

SIR Based Interference Modeling For Wireless Mesh Networks: A Detailed Measurement Study

Vishal Sevani

Dept. of Computer Science & Engg.
Indian Institute of Technology-Bombay
Email: vsevani@cse.iitb.ac.in

Bhaskaran Raman

Dept. of Computer Science & Engg.
Indian Institute of Technology-Bombay
Email: br@cse.iitb.ac.in

Abstract—Spatial reuse is a central aspect in the efficiency of wireless mesh networks. Accurate inter-link interference measurement and estimation is necessary for such spatial reuse. In this work, we explore in depth Signal to Interference Ratio (SIR) based interference modeling. We take a measurement centric approach, characterizing the SIR versus PDR (Packet Delivery Ratio) relationship in outdoor mesh network settings. Our significant findings are the following. (1) In outdoor environment there is a range of SIR values (*intermediate SIR region*), wherein it is difficult to predict the PDR accurately. (2) Fortunately, the width of this *intermediate SIR region* is small: about 4-5dB for most data rates. (3) The SIR vs PDR relationship depends significantly on the modulations used by transmitter as well as the interferer: we characterize this dependence.

The use of an outdoor 802.11g-based testbed is a significant aspect of our measurements. The above findings have important implications for the design of interference measurement and prediction schemes. We quantify the accuracy of our SIR-based technique by applying an offline prediction model in our outdoor mesh testbed and observe that of the 44 pairs of links the SIR technique predicts performance with less than 10% error for over 85% of the links for which the SIR values lie outside the *intermediate SIR region*. An evaluation of existing interference modeling techniques reveals that these make certain incorrect assumptions which make them perform significantly worse than our proposed technique.

I. INTRODUCTION

Outdoor multi-hop wireless mesh networks are cost-effective solutions to provide last mile connectivity [4]. Since multi-hop wireless networks using the conventional 802.11 based CSMA/CA protocol are known to perform poorly in terms of throughput, QoS, fairness, etc, researchers have considered the use of TDMA based mesh networks [5], [13], [17]. To further improve the performance of such networks, there is a need to incorporate techniques such as *spatial reuse*. Spatial reuse requires *estimating the interference* in the network, and predicting the performance when two links in the network are active simultaneously. Note that this interference is *internal* to the network¹. The goal of this work is to estimate such internal inter-link interference for static outdoor TDMA based multi-hop networks.

¹Internal interference is interference caused by sources within the network, while external interference is uncontrolled interference from sources outside the network.

Interference estimation is a particularly challenging problem, primarily due to the inherent unpredictability of the wireless environment. This work uses a *Signal to Interference Ratio* (SIR) based technique for modeling and estimating hidden terminal interference, particularly for outdoor TDMA mesh networks. We term this the *SIR technique*; it correlates SIR with PDR, and requires only RSS (Received Signal Strength) to be measured from individual senders. This can be accomplished using $O(n)$ measurements for an n -node network.

In the recent past, researchers have proposed various measurement based (e.g. [21], [3], [2]) as well as model-based approaches (e.g. [9], [19]) to address this issue. In light of these, our contributions in this work are four-fold.

- 1) *Outdoor testbed-based measurements*: We are not aware of any prior work which has used outdoor measurements for validation of interference estimation. In fact, most are intended for indoor infrastructure networks. Since the intended deployment for a wireless mesh network is almost always *outdoors*, it is important to validate the mechanism in outdoor settings. Radio propagation in outdoor settings can be very different from indoor settings: due to the larger distances (hundreds of metres) in outdoor settings, path loss and delay spread due to multi-path fading could be much higher [1]. Indeed we find in our experiments that the effect of frequency-selective fading in OFDM is very different outdoors, compared to prior indoor results in [8] (Sec. VI).
- 2) *Measurement using 802.11g*: Most prior work has used either 802.11a or 802.11b measurements, whereas we have considered the more popular 802.11g. Now, 802.11a has significantly shorter radio range than 802.11g, while 802.11b uses different modulation schemes from 802.11g; so measurements using 802.11g are significant.
- 3) *Characterization of interference dependence on interferer's modulation*: Most prior work has assumed that the level of interference depends only on the interference power level (e.g. [12], [19], [9], [3]). While the dependence of interference on the interferer's modulation was explained analytically in [10], we believe we are the first to show it in practice, as well as characterize it quantitatively.

4) A *simplified yet effective SIR-based model*: Prior measurement based approaches typically involve $O(n^2)$ measurement overhead. And most prior model based approaches (e.g. [9], [19]) involve significant complexity, most of which comes from modeling carrier sensing in the 802.11 CSMA/CA protocol. Our focus is on TDMA networks, where links are scheduled, and carrier sense is turned off [5], [18]. This allows a simple yet effective model.

As part of our measurement work we characterize the SIR vs PDR relation. Although several prior works have studied in detail the SNR vs PDR correlation for outdoor networks [1], [6], [7], there are practical differences between measured SIR and measured SNR and the results for SNR vs PDR correlation cannot be applied directly (Sec. III). We also quantify the prediction accuracy of our SIR technique by integrating it with a multi-hop TDMA MAC, in our outdoor testbed.

The significant findings and implications of our measurements are as follows. (1) In outdoor settings, some links may have an SIR value within what we term the *intermediate SIR region*, wherein it is difficult to predict the PDR accurately for a given SIR. This implies that we cannot rely upon SIR to choose between two such links. (2) Fortunately, the width of this *intermediate SIR region* is small: about 4-5dB for most data rates. This implies that scenarios for which SIR-based prediction is error-prone, will be rare. (3) The accuracy of our SIR technique is good overall (less than 10% error for over 85% of the links outside the *intermediate SIR region*), despite minimal measurement overhead. (4) The SIR vs PDR correlation depends significantly on the modulations used by transmitter as well as the interferer. This means that such dependence must be taken into account while deciding which link pairs can or cannot operate in parallel in the network.

The rest of the paper is organized as follows. The next section presents related work. Sec. III describes the SIR technique. Subsequently, Sec. IV presents detailed experiments to study the SIR vs PDR relationship in outdoor as well as controlled settings. Sec. V quantifies the prediction accuracy of the SIR technique, while Sec. VI presents a comparison of our results with prior studies. Sec. VII concludes the paper.

II. RELATED WORK

Table I is a succinct summary of comparison with prior work. As can be seen immediately in the last column, prior work on interference estimation has not considered *outdoor* settings, or used *802.11g* for measurements. As indicated earlier, these are significant aspects of our work. The rest of this section presents further detailed comparison.

[15], [14] represent some of the early works in building interference maps. They make use of throughput based experiments, wherein pairs of nodes transmit simultaneously while other links infer inter-link interference from the measured throughput. For an n node network, the measurement overhead is $O(n^2)$, which is impractical for large networks, especially if the interference pattern changes frequently.

Technique	Overhead	Focus on hidden terminals	Predictive capability	Testbed, 802.11 variant used
Throughput based [15], [14], [3]	$O(n^2)$	No	None	Indoor, 802.11b
RSS modeling based [19], [11], [9], [16]	$O(n)$	No	Yes	Indoor, 802.11a/b
Micro-probing [2]	$O(n^2)$	No	None	Indoor, 802.11a
Passive scanning [23], [21]	None	Yes	None	Indoor, 802.11a
SIR based [20], [22]	$O(n)$	Yes	Yes	Indoor, 802.15.4
TDMA based [13]	$O(n)$	Yes	Yes	Indoor, 802.11a-6M only
SIR based (our work)	$O(n)$	Yes	Yes	Outdoor, 802.11b/g

TABLE I: Summary of comparison with prior work

A further disadvantage of throughput based measurements is that they cannot predict interference relations in alternative scenarios. For instance, they require the throughput measurements to be carried out again if the transmit power, data rate or channel of any of the nodes were to be changed, thereby incurring huge overhead.

A model based, as opposed to measurement based, approach overcomes the above limitations. The work in [19], [11] models the PHY receiver as well as the CSMA/CA MAC. The models are fed by broadcast measurements from each node in the network, constituting $O(n)$ overhead. Subsequent work [9], [16], has extended such modeling to the case of multiple interferers, as well as non-backlogged interferers.

The focus of most modeling work is on CSMA/CA settings. Much of the MAC modeling as well as PHY layer carrier sense modeling is not applicable to TDMA settings, since in TDMA networks, carrier sensing is disabled and all interference shows as hidden terminal interference. Furthermore, these models make two underlying assumptions: (a) interference can be estimated by measuring the variation in RSS [19], [16], and (b) the SIR vs PDR curve can be estimated using the SNR vs PDR curve [19], [9]. As we will evaluate in Section VI, these assumptions do not seem to hold in practice, at least in our measurements.

The work in CMAP [23] takes advantage of parallel transmissions by detecting and enabling exposed terminal situations. Hence, like in our work, the conflict map built by CMAP also consists entirely of hidden terminal cases. The advantage of CMAP is that it is a passive technique, without any extra overhead. Micro-probing is another technique developed to drastically reduce the overhead of building interference maps [2]. It works by having Access Points synchronize with each other, and by sending probes in overlapping time-windows. [3] builds upon micro-probing and seeks to optimize the number of measurements when parameters such as the transmit power or data rate have to be varied. Likewise, [21] also employs similar technique as [2], but instead makes use of real traffic for deriving interference relations, so that there is no overhead of measurements.

However, like throughput-based interference map building, [23], [2], [21] lack the predictive ability of model-based techniques. Moreover, [2], [21] have been designed and evaluated only for infrastructure WiFi deployments, where nodes can rely upon a backbone ethernet either for fine-grained

clock synchronization [2] or for collecting the information across the different access points [21].

SIR based interference measurement for TDMA based networks has been considered by [13], [12]. The focus of [12] is on the use of the directional antennas, and it assumes the efficacy of the SIR based interference estimation. So our study of the various aspects of SIR vs PDR correlation, and the validation of the SIR technique are complementary to [12]. Though being closest to our work in terms of the goal, [13] has also been evaluated only on an 802.11a indoor testbed, using only the 6M transmit rate. The work in [20], [22] also propose the use of SIR based technique for enabling concurrent transmissions; while these have carried out measurements for 802.15.4, our work focuses on 802.11b/g, which uses different PHY layer modulations as compared to 802.15.4.

III. SIR-BASED INTERFERENCE MEASUREMENT

In this section, we briefly explain the SIR technique to measure internal inter-link interference in a mesh network².

A. The SIR technique

At the symbol level, it is fundamentally known that signal to interference plus noise ratio (SINR) has a strong correlation with bit error rate (BER). Intuitively, we *speculate* that, given the averaged packet-level SINR for the packets from a given sender and interferer at a particular receiver, it is possible to predict PDR for the sender when both sender and interferer are active simultaneously. Packet-level SINR estimates can be easily obtained by measuring the packet-level *Received Signal Strength* (RSS) from sender and interferer, individually, at the receiver. Moreover for an n -node network, this can be easily accomplished in $O(n)$ measurements, if broadcast measurements are used from each node in turn. Since for commodity hardware noise is significantly less than the signal strength, it can be neglected and hence we term the technique as SIR technique.




Fig. 1: SIR technique: key components

The key components of our SIR technique, as shown in Fig. 1, are the *RSS profile* i.e. packet-level averaged RSS from each senders at each of the receivers and the *prediction model* i.e. the SIR vs PDR curve to be used for prediction of performance. Given these, to predict the PDR of link A when link B is also operating simultaneously, the RSS profile is used to compute the SIR of link A . And the SIR vs PDR curve is used with the computed SIR to predict the PDR.

The methodology to generate the RSS profile depends on the TDMA MAC protocol design, and we outline a methodology

²The SIR technique was initially outlined in our earlier work [18], but not studied or evaluated in detail.

using broadcast packets that takes $O(n)$ measurements in Sec. V. For deriving the prediction model to be used, it is essential to study how the radio signal propagation and channel characteristics in outdoor environment affect the SIR vs PDR plot. This is the focus of the next section (Sec. IV).

B. Is SIR same as SNR?

In the SIR technique (Fig. 1), can we use the SNR vs PDR curve in place of the SIR vs PDR curve? The SNR vs PDR relation for outdoor networks has been studied in detail [1], [6], [7]. But we cannot use these results directly, since there are practical differences between measured SIR and measured SNR.

Essentially SIR is estimated by measuring the RSS from sender and interferer individually. Now the RSS as reported by commodity hardware, for a given packet, is that measured during preamble of the receiving frame. So the implicit assumption is that SIR for all the bits in a given packet is same. But this is not true practically, since when sender and interferer are active simultaneously, the packet from interferer may not overlap completely with that from sender. So in that case, some of the bits in a given packet for the sender will be affected by interference, while others will be not i.e. SIR for different bits, in a given packet, will likely be different.

Further even though we have considered static networks, the environment is not entirely static due to movement of foliage, vehicles on the road, etc. So the multi-path fading experienced by different bits in a given packet may be different, again causing the SIR to vary across a given packet. For the same reason the SNR for different bits in a given packet will also be different. But since in the case of SIR measurement as both signal and interference are varying, the variability is much more than for the case of SNR measurement.

Given these differences, it is essential to answer the following questions to study the applicability of SIR technique for interference modeling

- 1) Will the symbol level SIR vs BER correlation necessarily mean packet-level SIR vs PDR correlation?
- 2) Is the packet-level SIR vs PDR relation uniform across different links? How does it compare with a controlled setting?
- 3) What effect does (frequency selective) fading have on SIR vs PDR relation?
- 4) Finally, what good is the measured interference relation, for spatial reuse in a TDMA MAC?

The rest of the paper sheds light on the above practical aspects. Henceforth we use the term *SIR curve* interchangeably with *SIR vs PDR plot*.

IV. SIR CURVE CHARACTERISTICS IN OUTDOOR SETTINGS

We now experimentally characterize various aspects of the SIR vs PDR relationship in an outdoor environment. We first compare the SIR curve in outdoor settings with that from controlled settings (Sec. IV-B). We then study variation of the SIR curve with time (Sec. IV-C), as well as its dependence

on the interferer’s modulation (Sec. IV-D). This section ends with a discussion of the implications of these measurements (Sec. IV-E).

A. Experimental methodology

We have carried out experiments in our campus wide outdoor testbed (Fig. 2) consisting of 6 nodes, which gives us a total of 30 different links to carry out the measurements. The distance between the nodes varies from few hundred meters to over a kilometer. Many links in our network have LoS (line of sight); some have a few obstacles in the form of large trees and buildings in-between and around the nodes, which likely cause some amount of multi-path fading.

To study how the fading effects associated with the outdoor environment affect the SIR curve, we first carry out experiments on 10 *good quality* links of the total of 30 links in our testbed. These links are chosen such that the PDR for these (without internal interference), is above 95%. This ensures that the effect of external interference on the observed results is negligible. (Our final evaluation in Sec. V however considers all the links in the testbed).

For controlled measurements, we connect the nodes using RF cables (within an indoor lab) to exclude fading effects like multi-path reflection that are associated with outdoor links. To make the signal strength from the nodes comparable to that for outdoor links, we attenuate the signals from sender and interferer which is then fed as input to the receiver via a directional coupler. Fig. 2 shows the controlled experimental set-up and the outdoor testbed.




Fig. 2: Experimental set-up: (a) controlled & (b) outdoor

The hardware we have used consists of Mikrotik RB43-3AH single board computers. We used atheros chipset based Mikrotik R52-350 wireless cards to carry out detailed experiments. We have also verified the results for atheros chipset based Ubiquiti XR2 & SR2 cards. We have used Madwifi driver code (v0.9.4, release 3314) that we have modified to implement our TDMA MAC protocol, which provides synchronization accuracy of 10-15 μ s at single hop [5].

The experimental methodology we make use of, is

- Each experiment involves a chosen sender, receiver, and interferer.
- The experiment progresses in *rounds*. In the first round the sender alone transmits 200 UDP packets of 1400 bytes

each, while in the second round the interferer transmits alone. In the third round both sender and interferer transmit simultaneously. We term the set of three such rounds, as a *bin*.

- The receiver measures the RSS from the sender and interferer, in the first and second rounds of a bin respectively. The SIR³ for the bin is then calculated as,

$$SIR(dB) = Avg. RSS(dBm) \text{ of sender in first round} - Avg RSS(dBm) \text{ of interferer in second round}$$

The PDR from the sender in the third round is then used to get a sample point for the SIR curve.

- This set of rounds (bin) is repeated by varying the transmit power of sender and interferer to obtain the range of SIR values.




Fig. 3: Experiment methodology

The experiment methodology described above, is illustrated in Fig. 3. We vary the transmit power after 20 bins each. We measure PDR as average across all the bins for which the SIR is the same. Since the granularity at which the commodity hardware reports the signal strength is 1dBm, we round off the SIR value to the nearest integer for reporting the results.

B. SIR curve: outdoor vs controlled environments

To study the SIR vs PDR correlation for both 802.11b & 802.11g, we have carried out experiments for two 802.11b data rates of 5.5 and 11Mbps, and three 802.11g data rates of 6, 9 and 12Mbps. We describe the results below.

We first carried out experiments to measure the SIR curve in controlled set-up, for each of the five data rates. To obtain the SIR curve for a given data rate, we repeated the experiment described above in Section IV-A, 10 times. We then used the SIR curve obtained by taking average across these 10 experiments as a basis for studying how the SIR curve varies in the outdoor testbed.




Fig. 4: SIR curves in controlled set-up

Accordingly Fig. 4 shows these SIR curves, measured in controlled set-up, for each of the five data rates. In the SIR curve, we define the *transition region* as the range of SIR

³Ratio of actual power values in mW = difference in dBm.




Fig. 5: Average, std. deviation, and maximum of PDR prediction error across 10 outdoor links

values with PDR between 10 and 90%. And we term the required SIR for 90% PDR as the *SIR threshold* (this is the outer end-point of the transition region). As can be seen from Fig. 4, SIR curves are similar for lower data rates and as the data rate increases the *transition region* shifts towards the right with *SIR threshold* also increasing. For clarity, we have not plotted the standard deviation in the graph, but we note that the standard deviation across these 10 experiments was small: under 5% for most cases, and about 10% for some of the readings.

We then carried out experiments to measure SIR curves for each of the good quality 10 outdoor links for all the five data rates. For comparing the SIR curves measured across the outdoor links with that obtained in controlled set-up we compute the *PDR prediction error* for a given SIR value as, absolute value of difference in the *measured PDR* value for outdoor link and the *expected PDR*. The expected PDR for a given SIR value for a given data rate, is that obtained via the SIR curve (Fig. 4) measured in the controlled experiment described above.

Fig. 5 shows the maximum, average and standard deviation of the PDR prediction error, for each of the five data rates across the 10 outdoor links. A small PDR prediction error implies a good match between the SIR curve of the outdoor link with that of the controlled set-up. As can be seen from Fig. 5, SIR curves are consistent across outdoor and controlled set-up, except for small range of SIR values for which PDR prediction error is high. We hence define the *intermediate SIR region* as follows. For a particular link, we term the range of SIR values for which the PDR prediction error is more than 10% as *intermediate SIR region*. If we consider the average PDR prediction error across outdoor links, then the *intermediate SIR region* is about 4-5dB wide for the data rates of 5.5, 11 and 12Mbps, while it is only about 3dB for data rates of 6 and 9Mbps as can be seen from Fig. 5.

Within the *intermediate SIR region*, we see that the average PDR prediction error can be 30-40%. And the maximum PDR prediction error can be quite high in the *intermediate SIR region*: 50% or higher.

Explaining the results: As is intuitive, the *intermediate SIR region* typically corresponds to the transition region of the SIR curve measured in the controlled set-up, as can be seen from Table II, which shows the range of the transition and *intermediate SIR region* for these five data rates. In fact the width of the *intermediate SIR region* is about 1dB less than the width of the transition region, for the data rates of 5.5, 9 and 12Mbps, while it is about 3dB less for 6M. Thus

this shows that the region of unpredictability in the outdoor environment lies within the transition region of the SIR curve in the controlled set-up.

Data Rate (Mbps)	5.5	11	6	9	12
Transition region	2 - 7 dB	4 - 8 dB	1 - 6 dB	3 - 6 dB	4 - 8 dB
Intermediate SIR region	3 - 7 dB	4 - 8 dB	3 - 5 dB	4 - 6 dB	5 - 8 dB

TABLE II: Transition and *intermediate SIR region*

Also, we observed that the SIR value for which PDR prediction error is maximum typically corresponds to the *Signal to Noise Ratio (SNR) threshold*⁴ for that data rate. For e.g. the average PDR prediction error for 5.5 and 6Mbps is maximum for SIR = 4dB, which is the SNR threshold for these data rates, for our wireless cards, as measured in our controlled set-up.

What causes PDR prediction error in the *intermediate SIR region*? Now, there are three sources of inaccuracy as regards RSS measurement: (1) the granularity at which commodity hardware reports RSS is 1dB, (2) the RSS reported for the received packet is that averaged over the preamble of the packet, and (3) multi-path fading. While the first two are due to limitations of commodity hardware, the third is fundamental to wireless channels.

What effect can RSS granularity have on PDR prediction error? Now, a 1dB error in estimating RSS could cause a 2dB error in estimating SIR. As mentioned earlier, the *intermediate SIR region* typically corresponds to the transition region. And as can be seen from Fig. 4, a small error in SIR in this transition region causes a large error in predicted PDR.

Secondly as explained in Section III, signal strength as well interference for different bits in a given packet can be different. Thereby approximating the RSS for each bit with that measured during preamble, can also lead to error in estimating SIR. Thus, given these errors in estimating SIR, it is intuitive that prediction error is high in transition region wherein the sensitivity to SIR estimation is also more.

The third reason for PDR prediction error is multi-path fading. Indeed, in Fig. 5 we see high standard deviation across links, of the PDR prediction error in the *intermediate SIR region*. In comparison, we observed that the standard deviation of PDR across different experiments in the controlled set-up was small. This suggests that fading plays a significant role here. Multi-path fading can be different for the different outdoor links, and this affects the measured RSS and hence the measured SIR. We also observed that the different outdoor

⁴SNR threshold (dB) = Receive sensitivity (dBm) - Noise level (dBm)

links exhibited slightly different SIR curves (not shown due to lack of space), especially in the *intermediate SIR region*.

In light of the above, the immediate question is: what effect does time dependent fading have on the same link? This is what we study next.

C. Time varying behaviour of the SIR curve

To study the time varying behaviour of the SIR curve, we carried out measurements across four outdoor links over five different time instances, for 802.11g rate of 6Mbps, and 802.11b rate of 11Mbps. We repeated these experiments across two different time scales, one wherein the difference between consecutive measurements was half an hour and another wherein the difference between the measurements was 2-3 days.

For comparing the SIR curves across different experiment instances, we take the SIR curve of the first experiment as the basis. We compute the PDR prediction error for a given SIR value, as the absolute value of difference in corresponding PDR of the subsequent experiment instance compared with the first experiment instance. Thus we get error curves for 4 experiment instances, across the two time scales, which are shown in Fig. 6. Fig. 6 shows the PDR prediction error for 6 Mbps for the link KR-SOM across the shorter and longer time scales.




Fig. 6: PDR prediction error, at 6M, across the shorter and longer time scales respectively

As can be seen from Fig. 6, for the experiments carried out across shorter time scale, the PDR prediction error is less than 10% for the most of the SIR values. However for the experiments carried out across longer time scale, the maximum PDR prediction error is much more, as much as 30-40%. We observed similar results across other outdoor links as well.

Thus the results above, and those in the previous subsection (Sec. IV-B) suggest that the SIR vs PDR curve is a characteristic of an outdoor link, especially in the *intermediate SIR region* and is different across different links. And even for a given link, the relationship is variable across time and is consistent only across shorter time scale (a few hours). We discuss the implications of these results later in Sec. IV-E. We next study how the SIR curve varies when the sender and interferer are operating with different modulation schemes.

D. SIR curve: effect of interferer's modulation

Most prior interference estimation work implicitly assumes that it is only the interference power level, and not the interferer's modulation, which determines the PDR [12], [19], [9], [3]. In this section, we present experimental results to quantify the dependence on the interferer's modulation.

Experiments: For this, we first carried out experiments on one of the outdoor links by keeping the sender data rate at 6Mbps and varying the interferer data rates to 1, 6 and 18Mbps. We observed that the SIR curve is distinctly different when the interferer is at 1Mbps than when the interferer is at 6 or 18Mbps as can be seen from Fig. 7. In-fact the SIR threshold in the former case is about 7-8dB more. Also, as can be seen from Fig. 7 that though SIR curves for the case of interferer data rates of 6 and 18Mbps are also not exactly same, but they vary only in the narrow transition region from 2-6dB. We carried out similar experiments across other outdoor links and in the controlled set-up as well, with similar results.




Fig. 7: Variation in SIR curve with interferer data rate in outdoor testbed (sender at 6Mbps)

Understanding the variation in SIR curves: To further study this variation in SIR curve with interferer data rate in detail, we carried out measurements in controlled set-up by varying the interferer data rate from 1 to 18Mbps, for each of the five data rates (5.5, 11, 6, 9 and 12Mbps) for the sender. Table III summarizes the results by listing all the interferer data rates, for a given sender data rate for which the SIR curves are similar. Fig. 8 shows the SIR curves for the case of sender at 6Mbps and five different interferer data rates.

#	Sender rate	Intrfr rate	Remarks
1	g	6M,9M,12M	Similar to one another
2	g	1M,2M	SIR curve is distinctly different than #1 with the SIR threshold being 7-8dB more than #1.
3	g	5.5M,11M	SIR curve is similar to #1 in the transition region till 80-90% PDR and then varies by 10-20% for further 6-7dB compared with #1.
4	b	5.5M,11M	Similar to one another
5	b	1M,2M	SIR curve is distinctly different than #4 with the SIR threshold being 7-8dB more than #4.
6	b	6M,9M,12M	SIR curve is different than #4 with the SIR threshold being 2-3dB less than #4.

TABLE III: Summary of results (g:= A 802.11g rate of 6,9 or 12M; b:= A 802.11b rate of 5.5 or 11M)




Fig. 8: Variation in SIR curve with interferer data rate in controlled setting (sender at 6Mbps)

As can be seen from Table III, for a given 802.11g data rate for the sender the SIR curves are similar for any of the

802.11g data rate for the interferer. As against this, the SIR curve varies when the data rate of the interferer is changed to 802.11b rate. Further, for a given 802.11g data rate for the sender, the SIR curves are also different for different 802.11b rates for the interferer (results #2 & #3 in Table III). Likewise is the case when sender is at a given 802.11b data rate of 5.5 or 11Mbps.

Now, 802.11b data rates of 1 and 2Mbps employ DBPSK/DQPSK + DSSS (DBPSK for 1Mbps, DQPSK for 2Mbps), 5.5 and 11Mbps employ CCK, while 802.11g data rates employ OFDM as modulation schemes⁵. So these results indicate that for a given sender data rate the SIR curve varies when the data rate of the interferer is changed such that it employs different modulation scheme than the sender.

Furthermore, as indicated by results #2 and #5 in Table III, the DSSS modulated signals of 1 and 2Mbps have more potential for causing interference, since SIR threshold increases by as much as 7-8dB in that case. As against this, as indicated by result #6, OFDM modulated signal of 802.11g is less interfering compared with 802.11b signal.

Reason for dependence on interferer's modulation: The Power Spectral Density (PSD) for different modulation schemes is different. Since a given frequency across different modulation schemes have different power levels, a given frequency of the sender's signal is affected differently when the modulation scheme for the interferer is varied, causing the overall BER to vary.




Fig. 9: PSD for DSSS and OFDM modulated signals

For e.g. consider the PSD for DSSS modulated signal of 802.11b and OFDM modulated signal of 802.11g as shown in Fig. 9. For DSSS modulated signal the power level for the frequencies around f_o is more than at other frequencies, while all the OFDM tones have same power level. For the case of sender and interferer both with OFDM modulated signal, all the OFDM tones for the sender get affected equally. However when the interferer's signal is changed to DSSS modulated signal, the OFDM tones for the sender around f_o are affected more causing the overall BER to increase [10]. This explains why the DSSS modulated signal has more potential for causing interference to other signals⁶.

E. Implications of the results: the offline prediction model

The various results above are significant from the point of view of design of an interference measurement technique in a mesh network. In particular, the results suggest that for

⁵DBPSK/DQPSK - Differential Binary/Quadrature Phase Shift Keying, DSSS - Direct-Sequence Spread Spectrum, CCK - Complimentary Code Keying, OFDM - Orthogonal Frequency Division Multiplexing

⁶While [10] predicts dependence on interferer's modulation analytically, we have shown the same experimentally, as well as quantified the same.

the links that have SIR value outside the narrow *intermediate SIR region*, the performance can be predicted with reasonable accuracy using the SIR curve from the controlled set-up. Further, the *intermediate SIR region* is quite narrow, at most 4-5dB wide for the data rates for which we have carried out the experiments (see Table II). This means that we can expect most of links to be outside this region (we validate this in Sec. V). This then suggests that there is little need for any per-link measurement of the SIR curve. The SIR curve from the offline controlled set-up can be used (see Fig. 1) for predicting the performance in the network, requiring no additional overhead of measuring the SIR curve in real time. We term this approach as the *offline model* of interference prediction.

Apart from the fact that links in the *intermediate SIR region* are few, another reason why the *offline model* makes sense is the following. For links in the *intermediate SIR region*, PDR also is intermediate. Now, intermediate PDR links (e.g. PDR is 50% or lower) are not preferable from an application perspective. So, typically it would be preferable to simply avoid operating such links in parallel anyway (unless the MAC protocol incorporates robust error recovery mechanisms to counter the low PDR in the *intermediate SIR region*).

The SIR curve dependence on the interferer's modulation (Sec. IV-D) is significant for predicting interference for the case when the data rate of the interferer is changed. Particularly these results indicate that SIR curves need to be measured for several combinations of sender and interferer data rates. However, note that under the offline model, these measurements need only be done offline, not in real-time.

A note on the testbed size: We note that though our testbed size is relatively small, we have carried out experiments for five different data rates, spread over a duration of 2-3 months and under varying weather conditions of rain, high humidity, etc. The consistency of our overall results across time, data rates and different links, gives us confidence in the overall conclusions and implications above. The effectiveness of the SIR technique is further evaluated in the next section.

V. PERFORMANCE EVALUATION OF THE SIR TECHNIQUE

While the SIR curves were characterized using only a subset of our testbed links, we now use all of the testbed links to evaluate the offline prediction model. We first outline the mechanism to integrate the SIR technique with a TDMA MAC protocol and then evaluate the performance.

A. Integrating the SIR technique with TDMA MAC protocol

As mentioned in Sec. III, the SIR technique requires the SIR curve (i.e. the prediction model) and the RSS profile. We focus on the evaluation of the offline model, where we make use of the SIR curve measured in the controlled set-up.

The mechanism to generate the RSS profile depends on the TDMA MAC protocol. We integrate the SIR technique along with the TDMA MAC protocol implementation outlined in [5], wherein a central root node carries out the scheduling of time slots to other nodes in the network. So, the RSS profile has to be constructed at the root node.

The protocol attains time synchronization and schedule dissemination using periodic broadcast of control packets by each node. As the control packets are broadcast periodically, we make use of these for construction of the RSS profile. Every node records the RSS of the control packets that it hears from its neighbours. Likewise, it maintains average RSS of the last 50 control packets (this is to account for any short term channel variations) that it hears from its neighbours. Each node then propagates this average RSS information of its neighbours in its control packet. This information propagates towards the root in a multi-hop fashion. The root then constructs the RSS profile for the entire topology. So even though we require $O(n)$ measurements to generate the RSS profile, since we make use of control packets to record RSS information the overhead of incorporating the SIR technique in the TDMA MAC protocol is minimal: control packets need to be sent periodically anyway.

B. PDR prediction accuracy of the SIR technique

For evaluating the PDR prediction accuracy, we carried out experiment similar to that in Sec. IV. However, the essential difference of our experiments in this section is that for prediction of PDR we obtain SIR values from the RSS profile generated at the root node. As against this, for our earlier experiments we obtained RSS values from the packets received from sender and interferer. This also enables us to verify the design mechanism to generate the RSS profile outlined above. Also, as we intend to evaluate prediction accuracy we do not vary the transmit power for the nodes to obtain range of SIR values, but predict performance for the SIR value at the fixed transmit power for the nodes.

In each experiment a pair of nodes transmit, and all other nodes in the network measure throughput from each of the two transmitting nodes. Each transmitting node acts as sender and interferer interchangeably. Since there are six nodes in our network, for each experiment (sender-interferer pair) we get 4 pairs of links at 4 different receivers, which gives us 8 SIR values for each experiment. Likewise, we carried out experiments for 11 different (sender-interferer) pairs of nodes, which gives us 44 pairs of links i.e. total of 88 SIR values for a total of 30 links in our testbed.

The duration of each experiment is 90 seconds and we repeated each experiment 10 times. The PDR prediction error is then taken as the average across the 10 runs.

We carried out experiments for 802.11g rate of 6M and 802.11b rate of 5.5M. We classify the links as *intermediate SIR* and *non-intermediate SIR* links depending on whether the SIR value for the links lie in the *intermediate SIR region* or not. If we consider the average PDR prediction error across the outdoor links in our testbed, then as can be seen from Table II, the *intermediate SIR region* extends from 3-7dB for 5.5M and 3-5dB for 6M.

We calculate the CDF of the PDR prediction error for these two types of links for both the data rates: these are shown in Fig. 10. As can be seen for *non-intermediate SIR* links, for 6Mbps over 90% of the links and for 5.5Mbps over 85% of the

links, have PDR prediction error less than 10%. However, for *intermediate SIR* links, the PDR prediction error is noticeably higher. But even then the PDR prediction error is less than 30% for about 50% of the links for both the data rates.




Fig. 10: CDF of PDR prediction error for *intermediate SIR* and *non-intermediate SIR* links respectively

Low fraction of intermediate SIR links: Since the PDR prediction error is high for the *intermediate SIR* links, it is important to know the fraction of such links. We carried out analysis to find out what fraction of links have SIR values in the *intermediate SIR region*. Of the 44 pairs of links (i.e. 88 links) for which we carried out measurements, we observed that the number of *intermediate SIR links* is 10 for 6Mbps, and 14 for 5.5Mbps. Thus the fraction of *intermediate SIR links* is only about 11% for 6Mbps and 16% for 5.5Mbps, which is quite less.

To further verify this we analyzed the wireless traces made available by the MIT Roofnet project [1]. We present the results of wireless trace analysis in Table IV. As can be seen from the table, the percentage of links with SIR values in *intermediate SIR region* is only about 10-12% for the MIT Roofnet testbed, which is similar to the numbers in our testbed.

Data Rate (Mbps)	Total no. of links	No. of intermediate-SIR links	% of intermediate-SIR links
5.5	5636	698	12.4%
11	3884	405	10.4%

TABLE IV: Number of *intermediate SIR* links

These results suggest that the fraction of instances when the *offline model* will have high prediction error will be low.

C. Performance improvement due to spatial reuse

The primary use of interference measurement is for enabling spatial reuse among non-interfering links. How does our SIR technique fare in this regard? For evaluating this, we measure throughput improvement over four 3-hop paths in our outdoor testbed over which we can schedule two of the non-interfering links (i.e. *non-intermediate SIR links*) simultaneously.

Path	Throughput (Kbps)		Jitter (ms)		Loss (%)		Throughput improvement (%)
	NSR	WSR	NSR	WSR	NSR	WSR	
1	1169.1	1735.1	19.5	13.8	17.1	16.4	48.4
2	1096.2	1585.2	20.8	13.3	22.3	24.0	44.6
3	1169.3	1590.3	18.5	14.1	17.1	23.8	36.0
4	1162.2	1517.1	23.5	15.5	17.6	27.2	30.5

TABLE V: Throughput improvement with spatial reuse (NSR: no spatial reuse, WSR: with spatial reuse)

Accordingly Table V shows the throughput, jitter, and losses obtained for uni-directional UDP transfer, over each of these four 3-hop paths. We have used iperf to carry out these measurements and these values are the average of 10 experiments with duration of each experiment being 30 seconds. The packet

loss numbers are high, even in the absence of spatial reuse, due to external interference.

Now the expected throughput improvement over a 3-hop path, when two of the non-interfering links are scheduled simultaneously is 3/2 i.e. 50% throughput improvement compared with when there is no spatial reuse. As can be seen from Table V, the throughput improvement with spatial reuse for the first two paths is close to 50%. For the other two paths it is lower than ideal.

We see that for these two latter paths, the packet loss (as seen from Table V) also slightly increases for TDMA with spatial reuse, compared to that of TDMA without spatial reuse. This increase in error corresponds to a case of erroneous PDR prediction by our SIR technique. We however note that even in these cases, the throughput improvement is more than 30%.

Another observation from Table V is that the jitter values are also about 5-6ms less for TDMA with spatial reuse. Thus apart from throughput improvement, decrease in jitter is another benefit of incorporating spatial reuse, particularly for applications such as real time voice and video.

VI. QUANTITATIVE COMPARISON WITH PRIOR STUDIES

In this section we compare with the SIR based interference modeling techniques in literature [19], [9], [13]. We also make quantitative comparisons of our work with link characterization studies that have investigated SNR correlation with PDR [1], [6], [7], [8].

A. Comparison with interference modeling techniques

[19], [9], [13] are some of the interference modeling techniques in literature that make use of SIR based method. However, none of these works have studied in detail the SIR vs PDR correlation, particularly in outdoor environment. So accordingly [13] assumes strict correlation of PDR with SIR in the transition region and outlines a scheduling mechanism to take advantage of the same. However, our results indicate that in the transition region of the SIR curve, it is difficult to predict performance and the PDR prediction error can be as high as over 70%.

Likewise [19], [9], also do not take into account the high PDR prediction error in the *intermediate SIR region*. On the other hand these studies assume that SIR vs PDR correlation is same as SNR vs PDR correlation. However as pointed out in Sec. III, there are essential differences in the SIR correlation with PDR compared with SNR correlation with PDR.

We now apply the model proposed in [19] to our outdoor mesh network and evaluate the performance.

Model Description: [19] proposes a model for predicting interference in CSMA based networks, wherein one component of the model is for measuring hidden terminal interference. For predicting performance, the model makes use of SNR vs PDR curve measured separately at each receiver.

The model also accounts for the effect of external interference. It observes that the RSS reported by the commodity hardware corresponds to the sum of the signal strengths from sender and external interference. So, for estimating external

interference at a receiver, it first measures excess RSS (*measured RSS for a packet from sender A - min observed RSS for sender A*) for each of the packets received from each of the sender. It then estimates external interference as the mean of the excess RSS variation across all the senders.

For predicting PDR when sender and interferer are transmitting simultaneously, it first computes equivalent RSS for sender, in presence of interferer, at the receiver as,

$$RX_{sr}^t = \bar{R}_{sr} - \delta_r(\bar{R}_{tr} - \bar{I}_r) \quad (1)$$

where \bar{R}_{sr} and \bar{R}_{tr} are the mean observed RSS from sender s and interferer t at receiver r respectively, \bar{I}_r is the estimate for external interference, while δ_r is the SINR threshold for successfully decoding the packet. It then predicts PDR for this equivalent RSS for sender, RX_{sr}^t , from the SNR vs PDR curve measured at the receiver r.

Results: We applied the model to our experiments in Sec. IV. However we observed that the model grossly overestimates the PDR. Even when SIR value is 1dB, the model predicts the PDR to be over 91% whereas the measured PDR is only 10% as can be seen from Table VI below, that shows the actual PDR and the PDR as predicted by the model for one of the outdoor links at 6M.

SIR(dB)	1	2	3	4	5	6	7	8	9
Actual-PDR(%)	10	13	41	55	83	95	93	95	96
Predicted-PDR(%)	91.9	91.9	93.1	95.2	92.5	100	100	100	100

TABLE VI: Actual and predicted PDR values using [19]

Analysis: On analyzing, we observed that the model in [19] performs poorly due to following incorrect assumptions,

(i) Variation in RSS quantifies external interference

- Firstly given that granularity of RSS measurement for commodity hardware is 1dBm, a small error in estimating the RSS variation can cause large error in estimating external interference.
- Secondly we note that even in static networks, RSS can vary due to factors other than interference such as multipath fading.

Further, we compared the estimate for external interference i.e. \bar{I}_r in Eqn 1 for one of the bins in our experiments of Section IV, using two approaches. First we computed \bar{I}_r using only the packets received from sender. Then we computed \bar{I}_r for the same bin, using only the packets received from the interferer. Now \bar{I}_r using the two approaches should be similar. However, we observed there was significant difference in \bar{I}_r estimated using the two approaches. For e.g.. for one of the bins, using the first approach \bar{I}_r is -69.3 dBm while using the second approach it comes out to be -80.6 dBm for the same bin. This indicates that the method for estimating external interference by the model does not apply in general.

(ii) SIR vs PDR is the same as SNR vs PDR

When sender and interferer are active simultaneously, the model in [19] first calculates the equivalent RSS and then predicts the performance using the SNR vs PDR curve. However, as pointed out earlier, measured SIR is essentially different from measured SNR and it is incorrect to approximate the SIR vs PDR curve with the SNR vs PDR curve.

Thus the model in [19] fails as it makes assumptions which do not apply in general. Likewise as [9] also assumes measured SIR is same as measured SNR and as [13] assumes strict correlation of SIR with PDR in transition region, these models are also unlikely to apply in our setting.

B. Link characterization studies

[1], [6] have studied in detail SNR vs PDR correlation for outdoor networks. Accordingly, [1] observes that *intermediate SNR region* in the controlled experiments is about 3dB wide, whereas in outdoor environment this region is much more wider (about 15-20dB). They attribute multi-path fading to be the reason for this behaviour. [6] on the other hand showed that wider *intermediate SNR region* as observed by [1], in outdoor environment, is likely due to interference rather than multi-path fading. For the experiments carried out in interference free outdoor environment, [6] reports that the width of the *intermediate SNR region* is only about 1-2 dB more than that measured in controlled experiments.

Similar to [6], in context of SIR vs PDR correlation, our results show that *intermediate SIR region* is also quite narrow (at-most 4-5dB wide) in outdoor networks. But in addition we observe that the SIR vs PDR correlation is different across different outdoor links and even for a given outdoor link the correlation varies across time. Though [7] reports that SNR vs PDR correlation varies across different outdoor links, but they have not carried out measurements to study the time varying behaviour of SNR vs PDR correlation.

[8] reports that the transition region for SNR vs PDR correlation in indoor testbed setting can exceed 10dB whereas in the controlled set-up it is only 1-2dB. They attribute frequency selective fading, experienced by OFDM modulated signals, as the cause for this widening of transition region in realistic setting of indoor testbed. However our results show that the *intermediate SIR region* for outdoor links is about 4-5dB, which is much less than the 10dB in [8]. More importantly, we observed similar results for DSSS-modulated 802.11b as for OFDM-based 802.11g. Now, 802.11b does not have multiple sub-carrier frequencies link 802.11g. This suggests that the presence of frequency selective fading in our setting of outdoor links, if at all, is minimal.

VII. CONCLUSION

Spatial reuse and interference estimation are important aspects of wireless mesh networks. In this work we have investigated in detail the SIR technique for quantifying hidden terminal interference in outdoor TDMA mesh networks. The technique is more efficient than direct measurement based approaches, and also has the predictive capability of model-based techniques. Distinct from prior studies, we have explored and evaluated the SIR vs PDR relation for *outdoor, 802.11g* links. We observed through experiments that the SIR vs PDR curve needs to be measured for several combinations of sender and interferer data rates that employ different PHY layer modulation schemes of 802.11b/g; in contrary, prior work has assumed that PDR is dependent only on interferer's transmit power, not on its modulation [12], [19], [9], [3].

Our results show that with an *offline prediction model* that uses the SIR vs PDR curve from a controlled set-up, our SIR technique can predict performance in outdoor network with good accuracy: less than 10% error for more than 85% of links outside the *intermediate SIR region*. This offline model does not require any per-link real-time measurements of the SIR vs PDR relation, and thus has negligible overhead. However the offline model has a high prediction error for links that have SIR values in the *intermediate SIR region*. Fortunately, this region is at most 4-5dB wide, and hence there are only a small fraction of links which fall in this region. We have also integrated the SIR technique with a TDMA MAC protocol and shown throughput improvement through spatial reuse.

REFERENCES

- [1] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11 b mesh network. In *SIGCOMM 2004*.
- [2] N. Ahmed, U. Ismail, S. Keshav, and K. Papagiannaki. Online estimation of RF interference. In *CoNEXT 2008*.
- [3] N. Ahmed, U. Ismail, S. Keshav, and K. Papagiannaki. Measuring multi-parameter conflict graphs for 802.11 networks. *ACM SIGMOBILE Mobile Computing and Communications Review*, 13(3), 2009.
- [4] I. Akyildiz, X. Wang, and W. Wang. Wireless mesh networks: a survey. *Computer networks*, 47(4), 2005.
- [5] A. Dhekne, N. Uchat, and B. Raman. Implementation and evaluation of a tdma mac for wifi-based rural mesh networks. In *NSDR 2009*.
- [6] D. Gokhale, S. Sen, K. Chebrolu, and B. Raman. On the feasibility of the link abstraction in (rural) mesh networks. In *INFOCOM 2008*.
- [7] P. Gupta, B. Jain, B. Raman, and P. Kulkarni. Link-level measurements of outdoor 802.11 g links. In *WiMesh 2009*.
- [8] D. Halperin, W. Hu, A. Sheth, and D. Wetherall. Predictable 802.11 packet delivery from wireless channel measurements. *SIGCOMM*, 2010.
- [9] A. Kashyap, S. Ganguly, and S. Das. A measurement-based approach to modeling link capacity in 802.11-based wireless networks. In *MOBICOM 2007*.
- [10] S. Kawade and T. Hodgkinson. Analysis of Interference Effects between Co-Existent 802.11 b and 802.11 g Wi-Fi Systems. In *VTC, 2008*.
- [11] J. Lee, S. Lee, W. Kim, D. Jo, T. Kwon, and Y. Choi. RSS-based carrier sensing and interference estimation in 802.11 wireless networks. In *SECON 2007*.
- [12] X. Liu, A. Sheth, M. Kaminsky, K. Papagiannaki, S. Seshan, and P. Steenkiste. Dirc: increasing indoor wireless capacity using directional antennas. In *SIGCOMM 2009*.
- [13] R. Maheshwari, J. Cao, and S. Das. Physical interference modeling for transmission scheduling on commodity wifi hardware. In *INFOCOM 2009*.
- [14] D. Niculescu. Interference map for 802.11 networks. In *IMC 2007*.
- [15] J. Padhye, S. Agarwal, V. Padmanabhan, L. Qiu, A. Rao, and B. Zill. Estimation of link interference in static multi-hop wireless networks. In *IMC 2005*.
- [16] L. Qiu, Y. Zhang, F. Wang, M. Han, and R. Mahajan. A general model of wireless interference. In *MOBICOM 2007*.
- [17] B. Raman and K. Chebrolu. Design and evaluation of a new MAC protocol for long-distance 802.11 mesh networks. In *MOBICOM 2005*.
- [18] B. Raman and R. Jain. SIR-Based Interference-Maps for TDMA-Based Outdoor Mesh Networks. In *Invited paper, LANMAN 2010*.
- [19] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan. Measurement-based models of delivery and interference in static wireless networks. In *SIGCOMM 2006*.
- [20] M. Sha, G. Xing, G. Zhou, S. Liu, and X. Wang. C-mac: Model-driven concurrent medium access control for wireless sensor networks. In *INFOCOM 2009*.
- [21] V. Shrivastava, S. Rayanchu, S. Banerjee, and K. Papagiannaki. Pie in the sky: online passive interference estimation for enterprise w lans. In *NSDI 2011*.
- [22] D. Son, B. Krishnamachari, and J. Heidemann. Experimental study of concurrent transmission in wireless sensor networks. In *SenSys 2006*.
- [23] M. Vutukuru, K. Jamieson, and H. Balakrishnan. Harnessing exposed terminals in wireless networks. In *NSDI 2008*.