS displays a simple relationship with delivery probability.

Figure 2(a) suggests that RSSs averaged over thousands of packets are good predictors of delivery probability. Figure 2(b) examines a much smaller measurement interval. It plots the cumulative distribution function (CDF) of RSSs measured with a particular attenuation, corresponding to the circled points on Figure 2(a). Most reported values are quite close to each other, but there remains some variation¹. For our purposes, this means the RSS of a single packet cannot be taken in isolation and we must sample multiple packets for a measure that accurately reflects packet delivery.

Uncontrolled Setting Next, we run experiments on our testbed to study whether RSS is predictive of packet delivery in an uncontrolled environment as well. Here, individual nodes transmit large broadcast packets (1084 bytes) for two minutes, while all other nodes record the packets they receive. We repeat for three rounds, each conducted at night to minimize variation.

Figure 3 shows delivery probability (averaged over the two minutes) as a function of mean RSS. It includes points for all receivers in the testbed for all three rounds. With all receiver data combined, there is a weak correlation between the two quantities: some nodes exhibit high delivery probabilities at an RSS of -80 dBm, while others receive only 50% at an RSS of -70 dBm. However, there is a stronger correlation when viewing receivers individually, as shown by the lines for nodes 8, 14, and 15. While the thresholds differ between the nodes, each node exhibits a fairly clear relationship between delivery probability and RSS. The differences in the relationships are likely due to different distributions of external interference experienced by these nodes, which we investigate shortly.

We can also see from the scatter plot in Figure 4 that distance is a poor indicator of delivery probability compared to per-node RSS, as found by earlier work [14, 3, 17, 7]. To emphasize the problem,

¹We hypothesize the variation is external interference due to imperfect RF shielding. Similar experiments on less sensitive Prism cards showed almost no variation in RSS.




Figure 3: *Scatter plot of delivery probability as a function of mean RSS, for all receivers. Lines connect points for three particular receivers; despite the weak correlation between delivery probability and RSS across all receivers, the correlation is strong for individual nodes.*




Figure 4: *Scatter plot of delivery probability as a function of distance, for all receivers. Points for Node 1 are highlighted.*

we have highlighted the points for node 1; they present no clear pattern that could be used for prediction in distance based models.

In conclusion, we observe a relationship between measured RSS and delivery probability, but that the exact relationship between the two can vary substantially across nodes. Note that this is not inconsistent with RSS measurements from other studies, e.g., Roofnet operates in a setting where inter-node distances caused delay spreads to exceed the engineering margin [3], and most other studies combine RSSs across receivers, despite the fact that they may have different sensitivities [12].

3.2 Nature of External Interference

We now study external interference: energy that is not caused by packet transmissions in the system under our control. Specifically, we study the region over which it has a significant effect. Like all signals, this energy will be attenuated with distance.

External interference might have a measureable effect on packet delivery across multiple nodes or be confined to a single node. To assess this, we look for correlation in packet losses at different nodes. In multiple trials, each with a different sender, we had the sender transmit broadcast packets with increasing sequence numbers. All other nodes logged which packets they successfully received. We found loss between the majority of pairs of receivers to be roughly independent, such that loss at one receiver was not gen-




Figure 5: (a) Log-linear PDF of loss burst length for all receivers with loss probabilities less than one half. (b) Probability that packet i is lost, given that packet $i - k$ was lost. The gray line shows aggregate data for all receivers with loss probabilities less than one half, and the black line highlights Node 9. Dashed lines show the overall loss probabilities for each.

erally a good indicator of loss elsewhere. This is consistent with other studies [16].

We also performed a simple check to test whether external interference is a property of specific machines (as it is generated by their components) or specific locations (as it is generated by other sources in the environment). To do so, we focused on one of the pairs of nodes with asymmetric delivery probabilities, despite homogeneous hardware and software configurations. We first swapped just wireless cards between the two machines, and then the two machines themselves between their locations. We found that the asymmetry was tied only to the location: it still held, in the same direction, despite swapping equipment.

These experiments lead us to conclude that external interference is primarily a local phenomenon in our testbed. For our purposes, it has to be measured independently at each node.

3.3 Stability Across Time

To be effective, a measurement-based model must be guided by an understanding of the system behavior across time. This determines, for instance, how long the system needs to be measured to obtain representative values and how far into the future a set of measurements can be used to make predictions. Here, we study stability on different timescales.

3.3.1 Short Term Stability: Loss Events

It is well known that wireless networks tend to have bursty losses. The length and frequency of the bursts determine how well measurements taken over short time scales are likely to predict the immediate future. To study them, we broadcast packets from node 12 for one hour at night. Figure 5(a) shows a PDF of loss burst length, measured across all receivers with a loss probability under 50%. (Receivers with larger loss probabilities would appear as very long bursts of unpredictable size.) There is a somewhat less than geometric decline in burst length (as the y-axis is log-scale). This suggests that losses are slightly bursty, but only over relatively short intervals. Indeed, we find that 92% of all bursts have a length of 3 packets or less, though there are bursts that last up to 76 packets.

An alternative view of this data is given in Figure 5(b). It tries to uncover dependence among non-consecutive losses. The graph shows the aggregate data for all receivers with loss probabilities less than 50% (in gray), plus the data for node 9 (in black) as an individual example. The dashed lines show the overall loss probabilities. The solid line shows the auto-conditional loss probability:




Figure 6: The CDF of variability in delivery probability measurements at different time scales.

that packet i was lost, given that packet $i - k$ was lost. If losses were fully independent, this probability would equal the overall loss probability for all k . Instead, we see that for small values of k the auto-conditional loss probability tends to be higher than average (particularly for node 9), which reflects bursty losses. Moreover, it takes a separation of around 100 packets, many times the average burst length, before a current packet loss is irrelevant to the future. This indicates that it is likely that packet loss bursts are often followed by other bursts, separated by only a few received packets. In intervals larger than 100 packets, however, losses appear to be largely independent.

Finally, there is clearly a periodic effect in which loss probability is slightly increased. As hypothesized by Miu *et al.* [16], this is potentially due to interference from beacon frames emitted by the building’s official wireless network on a non-orthogonal channel. We find the peaks to occur approximately 10 times a second when we translate packets into time for our setup. This is a common frequency for beacon frames.

In conclusion, we observe that losses do tend to occur in small bursts, but can be treated as independent for larger time intervals of several seconds or minutes. This means that we must measure the network at least for such timescales, and our predictions will be stable at such timescales.

3.3.2 Longer-Term Stability: Stationarity

We now look at the variability among consecutive measurement intervals separated by differing amounts of time. Figure 6 is a CDF of the difference between the average delivery probability for all sender-receiver pairs measured for one time interval and the next; it would be a vertical line at zero if all intervals were identical. The solid line depicts a time scale of 10 seconds, for an experiment in which senders transmitted for two minutes at a time. The dashed line compares delivery probabilities averaged over two minutes between rounds separated by about two hours. Finally, the dotted line compares delivery probabilities, averaged over two minutes, between an experiment conducted late at night and a second experiment the following day.

We see variability to be small at the smallest time scale, and to rise noticeably over the longest time scale; the latter also includes day/night changes in patterns of activity that reduce predictability. To quantify the difference from vertical, we compute the root mean square error (RMSE) for each scale: it is 4%, 12%, and 18%, respectively. We conclude that there is enough similarity between measurements to make useful predictions over moderate time scales of minutes to hours, but that prediction accuracy will be degraded for longer periods of time.




Figure 7: Delivery probability (averaged over one second intervals) and RSS over time. The gray boxes indicate when a nearby microwave oven was turned on.

3.3.3 Atypical Events

We also observed periods of atypical network behavior in which measurements are not close to representative. We report on two causes and their effects below. Fortunately, these periods tend to be rare and do not detract from overall predictability. To quantify this, we computed how often delivery probability and RSS values varied from their 10 second average by more than 10% of their respective ranges (i.e., a delivery probability of 0.1 and a RSS value change of 6) for two hand-sampled hour-long experiments with a given sender. We found that delivery probability was atypical 4.6% of the time, and observed RSS values were atypical 3.7% of the time.

Non-802.11 energy sources In reviewing an early set of experiment results, we noticed a period during which delivery probabilities for some nodes fell dramatically, even though RSS values showed no similar trend. Because of the locations of the affected nodes, we guessed that the cause might be interference from a microwave oven. We succeeded in reproducing the effect with a controlled experiment in which the microwave was switched on and off while one node sent to another. Figure 7 plots the observed RSS and packet delivery probability at the receiver as a function of time. It shows that microwave activity significantly degrades delivery without a clear effect on RSS, for the packets that were still received.

Shadowing We also observed instances where RSS dropped together with delivery probability. We believe that these events result from signal obstructions, or “shadowing,” of the receiver due to macro-scale changes in the environment. We were able to experimentally support this hypothesis by introducing obstacles between the sender and receiver.

In conclusion, while measured RSS is generally predictive of delivery probability, various transient sources of interfering energy can distort that relationship. One needs to be careful that this relationship is not measured during such an event. We measure over multiple rounds because of this in Section 5. It may also be possible to develop online heuristics to detect atypical periods. Note that persistent sources of interference are not an issue since they should be reflected by measurements.

3.4 Summary

Overall, we conclude that there is a general relationship between measured RSS and delivery probability in real networks; the challenge is to make it precise enough to use for prediction. The major

source of external interference seems to be the local environment, which varies substantially across nodes. While wireless networks exhibit substantial variability, measurements of average behavior over even relatively short time periods tend to be stable, even for widely separated intervals. Finally, there may be infrequent atypical periods during which measurements are not representative.

4. MEASUREMENT-BASED MODEL

We now develop PHY models for wireless delivery with interference by recasting the classical notion of signal-to-interference-plus-noise-ratio (SINR) in terms of our observable measurements. These PHY models are combined with higher layer models, such as the simple MAC model we provide, to predict the performance of a static wireless network with arbitrary sets of interfering senders based on past measurements with individual senders.

4.1 Operation

Our models operate as follows:

1. The RF profile of the network is measured. Each of the N senders broadcasts packets in turn, while the other nodes record the number of received packets along with their RSS values, creating N^2 data points. We use special-purpose traffic for these measurements, though in a deployed network they could be gathered using application traffic and 802.11 sequence numbers.
2. The *PHY receiver model* we derive below is used with the RF profile to compute the probability a packet is correctly received from a given sender in the presence of competing transmissions.
3. The *PHY deferral model* we derive below is used with the RF profile to compute the probability that a sender will sense competing transmissions and defer its own transmission.
4. *MAC and traffic models* build on the PHY models to predict the performance of the network in a specified configuration. MAC models capture higher-layer protocol rules, e.g., CSMA/CA. Traffic models specify the sets of nodes that compete to send packets at the same time, with possibly different power levels than measured in the RF profile but the same packet size and transmit rate. This would be implemented as a packet-level simulator in the general case. For a concrete exposition in this paper, we provide closed-form approximations for the analytically tractable case of two broadcast senders running CSMA/CA, at a fixed power level and transmit rate.

Our key contribution is the PHY models. They are based on the classical signal-to-interference-plus-noise-ratio (SINR) model interpreted in our context. We build them in stages below and then summarize how they are used for prediction.

4.2 Recasting the SINR Model

We recast the SINR model that is used widely in the literature so that we can use it with our measurements. The version we use gives the probability p_r that a receiver node r can decode a packet transmitted from a sender node s :

$$p_r(\mathcal{A}_{sr}(P_{sr})) = Prob \left[\frac{\mathcal{A}_{sr}(P_{sr})}{I_r + n_r} \geq \delta_r \right] \quad (1)$$

In this equation², the sender node is visible at the receiver in terms of its signal power, P_{sr} , as attenuated along the path from s to r . The function $\mathcal{A}_{sr}(\cdot)$ models this attenuation so that $\mathcal{A}_{sr}(P_{sr})$ is the signal strength at the receiver. The interference from the environment experienced at r while trying to receive the packet is I_r . We take it to be a receiver-specific random variable, and it causes p_r to be a probability over packets. I_r does not include the thermal noise floor, n_r , which is generated by the node itself and assumed constant. There is no dependence on packet length because we take it to be fixed. Finally, for a given transmit bit-rate and modulation, δ_r is the SINR threshold of the radio at r above which it can successfully decode a packet.

Unfortunately, we cannot use equation (1) directly to compute delivery probabilities. There are two difficulties. The first is that attenuation is a hopelessly complicated function that depends on many details of the environment. To avoid this complexity, much work uses generic distance-based models in which $\mathcal{A}_{sr}(P) \approx d_{sr}^{-\alpha} P$, for $2 \leq \alpha \leq 4$, where d_{sr} is the distance between s and r . $\alpha = 2$ corresponds to free space, while higher values reflect denser and more irregular environments such as office buildings. But in real networks, it is a poor predictor of packet delivery [14] and hence not appropriate for our purposes.

The second difficulty is that we cannot obtain the parameters of the SINR model directly using commodity hardware. We need the strength of the incoming signal S and the distribution of interference I . However, the information widely accessible across wireless cards is the reported RSS value³. RSS is a measure of the energy at the receiver during decoding, and so conflates signal strength and interference: $RSS = S + I$. Additionally, RSSs are available only for successfully received packets, increasing the challenge of estimating the interference distribution.

Despite these problems, we can use measurements to leverage the SINR model. Our experiments in Section 3 indicate a strong relationship between RSS and delivery probability. The exact form varies with receivers and network deployments, but we still expect the signal strength, interference and noise relationships to be consistent with the SINR model. To capture these relationships, we use the measurements described above to create an RF profile of the network. The measured RSS values and packet counts allow each receiver to compute several parameters:

- The mean RSS for packets received at r from another node s , which we denote \bar{R}_{sr} , can easily be computed by averaging.
- We can also estimate the average external interference at a node r , which we denote \bar{I}_r , from our earlier observation that most of the variation in RSS values stems from interference. If we assume that at least one packet across all senders was received when the external interference was almost zero, \bar{I}_r can be estimated by the mean excess of the RSS values from individual senders above their minimum observed values.
- Finally, we obtain the curves of the delivery probabilities associated with the mean RSS levels of the different senders. We denote this with $\hat{p}_r(\bar{R}_{sr})$, the probability of correctly decoding packets with an RSS of \bar{R}_{sr} at r . This is an approximation of the actual delivery probability because of the averaging that has been performed, and because RSS is only reported for packets that are successfully received. Nonetheless, it is appealing to think of $\hat{p}_r(\cdot)$ as a sur-

rogate for $p_r(\cdot)$ where the function domain has been transformed from signal power to RSS.

To work with these measurements, we observe that incoming signal strength is approximately constant for stationary nodes at a given power level. Thus, we can estimate the true incoming signal energy at the receiver, which we denote S_{sr} :

$$S_{sr} = \mathcal{A}_{sr}(P_{sr}) \approx \bar{R}_{sr} - \bar{I}_r \quad (2)$$

We can now recast the SINR model to fit our measurements. Substituting (2) into (1) and using $\hat{p}_r(\cdot)$ to estimate $p_r(\cdot)$ we have:

$$p_r(S_{sr}) = Prob \left[\frac{\bar{R}_{sr} - \bar{I}_r}{I_r + n_r} \geq \delta_r \right] \approx \hat{p}_r(\bar{R}_{sr}) \quad (3)$$

All of the terms here are constants except for I_r . We rewrite the equation to expose the distribution of interference in a form we will use shortly:

$$p_r(S_{sr}) = Prob \left[I_r \leq \frac{\bar{R}_{sr} - \bar{I}_r}{\delta_r} - n_r \right] \approx \hat{p}_r(\bar{R}_{sr}) \quad (4)$$

That is, the SINR model enables us to use RSS measurements to estimate the delivery probability as a function of interference, which is the prime cause of variation in packet delivery.

4.3 PHY Receiver Model

The PHY receiver model predicts the probability that a receiver r will correctly decode a packet transmitted by a sender s while packets are being sent by other, competing senders. We can apply the SINR model to this situation by treating the energy from the competing senders as adding to the external interference. This may seem surprising because it ignores temporal considerations such as whether the competing packet starts its transmission slightly before or after the sender. But this reflects “capture effects” in which real radios lock onto stronger signals regardless of when they occur [13]. Thus, we can restate (3) and (4) to give delivery probabilities when there is a competing sender t as follows:

$$\begin{aligned} p_r(S_{sr}, S_{tr}) &= Prob \left[\frac{\bar{R}_{sr} - \bar{I}_r}{\bar{R}_{tr} - \bar{I}_r + I_r + n_r} \geq \delta_r \right] \\ &= Prob \left[I_r \leq \frac{\bar{R}_{sr} - \bar{I}_r}{\delta_r} - (\bar{R}_{tr} - \bar{I}_r) - n_r \right] \end{aligned} \quad (5)$$

A key insight is that we can evaluate this probability with the RF profile we already have available. According to (4), each measured single-sender delivery probability, $\hat{p}_r(\bar{R}_{sr})$, gives the probability that the interference I_r is below a corresponding threshold. We can now compute this threshold for the case of competing senders by using (5), but we do not have measured delivery probabilities for multiple senders. Instead, we can find a hypothetical single-sender RSS, RX_{sr}^t , that corresponds to the interference threshold for competing senders. Once we have this new RSS, we can predict the delivery probability for competing senders using the single-sender probabilities $\hat{p}_r(\bar{R}_{sr})$.

From (4) and (5), we want $\hat{p}_r(RX_{sr}^t)$ such that:

$$\frac{RX_{sr}^t - \bar{I}_r}{\delta_r} - n_r = \frac{\bar{R}_{sr} - \bar{I}_r}{\delta_r} - (\bar{R}_{tr} - \bar{I}_r) - n_r \quad (6)$$

Solving for RX_{sr}^t , we obtain:

$$RX_{sr}^t = \bar{R}_{sr} - \delta_r (\bar{R}_{tr} - \bar{I}_r) \quad (7)$$

²We abuse notation slightly by reporting specific I_r , n_r and RSS values in (log-scale) dBm units, as is common practice, but writing all formulas in terms of (non-log scale) quantities such as mW.

³Cards actually provide manufacturer specific RSSI values. As explained in Section 2.1, we transform these into somewhat more universal RSS values.

Our prediction for the delivery probability of packets from s when t is also transmitting is then the RF profile value $\hat{p}_r(RX_{sr}^t)$. As a practical matter, we will not have measured $\hat{p}_r(\cdot)$ at the precise RSS that is needed. To estimate it, we piecewise interpolate the $\hat{p}_r(\bar{R}_{sr})$ data points. It is also the case that \bar{R}_{tr} will not be available if r received no packets from t . However, we can then treat it as zero without affecting accuracy in practice. This is because r is likely to receive at least some packets from t , given that it receives packets from s , unless \bar{R}_{tr} is significantly smaller than \bar{R}_{sr} . So, if \bar{R}_{tr} is not available, then the term containing it is likely to be negligible and can be omitted.

We have considered one competing sender above for simplicity, but it is straightforward to extend the equations for multiple simultaneous senders by introducing additional interference terms. We can also factor in changes in the transmit power level of the senders relative to the single-sender measurements at the same time. This is because changing the power of the transmitter by some factor causes the same change in the power of the received signal. If we change the sender power by a factor α , and we have competing senders $t_i, i \geq 0$, each of which changes their power by a factor α_i , then we can restate (5) and (7) as:

$$p_r(S_{sr}, S_{t_0r}, \dots) = Pr \left[I_r \leq \frac{\alpha(\bar{R}_{sr} - \bar{I}_r)}{\delta_r} - \sum_i \alpha_i (\bar{R}_{t_i r} - \bar{I}_r) - n_r \right] \quad (8)$$

$$RX_{sr}^{t_i} = \alpha \bar{R}_{sr} + (1 - \alpha) \bar{I}_r - \delta_r \sum_i \alpha_i (\bar{R}_{t_i r} - \bar{I}_r) \quad (9)$$

Our complete PHY receiver model predicts the delivery probability by computing $RX_{sr}^{t_i}$ using (9) and then looking up the associated delivery probability in the RF profile for receiver r . That is:

$$p_r(S_{sr}, S_{t_0r}, \dots) = \hat{p}_r(RX_{sr}^{t_i}) \quad (10)$$

As a caveat, we cannot scale up the power for senders for which we did not receive sufficient packets to estimate their mean RSS, \bar{R} , in the RF profile.

4.4 PHY Deferral Model

The PHY deferral model gives the probability that a node will defer its own transmission to competing transmissions because it senses that the channel is busy. We use a similar approach as above to derive it.

We assume that a node senses the channel busy when the total energy it receives is above the CCA (clear-channel assessment) threshold, β_s , which depends on the radio. Now consider a node s preparing to send to some other node when there are competing senders $t_i, i \geq 0$ with scaled power α_i as before. The probability $p_s(\cdot)$ that s senses the channel busy and defers is:

$$p_s(S_{t_0s}, \dots) = Prob \left[\sum_i \alpha_i (\bar{R}_{t_i s} - \bar{I}_s) + I_s > \beta_s \right] \quad (11)$$

We can rewrite this in terms of interference as we did with (4):

$$p_s(S_{t_0s}, \dots) = 1 - Prob \left[I_s < \beta_s - \sum_i \alpha_i (\bar{R}_{t_i s} - \bar{I}_s) \right] \quad (12)$$

As before, we can use (4) to find a hypothetical single-sender RSS $TX_{t_i s}$ that produces deferrals equivalent to those in (12). We start with:

$$\frac{TX_{t_i s} - \bar{I}_s}{\delta_s} - n_s = \beta_s - \sum_i \alpha_i (\bar{R}_{t_i s} - \bar{I}_s) \quad (13)$$

Then, our prediction for the likelihood that node s will defer its transmission is the complement of the delivery probability associated with $TX_{t_i s}$ in the RF profile:

$$TX_{t_i s} = \delta_s \left(\beta_s - \sum_i \alpha_i (\bar{R}_{t_i s} - \bar{I}_s) + n_s \right) + \bar{I}_s \quad (14)$$

$$p_s(S_{t_0s}, \dots) = 1 - \hat{p}_s(TX_{t_i s}) \quad (15)$$

4.5 MAC and Traffic Models

To apply our PHY models, we need higher layer MAC and traffic models that capture the rest of the system by specifying which nodes have packets to send at what times, and the protocol rules by which they attempt to transmit packets. In the general case, this might be done with a simulator that models a traffic workload and higher-layer protocol rules, such as 802.11 CSMA/CA, exponential backoff, and acknowledgements with retransmissions. Here, we give a model for the analytically tractable case of two competing broadcast senders that run CSMA/CA. This model is only approximate as it ignores various corner-cases. However, it provides a self-contained example for this paper, and our evaluation shows that it already has sufficient predictive power to validate the base case of wireless interference.

In 802.11, CSMA/CA with broadcast senders works as follows. Each node senses whether the channel is busy, i.e., whether the received energy is above the CCA threshold. It defers if so. To avoid collisions when the channel becomes free, each node randomly picks a number of fixed-time slots in the range $[0, W - 1]$.⁴ The node counts down this many free slots to pass before transmitting the packet, pausing the countdown when the channel is busy. Ideally, transmission should occur immediately when the countdown stops. Our observations, though, suggest that in practice there is a turnaround time equal to about a slot time.

The countdown operations to avoid collisions are effectively races among the nodes. We are interested in average behavior when two nodes, s and t , continuously send fixed-size packets. To estimate the probability of a collision, consider a moment when a node, say s , has just finished a transmission. At that point t has some remaining countdown time $C_t < WT$, where T is the slot time, and s picks a new countdown time, C_s . Because slot boundaries are not synchronized, we model C_s as uniformly distributed over $(0, WT)$. A collision occurs on the next packet transmission if C_s and C_t are within a turnaround time, that is, if C_s falls in a window of length $2T$ centered around C_t . Because the nodes are in a symmetric situation, each wins the race half the time when there are no collisions. Thus, for $n \in \{s, t\}$:

$$\begin{aligned} Prob[n \text{ collides}] &= \frac{2T}{WT} = \frac{2}{W} \\ Prob[n \text{ wins}] &= \frac{1}{2} - \frac{1}{W} \\ Prob[n \text{ loses}] &= Prob[n \text{ wins}] \end{aligned} \quad (16)$$

We are interested in the number of packets sent by each combination of senders, and the number of packets received from each sender by each other receiver. To obtain this performance data, we

⁴There is no exponential backoff for broadcasts in 802.11.

Inputs to the model		Derivation method
$\hat{p}_r(\cdot)$	Function that maps RSS to delivery probability at r	RF profile measurements
\bar{R}_{sr}	Average RSS observed at r when s sends alone	
\bar{I}_r	Average external interference (variation in RSSs) at r	
n_r	Thermal noise inside the card	Hardware-specific (constant)
δ_r	SINR threshold for successful reception	
β_r	CCA threshold for deferral	

Table 1: Summary of the inputs to the PHY models. All inputs are per-node.

combine the race outcomes with the PHY models. A node s will send alone when it both wins the race and the other node t defers to the energy that s transmits. A node s will send at the same time as a competitor t when it either collides, or loses the race but fails to defer to the energy that t transmits. Assuming that a packet transmission time is much larger than the largest countdown delay, the fraction of time that packets will be transmitted by different combinations of senders s and t is approximately:

$$\begin{aligned}
Frac[s \text{ sends alone}] &= Prob[s \text{ wins}] (1 - \hat{p}_t(TX_{st})) \\
Frac[t \text{ sends alone}] &= Prob[t \text{ wins}] (1 - \hat{p}_s(TX_{ts})) \\
Frac[s, t \text{ send}] &= Prob[s \text{ collides}] \\
&\quad + Prob[s \text{ wins}] \hat{p}_t(TX_{st}) \\
&\quad + Prob[t \text{ wins}] \hat{p}_s(TX_{ts})
\end{aligned} \tag{17}$$

A node r receives a packet from a sender when it sends, either alone or in combination with other nodes, depending on the associated delivery probability. The fraction of time r receives from s and t when both nodes attempt to transmit continually is then:

$$\begin{aligned}
Frac[r \text{ receives } s] &= Frac[s \text{ sends alone}] \hat{p}_r(RX_{sr}) \\
&\quad + Frac[s, t \text{ send}] \hat{p}_r(RX_{st}^t) \\
Frac[r \text{ receives } t] &= Frac[t \text{ sends alone}] \hat{p}_r(RX_{tr}) \\
&\quad + Frac[s, t \text{ send}] \hat{p}_r(RX_{tr}^s)
\end{aligned} \tag{18}$$

We can then use these fractions to compute other performance metrics of interest. For example, the delivery probability from s to r is $Frac[r \text{ receives } s]$ divided by the sum of $Frac[s \text{ sends alone}]$ and $Frac[s, t \text{ send}]$. Similarly, throughput is the product of the capacity of the channel and $Frac[r \text{ receives } s]$.

4.6 Summary

We have derived two PHY models for reception and deferral by combining measurements with an adaptation of the classic SINR model. The inputs required by our PHY models are listed in Table 1. They use an RF profile that is measured by having each node send by itself and observing the RSSs and delivery probability of packets received at other nodes. At each receiver, this gives an average RSS for each sender, an RSS to delivery probability curve across senders, and the average external interference as the variation in RSSs from the same sender. This requires N trials, one for each sender in an N node network, and results in $N - 1$ tuples at each receiver.

Given the single-sender RF profile and hardware constants, the PHY receiver and PHY deferral models (Equations 10 and 15) then predict the likelihood of reception and deferrals in the same network when multiple nodes send at once. These likelihoods are fed into higher-layer MAC and traffic models that complete the rest of the system under study. MAC models fold in higher-layer protocol

behavior, and traffic models specify the workload placed on the network. We have given closed-form equations for the simple system of two sender broadcast with CSMA/CA. The output of the MAC and traffic models is a prediction of the performance of the system.

5. MODEL EVALUATION

In this section, we test an instantiation of our models on our wireless testbed, to show that the predictions are accurate enough to be useful, despite the inherent variability of a wireless environment and the inaccuracies of measuring it using commodity hardware. We evaluate both overall and component prediction accuracy, and show how our models can be applied, by using them to infer conflict graphs.

Instantiation of the Model As input to the PHY models, we gather the RF profile of our network by having nodes take turns to broadcast large (1084 bytes) packets for two minutes each. While one node sends, the rest log packets and their observed RSSs. This takes $O(N)$ time and gathers $O(N^2)$ parameters for an N -node network. To filter out any atypical events, we conduct three such rounds, separated by roughly 4 hours, and use the median values as input to the model. This data gives us \bar{R}_{sr} , \bar{I}_r , and $\hat{p}(\bar{R}_{sr})$ directly. As mentioned in Section 4, such data could potentially be derived from application traffic in a deployed network.

To obtain the RSS versus delivery function $\hat{p}(\cdot)$ for each receiver, we simply piecewise linear interpolate the measured delivery probability and mean RSS pairs that were observed from each of the other senders. The other PHY model parameters depend on the card and driver, and can be computed once per type of radio; they could also be supplied by the manufacturer. Atheros cards report a constant noise floor n of -95 dBm. To approximate the SINR threshold δ for our Atheros cards, we use our attenuator experiments, where interference is near zero and δ is thus $\frac{S}{n}$. The knee of the delivery curve in Figure 2 gives the value of δ as 2.5 dB. For the CCA threshold β , we first measure the deferral probability between each pair of senders, as described in Section 5.3. We then plot the deferral probability between all sender pairs against RSS. The knee of the curve represents the threshold above which the radios in our testbed defer. Our results yield a β of -81 dBm.

The other parameter we need is for our MAC model. For window size W in Equation 16, we use 802.11 standard values of 16 for 802.11a and 32 for 802.11b.

Experimental Methodology We use the closed-form expressions for a two sender broadcast CSMA/CA system given in Section 4.5 to evaluate our PHY models. These expressions predict the throughput and delivery probability at a third receiver when two senders compete to send packets. Because there are a large number of possible pairs of senders, we select a subset of 12 pairs that are representative of senders with high (i.e., over 80%), middle, and low (i.e., under 5%) delivery probabilities to each other; each pair is used to make predictions to 14 possible receivers.⁵ The

⁵Node 5 in Figure 1 was excluded due to a hardware problem.

two-sender measurements are taken over three separate rounds, interleaved with RF profile collection. The entire experiment, including single and two sender measurements for both 802.11a and 802.11b, took approximately 12 hours to run. We compare our predictions to values observed in each of these instances, aggregated over the three rounds. After removing uninteresting cases where the receiver is out of range of both senders, we make a total of 528 predictions for 802.11a and 828 predictions for 802.11b.

Other MAC models and traffic mixes are clearly of interest, but we leave experiments with them for future work. The number of multiple sender experiments is combinatorial, and two senders are the most important base case for interference effects. We also use a bitrate of 6 Mbps for 802.11a and 1 Mbps for 802.11b; our current experiments do not explore variations in transmission bit rate or packet size.

We judge the accuracy of our predictions by comparing them to actual two-sender measurements. To put these results in context, we compare them with results for two other models. The first is a *naive* model that optimistically ignores the presence of the second sender, assuming it will not affect either the first sender or the receiver. This model predicts that the outcome of the two-sender case will be the same as that of the single-sender case. The comparison against this model lets us quantify the advantage of explicitly modeling interference after having seeded the model with measurements. It is worth noting that even this naive model is likely to outperform many existing analytic models, as it incorporates measured information about signal propagation.

The second model is based on *history*. It predicts the current round using the corresponding, direct measurements made in the last round. This model cannot be used to predict network configurations which have not been previously recorded; instead, it quantifies the temporal variability inherent in the wireless network. All models other than those that predict changes in the RF environment itself are likely to be no more accurate than this model. Comparison here lets us isolate the impact of modeling inaccuracies from the impact of temporal variability.

To present our results, we use the root mean square error (RMSE) as a measure of the accuracy of predictions. This is a standard metric in model fitting that conveys how far off a given prediction is likely to be. We give RMSE values as percentages of the maximum possible value for that prediction. For probability predictions, the maximum is 1. For throughput predictions, the maximum is the bitrate: 6 Mbps for 802.11a and 1 Mbps for 802.11b. For brevity, we adopt the convention of listing 802.11a and 802.11b accuracy numbers side by side (e.g., x_a / x_b %).

5.1 Overall Accuracy

We evaluate the overall accuracy of our models by comparing the two-sender throughput and delivery probability predictions with measured values.

Throughput Prediction Figure 8 plots the CDF of the error in predicting throughput for each two-sender and receiver trial. The error is measured as the difference between predicted and actual throughputs. For an ideal model, the graph would show a vertical transition from 0% to 100% at the x -value of 0. We see that most of our model predictions are accurate, though a small portion are too low (at the left end) or too high. The naive model does worse, and in particular it often overestimates throughput by a large margin by not accounting for interference from other senders.

The RMSEs for the throughput predictions of our model are 11 / 9% for 802.11a/b. The RMSEs of the naive model are substantially higher, at 24 / 31% for 802.11 a/b; 802.11b fares worse because there is more interference due to its longer range. Surpris-




Figure 8: The CDF of error in predicting throughput.




Figure 9: The CDF of error in predicting throughput for the naive model, broken down by the sender range. There are no middle range senders for 802.11a in our testbed.

ingly, our accuracy is similar to that of the history based model, whose RMSEs for 802.11 a/b are 11 / 7%. This suggests that the accuracy of our model might be limited more by temporal variability inherent in a wireless network, rather than modeling inaccuracies. Nonetheless, we believe this accuracy level to be sufficient for many scenarios of interest, such as conflict graph prediction (Section 5.4).

To further understand the performance of our model, we looked for patterns in the cases where our predictions were poor. We found that the most challenging case was that of receivers listening to senders that can only partially hear each other. In these cases, the delivery and deferral probabilities are both highly variable, leading to throughput values that are very sensitive to small changes. However, these same factors affect the variability between successive rounds as well, such that any model will be partly impaired in this region unless it predicts the variability itself.

We further investigate where our predictions of interference are most beneficial. Figure 9 shows the results of the naive model throughput predictions, broken down by the range between the senders (as determined using delivery probabilities from the RF profile). We can see that the naive model does poorly at predicting middle range senders, due to the high variability mentioned above and the high potential for interference at a receiver. It does worst for close sender throughputs, as it does not account for deferrals. The RMSEs for 802.11 a/b are 37 / 32% for close senders and 19 / 25% for far senders. For 802.11b middle range senders, the RMSE is




Figure 10: The CDF of error in predicting delivery probability.




Figure 11: The CDF of error in predicting delivery probability for the naive model, broken down by the sender range. There are no middle range senders for 802.11a in our testbed.

36%; because of 802.11a’s shorter effective range, we had no middle range senders in 802.11a. By comparing against our model results in Figure 8, we can see the substantial gains possible when accounting for interference and deferrals in predictions.

Delivery Probability Prediction Figure 10 shows a similar CDF of prediction error when delivery probabilities are compared. We see results for our model that are very similar to those of throughput, even though the two quantities may be quite different, e.g., two senders may halve their throughput by competing with little change in delivery probability if they sense each other clearly. We see that our model is again comparable to the history based model; the RMSEs of our model are 11 / 10%, while those of the history based model are 12 / 10%.

The naive model again overpredicts scenarios with interference, but its delivery probability predictions are not as poor as for throughput, because deferrals between senders are irrelevant here. Overall, its RMSEs are 19 / 29% for 802.11 a/b. Figure 11 confirms that the naive model does better at delivery probability than throughput for close senders, though it continues to neglect the impact of interference from senders that do not hear each other well. The RMSEs are 10 / 8% for close senders and 21 / 27% for far senders. For 802.11b middle range senders, the RMSE is 38%.

5.2 PHY Receiver Component

To evaluate the individual components of our model, we now look at them separately. The receiver component of our model predicts the fraction of packets a node will receive from a sender, in the




Figure 12: The CDF of error in predicting delivery probability for the receiver component.

face of simultaneous transmissions from another sender. To study this component in isolation, we first identify those triplets (consisting of a pair of senders and a receiver) that will actively test this component. In such triplets, the two senders are sufficiently far from each other that they transmit simultaneously, rather than deferring to each other, and the receiver can hear packets from both senders, at least partially. This second condition filters out cases where one of the senders is so far as to not make a difference. For 802.11 a/b, there are 12 / 30 such triplets in our dataset, implying that we test our model on 24 / 60 predictions. To ensure that our predictions here are not influenced by errors in the deferral component of our model, we use the measured, not predicted, number of packets sent by each of the senders.

Figure 12 shows the CDF of prediction error for delivery probability for all models. We see a stark difference from the earlier CDFs, since we are focusing on only those cases where interference is present. The RMSEs of our model remain low at 12 / 18% for 802.11 a/b, comparable to the history model RMSEs of 18 / 11%. The RMSEs for the naive model are considerably poorer at 67 / 60%. These results imply that our model is quite effective at both predicting situations when effects such as capture [13] arise and quantifying their impact.

5.3 PHY Deferral Component

We now evaluate our deferral component in isolation. The key metric here is the accuracy of the predicted probability of one sender deferring to another. We compare the predicted deferral probability to the measured deferral probability. The latter is computed using the packet transmission counts from each sender taken from the two-sender (validation) runs. Assume that the two senders, s and t , send P_s and P_t packets when sending simultaneously, and that there are a total of P opportunities to send packets. The difference for each node, i.e., $P - P_s$, is the number of packets that were deferred. Even if the senders hear each other perfectly, we will observe nodes sending slightly more than half the time because of collisions due to ties in the backoff race. We know from our analysis in Section 4.5 that the number of expected collisions, C , is $\frac{2}{16}P$ for 802.11a and $\frac{2}{32}P$ for 802.11b. Nodes are equally likely to win the countdown race and use the remaining opportunities. Thus, with perfect deferrals they defer $\frac{P-C}{2}$ packets to each other and send a minimum of $\frac{P-C}{2} + C$ packets. When the senders cannot hear each other, they never defer and send a maximum of P packets. The observed deferral probability is then the number of deferrals as a fraction of the maximum possible. That is:




Figure 13: The CDF of error in predicting deferral probability.

$$Prob[s \text{ defers}] = \frac{2(P - P_s)}{P - C}$$

Figure 13 shows the CDF of deferral probability prediction error. The RMSEs for our model in 802.11 a/b are 23 / 19%. The RMSEs for the naive model are 52 / 54%, while those for the history based model are 0 / 8%. We see that the errors for our predictions are higher than that of the history model but notably lower than those of the naive model, which predicts senders will never defer. (We found another naive model in which senders always defer to be even less accurate; senders clearly do defer to each other but only in particular cases.)

We do less well at predicting deferrals than predicting packet delivery. One possible cause is that the RSS versus delivery probability data used for interpolation is sparse and variable. To explore how much impact this has, we experimented by heuristically smoothing the RSS data in the RF profile. We did this by removing outlier points, then making the function monotonic increasing and ensuring that it reaches 100% delivery at a reasonable RSS. The result was to lower the RMSEs for our deferral probability predictions to 3 / 17% for 802.11 a/b, and slightly lower the RMSEs for our overall throughput predictions to 8 / 8%. This is promising, as it suggests that better modeling of the RSS relationship can further improve our results.

As a second possible cause, we observe slightly lower RSS measurements when a node is attempting to send as it listens, compared to when it only listens. We have not yet included this effect in our models.

5.4 Application: Conflict Graphs

We see our model as a foundation on which to predict the performance of a wide range of higher-layer behaviors. To illustrate its versatility, we show a simple application in which we predict the *conflict graph* [10] of the network. This graph captures the level of interference between different unicast conversations. It is of broad interest because knowledge of the conflicts can be used to improve performance, e.g., by using non-interfering links in parallel, rather than making overly conservative scheduling decisions in the manner of 802.11 CSMA/CS and RTS/CTS.

Padhye *et al.* quantify conflicts using link interference ratios (LIR) between two unicast conversations as a metric. LIR is the ratio of the sum of throughputs when both conversations are active to the sum when they are active individually. It thus ranges from 0 to 1, with 1 representing no interference. Because LIR is




Figure 14: The CDF of error in predicting BIR.

intractable in networks with many conversations due to the large number of combinations, Padhye *et al.* show that it can be estimated using a broadcast interference ratio (BIR) [17] that requires fewer trials. BIR is the ratio of the sum of packet delivery probabilities at respective receivers when both senders broadcast to when they broadcast individually.

We use our models to *predict* BIR as a proxy for the level of unicast interference. We do this simply by making delivery probability predictions for two-sender cases and combining them with single-sender data according to the BIR metric. Figure 14 shows that our model does well at this task. It plots the CDF of the difference between predicted and measured BIR, where measured BIR is obtained by using our two-sender experiments. For 802.11 a/b, RMSEs of our predictions are 14 / 10%. Both the CDFs and RMSEs indicate that we are quite close to the actual BIR on average.

Observe that the use of our model provides a substantial advantage: we require only N single-sender trials in an N node network to estimate BIR; measuring it directly requires N^2 trials, one for each pair of senders. In contrast, direct measurement of the conflict graph as LIR for pairs of conversations would have required roughly N^4 trials! Our method also generalizes easily to more conversations without additional measurements.

6. RELATED WORK

Much work on modeling wireless networks has relied on simple assumptions about signal propagation and interference [9, 19]. While such models provide important insight into the asymptotic behavior of large networks, the use of such models in real networks has been shown to be erroneous [14]. Our approach is fundamentally different: it is grounded in measurements of real networks to avoid simplistic assumptions about signal propagation.

There is relatively little work combining measurements with models. Padhye *et al.* [17] estimate the extent of interference between two unicast conversations with measurements of broadcast interference. Our model tackles a more general, and challenging, case for prediction. Woo *et al.* [20] uses measurements to construct link quality estimators in sensor networks. Similar to us, they find that measurements add significantly to realism, though they do not explicitly model effects such as interference. Cerpa *et al.* [7] present a tool for assessing connectivity in lossy environments. It is based on packet delivery statistics, rather than underlying causal models, and does not attempt to model packet delivery in new configurations. Also, Judd and Steenkiste emulate signal propagation in hardware to provide a tool for experimenting with wireless networks [12]. They notably improve realism and repeatability, but unlike predictive models they must evaluate each configuration of interest experimentally.

Other studies have characterized aspects of wireless behavior in practice for particular settings, e.g., the Roofnet project has investigated characteristics of packet loss, connectivity and throughput on a city-scale wireless network [3, 5]; work on Divert measures the burstiness and spatial patterns of losses on an indoor [16]; and work on SCALE includes an assessment of wireless conditions in sensor networks [7]. Our characterization work is largely complementary but consistent with the findings in these other studies.

Finally, some recent efforts have been made to use empirical observations to improve wireless protocols. Divert [16] attempts to reduce packet loss rates in WLAN systems by rapidly switching between APs to tolerate bursty losses. ExOR [6] leverages spatial loss independence to reduce packet transmissions in multi-hop networks by using opportunistic packet reception. These efforts indicate that there is much room to improve wireless protocols by adapting them to realistic conditions. Our work provides one tool to do so.

7. CONCLUSIONS

We present practical, measurement-based models for the physical layer behaviors of static wireless networks, including packet reception and carrier sense with interference. To improve accuracy over abstract RF propagation models, we seed our models with measurements of RSS values and delivery probability, both of which are easily obtainable. Our RF profile of a given network topology includes data for N nodes from N sequential transmission rounds, providing roughly N^2 data points. Our models can then predict packet delivery and throughput at receivers in the same network topology when multiple senders compete to transmit packets. This covers a much larger range of transmission configurations than were measured; there are $O(N^3)$ configurations when at most two nodes send at once. We evaluate our model for the base case of two senders that continually transmit broadcast packets, and find that it can accurately predict when there will be significant interference effects. Across many predictions, we obtain an RMS error for 802.11a and 802.11b of a half and a third, respectively, of a measurement-based model that ignores interference.

There is clearly much more work to be done to achieve our long-term goal of predicting the performance of a given network when it is used in a different configuration, for instance, with different routing, MAC and PHY design choices. Thus far we have evaluated our model for two simultaneous broadcast senders, the base case in which interference effects are seen, and not tested its accuracy across the much larger space of possible workloads. In future work we hope to extend our methodology in these directions; this would greatly increase the design space in which we can apply our models.

On the other hand, we view the importance of our work as being able to predict interference effects with any reasonable accuracy at all. Wireless networks are notorious for their variability and complex interactions. The tone of much work is that there is little rhyme or reason to the performance that will be achieved in practice for a given setting. We have shown that, to the contrary, measurements do have predictive power when they are carefully interpreted. We have shown how to measure an RF profile of a network to factor the complexity of RF propagation out of the modeling domain, so that RSS values averaged over tens of seconds can be useful predictors of performance for the majority of the time and well into the future.

We hope our work will provide a starting point for new methodologies and protocols as measurement-based prediction is improved in scope and accuracy. Measured RF profiles may be combined with simulators via our models to allow the realistic exploration of networks. New protocols can then account for interference by

using online models, rather than “playing safe” with overly pessimistic assumptions, as does 802.11 [8, 18].

8. ACKNOWLEDGEMENTS

We are grateful to Sergei Kaganovsky for helping with the attenuator experiments and Ed Lazowska for constructive discussions. We also thank the anonymous reviewers for feedback on earlier versions of this text. This work was supported in part by the NSF (Grants CNS-0133495 and CNS-0338837).

9. REFERENCES

- [1] R. L. Abrahams. Intersil - measurement of WLAN receiver sensitivity. <http://www.demartech.com/techsupport/rw-wireless-cards-support/wlan-receiver-test.pdf>, Feb. 2000.
- [2] D. Aguayo, J. Bicket, S. Biswas, D. De Couto, and R. Morris. MIT Roofnet implementation. <http://www.pdos.lcs.mit.edu/roofnet/design/>, Aug. 2003.
- [3] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b mesh network. In *SIGCOMM*, Aug. 2004.
- [4] J. Bicket. Madwifi stripped driver. <http://pdos.csail.mit.edu/~jbicket/madwifi.stripped/>.
- [5] J. Bicket, D. Aguayo, S. Biswas, and R. Morris. Architecture and evaluation of an unplanned 802.11b mesh network. In *MobiCom*, Aug. 2005.
- [6] S. Biswas and R. Morris. Opportunistic routing in multi-hop wireless networks. In *HotNets-II*, Nov. 2003.
- [7] A. Cerpa, NaimBusek, and D. Estrin. Scale: A tool for simple connectivity assessment in lossy environments. Technical Report 21, CENS, Sept. 2003.
- [8] D. De Couto, D. Aguayo, J. Bicket, and R. Morris. A high-throughput path metric for multi-hop wireless routing. In *MobiCom*, Sept. 2003.
- [9] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2), Mar. 2000.
- [10] K. Jain, J. Padhye, V. N. Padmanabhan, and L. Qiu. Impact of interference on multi-hop wireless network performance. In *MobiCom*, Sept. 2003.
- [11] K. Jamieson, B. Hull, A. Miu, and H. Balakrishnan. Understanding the real-world performance of carrier sense. In *E-WIND workshop*, Aug. 2005.
- [12] G. Judd and P. Steenkiste. Using emulation to understand and improve wireless networks and applications. In *NSDI*, May 2005.
- [13] A. Kochut, A. Vasani, A. Shankar, and A. Agrawala. Sniffing out the correct physical layer capture model in 802.11b. In *ICNP*, Nov. 2004.
- [14] D. Kotz, C. Newport, R. S. Gray, J. Liu, Y. Yuan, and C. Elliott. Experimental evaluation of wireless simulation assumptions. In *MSWiM*, Oct. 2004.
- [15] Mesh networks. <http://research.microsoft.com/sn/mesh/>.
- [16] A. K. Miu, G. Tan, H. Balakrishnan, and J. Apostolopoulos. Divert: Fine-grained Path Selection for Wireless LANs. In *MobiSys*, June 2004.
- [17] J. Padhye, S. Agarwal, V. N. Padmanabhan, L. Qiu, A. Rao, and B. Zill. Estimation of link-interference in static multi-hop wireless networks. In *IMC*, Oct. 2005.
- [18] A. Raniwala and T. Chiueh. Architecture and algorithms for an IEEE 802.11-based multi-channel wireless mesh network. In *IEEE INFOCOM*, Mar. 2005.
- [19] T. J. Shepard. A channel access scheme for large dense packet radio networks. In *SIGCOMM*, Aug. 1996.
- [20] A. Woo, T. Tong, and D. Culler. Taming the underlying challenges of reliable multihop routing in sensor networks. In *Sensys*, Nov. 2003.