

Understanding the Limitations of Transmit Power Control for Indoor WLANs *

Vivek Shrivastava, Dheeraj Agrawal, Arunesh Mishra, Suman Banerjee

University of Wisconsin-Madison

Madison, WI 53726

{viveks,dheeraj,arunesh,suman}@cs.wisc.edu

Tamer Nadeem

Siemens Corporate Research

Princeton, NJ 07040

tamer.nadeem@siemens.com

ABSTRACT

A wide range of transmit power control (TPC) algorithms have been proposed in recent literature to reduce interference and increase capacity in 802.11 wireless networks. However, few of them have made it to practice. In many cases this gap is attributed to lack of suitable hardware support in wireless cards to implement these algorithms. In particular, many research efforts have indicated that wireless card vendors need to support power control mechanisms in a *fine-grained* manner – both in the number of possible power levels and the time granularity at which the controls can be applied. In this paper we claim that *even if fine-grained power control mechanisms were to be made available by wireless card vendors, algorithms would not be able to properly leverage such degrees of control in typical indoor environments*. We prove this claim through rigorous empirical analysis and then build a tunable empirical model (Model-TPC) that can determine the granularity of power control that is actually useful. To illustrate the importance of our solution, we conclude by demonstrating the impact of choice of power control granularity on Internet applications where wireless clients interact with servers on the Internet. We observe that the number of feasible power was found to be between **2-4** for most indoor environments. We believe that the results from this study can serve as the right set of assumptions to build practically realizable TPC algorithms in the future.

Categories and Subject Descriptors

C.4 [Computer Systems Organization]: Performance of Systems; C.2.1 [Computer Communication Networks]: Network Architecture—*Wireless Communication*

*Traces used in this paper will be available at www.cs.wisc.edu/~viveks/powercontrol

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.

IMC'07, October 24-26, 2007, San Diego, California, USA.

Copyright 2007 ACM 978-1-59593-908-1/07/0010 ...\$5.00.

General Terms

Algorithms, Experimentation, Measurement, Performance

Keywords

IEEE 802.11, Transmit Power Control, RSSI Modeling, Indoor WLAN, Fine-Grained, Limitations, Kullback-Leibler

1. INTRODUCTION

Power control mechanisms in wireless networks have been used to meet two different objectives — to reduce energy consumption in mobile devices, so as to conserve battery life, and to reduce interference in the shared medium, thereby allowing greater re-use and concurrency of communication. In this paper, our focus is to study power control as applicable to the interference reduction objective. As an example in this paper, we consider the impact of power control for WLAN clients interacting with servers on the Internet. Recent theoretical work has shown that ideal medium access protocol using optimal power control can improve channel utilization by up to a factor of $\sqrt{\rho}$, where ρ is the density of nodes in the region (using fluid model approximations) [9]. Power control mechanisms [11, 20, 18] typically try to optimize the floor space acquired by wireless transmissions by limiting the transmit power of control and data packets, thereby providing opportunity for multiple flows to coexist.

A number of research efforts have studied power control based on the theoretical abstraction of wireless signal propagation in free space and consider transmit power as a continuous variable (i.e., a fine grained parameter), that can be set per packet to yield optimal performance. Conventional power control mechanisms have exercised fine grained control in the two dimensions as shown in Figure 1 : 1) Time granularity at which power level is changed, 2) Magnitude granularity by which the power level is changed. We analyze both the dimensions of fine grained power control and provide guidelines for power control granularity in typical indoor environments.

Prior work [8] has pointed out that lack of vendor support for fine-grained power control mechanisms in the wireless cards inhibit deployment of these mechanisms. In this paper we ask the following questions: *Is fine-grained power control really useful and would lead to a better design of power-control algorithms? If not, what is the minimum granularity*

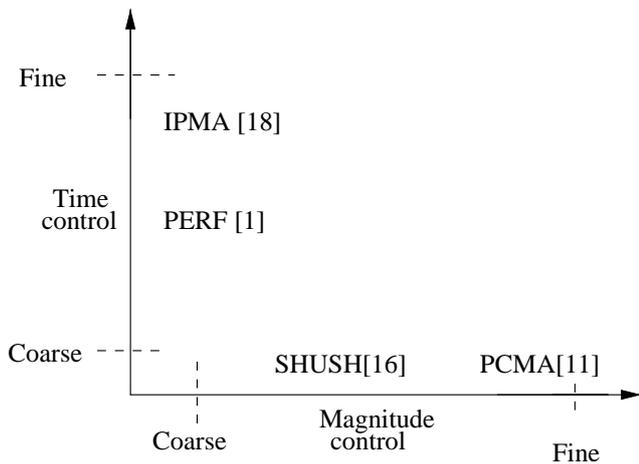


Figure 1: Two dimensions of transmit power control taken by prior approaches. PCMA, SHUSH rely on changing transmit power by small values (1dBm) and lie on the magnitude dimension. IPMA, Subbarao et. al. rely on changing the transmit power on a per packet basis and hence lie on the time dimension

of power control that is useful in different wireless environments, including Internet oriented wireless communication? We answer the first question in the negative. As we discuss in detail in the paper, in practical indoor wireless LAN (WLAN) deployments, multipath and fading effects, coupled with various sources of interference in the unlicensed bands, imply that power control algorithms cannot derive significant benefits from very fine-grained mechanisms. We demonstrate this through detailed experimentation in different indoor wireless network environments. We estimate the distributions of Received Signal Strength Indicator (RSSI)¹ for various transmit power levels at the transmitter and show that although more power at the transmitter on average translates to more power at the receiver, there is significant overlap between the RSSI distributions for nearby power levels, making them practically indistinguishable at the receiver. This can be attributed to dominant multipath and fading effects, that lead to significant signal strength variations in indoor environments.

Our answer to the second question is that a power control algorithm can make practical use of only a *few* (2-3) discrete number of power levels. The exact number and choice of power levels is a characteristic of the multipath and fading of a particular wireless environment and the presence of other interfering sources.

Our observations are true for both ad-hoc networks and Internet oriented wireless communications (WLANs), and in this paper we present our results from the latter setting. In particular, through this work we build an empirical model

¹Variations in RSSI typically correspond to variations in Signal to Noise Ratio (SNR) as shown by Reis et. al in their measurement based study of delivery and interference models for static wireless networks [7]. Moreover commodity wireless cards only report the RSSI values for each packet and hence we base our observations on the measurement for RSSI values. We further discuss this in detail in Section 3.

that allows us to characterize the specific set of power levels that is useful for a given environment and could be used to perform per packet power control.

Power control is also an important design consideration in cellular networks, where it is primarily used to counter fast fading. However, cellular networks primarily operate in outdoor environment, where we show that the effect of multipath is not significant enough to hinder fine grained TPC. We discuss more about relevance of our work in context of cellular networks in Section 6.

Key contributions

The following are the key contributions and the main observations from this work:

- **Measurement:** We collect extensive traces from multiple environments such as office building and university departments, to characterize Received Signal Strength Indicator (RSSI) variations in different indoor settings. Through rigorous statistical analysis of the traces using *Allan's Deviation* (for characterizing the burst size of RSSI fluctuations) and *Normalized Kullback-Leibler Divergence (NKLD)* (for characterizing RSSI distribution in real time), we observe that the number of feasible power levels that can be used in a transmit power control mechanism is few and discrete, and once identified, could be used to perform power control at small time scales (per packet).
- **Model:** Through this analysis, we propose an empirical model to determine the set of useful power levels in an online fashion, i.e., this model is computed and adjusted dynamically as wireless data communication is going on. Note, that the number and choice of such power levels would depend on individual wireless environment. In all our experimental scenarios, it was found to be *less than 4* and often much less.
- **Validation:** Through Internet-oriented wireless experiments, we demonstrate the usefulness of the our measurement based empirical model (Model-TPC) for improving performance of wireless clients interacting with servers on the Internet in our indoor WLAN deployment. In particular, we show that correct choice of power levels can lead to actual throughput gains in indoor environments.

We believe that our our experiments highlight some fundamental issues with transmit power control, that can help in design of future wireless interfaces that are used in laptops, PDAs and are widely used as a major Internet access mechanism.

The remainder of the paper is organized as follows. Section 2 motivates the infeasibility of fine grained power control in indoor WLANs and discusses various transmit power mechanism proposed in literature, with their respective evaluation in context of our practical models for transmit power control. In Section 3, we analyze the RSSI distributions under varying indoor scenarios and propose an online mechanism (Online-RSSI) to characterize the distribution in real time. We use the online mechanism to derive an empirical model for transmit power control (Model-TPC) described in Section 4. Section 5 highlights the impact of using our empirical model on end user experience through Internet oriented wireless experiments. We briefly discuss power control

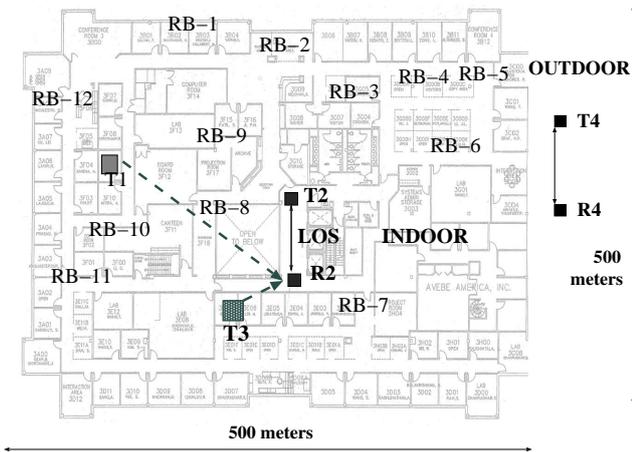


Figure 2: The wireless testbed, consisting of seven 802.11 a/b/g nodes (transmitters marked by T1, T2 and receivers marked by RB-1 - RB 12)). The dotted arrows indicate the transmitter-receiver pair T1-R2 and T3-R2 for our Internet oriented experiments.

in cellular networks in section 6 and some related work in 7. Finally we conclude in section 8.

2. MOTIVATION : POWER CONTROL APPROACHES AND LIMITATIONS

Implementation of fine grained power control mechanisms has been limited by the hardware support in current 802.11 wireless cards which have only limited number of discrete power levels. As described in [8], most of the wireless cards support only 4 to 5 power levels at the hardware, which is in stark contrast to the fine grained control preferred by most power control schemes like PCMA [11], SHUSH[18] and IPMA [20]. This being a limitation of current state of the art hardware, can be resolved in future wireless cards that may support fine grained power levels. However, we argue that *there are fundamental constraints to power control in indoor wireless environments, which limits the number of feasible power levels that is useful in such mechanisms.* We substantiate our claim through indoor WLAN and outdoor experiments in the following section, where we show that RSSI variations are present in both outdoor and indoor environments, but are especially dominant in indoor scenarios.

2.1 Infeasibility of Fine Grained Power Control

We present preliminary results from our detailed set of experiments explained in Section 3 to illustrate the fundamental constraints of fine-grained power control.

Outdoor Scenario

This sample experiment consists of a pair of outdoor transmitter-receiver pair (T4-R4) shown in Figure 2 operating using the 802.11a standard. At R4 we capture the packets transmitted by T4 for different power levels available at T4's Atheros based wireless chipset. Since low RSSI is more likely to cause a packet error, we have enabled Madwifi driver to receive packets in error and in order to prevent the bias towards

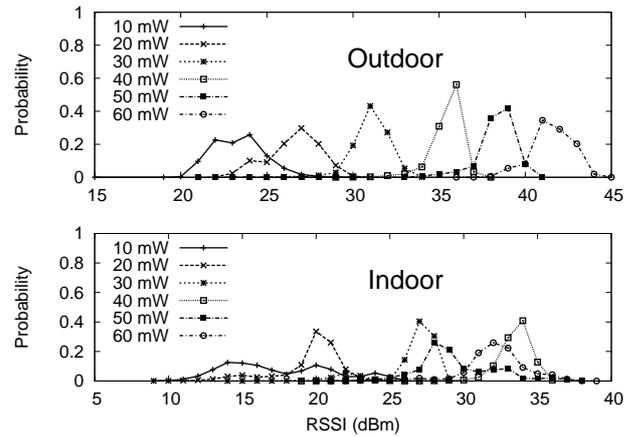


Figure 3: Probability Distribution of RSSI for varying power levels at the transmitter is shown in the figure. The top figure corresponds to outdoor scenario with 6 distinguishable power levels while bottom figure shows the effect of increased multipath and interference in the indoor WLAN scenario with the number of distinct power levels reduced from 6 to 3. Band:802.11g Data Packet Size:1Kbytes

high RSSI values, we include the RSSI of erroneous packets in our calculations for RSSI distributions. Figure 3 shows the probability density function of RSSI distribution for various power levels at the transmitter. The power levels are increased from 10mW to 60mW (max. transmit power), in steps of 10mW. For the sake of clarity, these power levels are chosen so that there is minimal overlap between their respective RSSI distributions. For example at a power level of 60 mW, the RSSI values vary from 35dBm to 45dBm, with 40 percent of the packets being received at 41dBm. The average variation in RSSI value over all power levels is approximately 7.5 dBm.

This variation can be attributed to the multipath and fading effects, due to which the packets transmitted at the same power level, may be received with varied signal strength at the receiver. A difference of an order of wavelength in the paths taken by the wireless signals from the transmitter, can lead to the two signals being out of phase [16], resulting in variations in the signal strength at the receiver. Even though more power at the sender translates to more power at the receiver, the distributions of the received signal power overlaps significantly, thereby making them hardly distinguishable. As we show next, this effect is even more pronounced in indoor environments than in outdoor environments where there are only a few strong paths that impact the signal.

Indoor Scenario

We repeat the aforementioned experiments for an indoor transmitter-receiver pair T2-R2 as shown in Figure 2. The resulting distribution of RSSI values is shown in Figure 3. As expected the RSSI variations increase, thereby increasing the overlap between RSSI of neighboring power levels. This observation indicates that in indoor settings, the number of power levels having non-overlapping RSSI distributions are further reduced, thereby making fine-grained transmit power

control much less effective. These experiments show that fine grained transmit power control mechanism are much more difficult to realize in indoor deployments.

It is evident from Figure 3 that in a collective fashion, the distribution of all the six power levels cover a wide range of RSSI values (20 - 45 dBm). Also note that for any single power level, its RSSI distribution overlaps significantly with that of neighboring power levels. The introduction of fine grained power levels at the hardware will imply significant overlap between the distribution of existing power levels (0,10,14,15,17,18)dBm and the new power levels. *A significant overlap between the RSSI distributions of two (successive) power levels correspondingly diminishes the practical effect of having the respective distinct power levels – they become practically indistinguishable at the receiver.* This can be considered analogous to the concept of channels in 802.11 band, where there are 11 channels available but only 3 channels are non overlapping and hence useful. Similarly, fine grained power levels cannot be distinguished easily at the receiver due to RSSI variations and hence may not be useful simultaneously.

We performed the same set of experiments at two different location, at the NEC Research Labs at Princeton and at the Computer Sciences Department at University of Wisconsin-Madison. We observed that, although the exact shape of the RSSI distribution may depend on the exact indoor environment and other interference effects, the general nature remains similar to Figure 3. In this paper, we report our measurements from the NEC Research Labs, which we believe are representative of a typical indoor WLAN scenario.

Next we summarize prior approaches proposed in the literature that rely on fine grained power control. We show why such approaches might face difficulty in a practical implementation. We also discuss how our proposed empirical model could act as an oracle to guide such algorithms to change transmit power that are effective in practice.

2.2 Implications on Existing Power Control Approaches

We categorize some of the prior power control methods applicable to WLANs into two categories : 1) fine-grained in magnitude of transmit power and 2) fine-grained in magnitude of time (per-packet). Existing power control approaches can be categorized in the two aforementioned categories as shown in Figure 1. We explain the implications of our observations on both categories of protocols:

Magnitude Dimension of Fine Grained Power Control

Monks et al. proposed a power controlled multiple access wireless MAC protocol (PCMA [11]), within the collision avoidance framework. PCMA generalizes the transmit-or-defer "on/off" collision avoidance models to a more flexible "variable bounded power" collision suppression model. Using PCMA, the transmitter-receiver pairs can be more tightly packed into the network by adjusting the power level of the transmitter to the minimum required for a successful transmission, thereby allowing a greater number of simultaneous transmissions (spectral reuse). In order to ensure successful packet delivery, each receiver in PCMA first calculates the extra noise that it can tolerate, such that the SNR for its own packets is above the threshold for correct reception. It then advertises this noise tolerance by sending a busy tone on the auxiliary channel, and all the transmitters in the

vicinity measure the received signal strength of the tone to determine the maximum power with which they can initiate their own transmissions. This mechanism requires exact calculations of received power, which may not be predictable under multipath and fading effects. Moreover, the authors treat transmit power as a continuous parameter, which may not be feasible in indoor environments due to significant RSSI variations.

Seth et al. propose a reactive transmit power control mechanism, called SHUSH [18], where nodes operate on the optimum (minimum) power required for communication. On detecting interference, SHUSH calculates the exact power required to send a RTS to the interferer and hence optimizes the "floor space" acquired by any flow. Unlike PCMA, however SHUSH transmits at a higher power only when a flow is interrupted by external interference. Again SHUSH assumes fine grained control on power levels and ignores RSSI variations which can make it difficult to infer the exact interference at the receiver, thereby complicating the calculation of target transmit power required to SHUSH the interferer. Our experimental observations suggest that such observations are too deviant from realistic scenarios. Using our empirically derived power control model (Section 4), the above mechanisms could dynamically determine an exact set of feasible power values to be used in an environment.

Time Dimension of Fine Grained Power Control

Many researchers in the past have proposed schemes which require change in the power level on a per packet basis.

Akella et al. [2] discuss some power control mechanisms in their work on wireless hotspots. They propose that APs should use the minimum transmit power required to support the highest transmission rate. In their scheme, the receiver sends the value of observed RSSI, averaged over some small number of packets, as a feedback to the transmitter. The transmitter on receiving the average RSSI value on the receiver side, decides the optimal power level suitable for use in the current channel conditions. However they do not provide exact values for power level granularity that should be used. As discussed earlier, a simple average of RSSI values at the receiver may not give a correct estimate of the actual SNR.

Subbarao [19] has proposed a dynamic power-conscious routing mechanism that incorporates link layer and physical layer properties in routing metrics. It routes the packet on a path that requires least amount of total power expended and each node transmits with the optimum (minimum) power to ensure reliable communication. This scheme requires per packet power control and also needs feedback from the destination regarding RSSI on a per packet basis.

Similar to PCMA approach, Yeh et al. [20] proposed an interference/power aware access control. They augment the normal RTS/CTS mechanism of IEEE 802.11 with provision for multi level RTS, where the transmit power of the RTS mechanism is set on the basis of the intended receiver. Such a dynamic per packet approach becomes difficult in the face of significant RSSI variations and become difficult to implement on real systems.

We analyze the stationarity (coherence time) of signal strength for various scenarios and propose a simple algorithm Online-RSSI, that can be used to determine the distribution of signal strength for a given transmit power level in any scenario. Once the set of feasible power levels (hav-

ing non overlapping signal strength distribution) is derived, the receiver can use this model to determine the transmit power of the transmitter for a packet received at any given signal strength and hence provide correct feedback to the transmitter on a per packet basis (or similar time scales).

3. CHARACTERIZING SIGNAL STRENGTH DISTRIBUTION

Our experiments serve three main purposes: (i) to gain an understanding of the characteristics of RSSI variations under varying practical scenarios (in terms of user movements, shadowing, multipath and external interference) (ii) as a learning data-set to build our empirical model for identifying the set of feasible power levels (iii) as an input to validate this model.

In this section, we characterize the distribution of RSSI under varying magnitudes of multipath, shadowing and other 802.11 and non 802.11 interference for a real WLAN deployment shown in Figure 2. By studying the RSSI distribution across different power levels and different channel conditions, we formulate mechanisms to dynamically predict and construct such distributions in real-time. Such mechanisms shall be used in the next section where we build a model to predict the useful power-levels in a given environment. We briefly describe various components of our experimental setup.

3.1 RSSI measurements

The performance of most wireless applications depends on the packet delivery probability. The SNR is widely used in the literature to model packet delivery probabilities: packets are successfully received if $S/(I+N)$ is above a certain threshold, and otherwise are not. Commodity wireless cards do not report the information required to compute SNR. For instance, our cards report only their version of RSSI, 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. It has been shown in a prior measurement based study by Reis et. al [7] that RSSI is generally predictive of delivery probability in static wireless networks and while wireless networks exhibit substantial variability, measurements of average behavior over even relatively short time periods tend to be stable. This phenomenon was also observed in our joint power and data rate adaptation experiments (described as an application of our model in Section 5), where the power levels with significant overlap in their corresponding RSSI distribution, perform similarly in terms of rate adaptation. Since rate adaptation again depends on packet delivery rate, we can infer that RSSI is a reasonable estimate for SNR and two power levels with significant RSSI overlap at the receiver will perform similarly for packet delivery probabilities. Hence we base our measurements and models on RSSI values that is readily available from the commodity wireless cards.

RSSI estimates signal energy at the receiver during packet reception, measured during PLCP headers of arriving packets and reported on proprietary (and different) scales. Atheros cards, for example report RSSI as $10\log_{10}(\frac{S+I}{n})$, where S is the signal strength of the incoming signal, I is the interfering energy in the same band, and n is a constant ($-95dBm$) that represents the "noise floor" inside the radio. Atheros

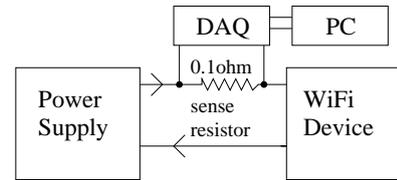


Figure 4: Figure shows the setup used to determine power drawn by wireless cards. The DAQ samples voltage across the WiFi device and sends it to a PC via USB. Performed at transmitter to validate the power levels available at the hardware.

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, RSSI is defined to be $S+I$. Note that these RSSI measurements are performed at the receiver and then provided as a feedback to the transmitter for constructing the empirical model for feasible power levels.

3.2 Validating Available Hardware Power Levels

To ascertain the available power levels in 802.11 WLAN cards, we measure the voltage across the wireless card of the transmitter by the setup shown in figure 4. The setup constitutes of a 0.1 ohm sense resistor, R , connected in series to the circuit of the wireless device (pcmcia card) that exposes the voltage supplied to the device. For the pcmcia based 802.11 card, we used the Sycard 140A cardbus adapter, to expose the voltage supplied to the card. A Data Acquisition Card (DAQ), DS1M12 Stingray Oscilloscope, samples the voltage through R at a rate of 1 million samples per second, thereby giving us voltage measurements on a per packet basis. The instantaneous power consumption, P_i can therefore be written as $P_i = V_d \times V_R/R$ where V_d is the voltage provided to the WiFi device and V_R is the voltage drop across R at a given moment. These measurements are performed at the transmitter and shows that indeed the right power levels are implemented at the hardware circuitry of the transmitter's wireless interface. On the basis of power consumed by the wireless interface, we validated that Cisco Aironet cards provide 6 different power levels for 802.11g and 5 different power levels for 802.11a respectively.

3.3 WLAN Trace Collection

In order to understand the behavior of RSSI under varying interference and multipath effects, we conduct detailed experiments to collect RSSI traces in an office building under varied indoor settings. In all our experiments, we use a fixed data rate of 1Mbps and fixed packet size of 1KB, so that the time intervals are directly translated into number of packets (modulo 802.11 DCF), which is the X axis for most of our plots. This facilitates easier packet based analysis of RSSI traces and their implications to power control mechanisms, which generally base their decisions on a per packet basis. For our experiments, 1 sec of receiver time window \approx 1000 packets of 1KB each (unless otherwise specified). *We repeated the same experiments with other wireless cards and found the results were consistent with the ones reported*

here. We discuss the exact set up for each of these scenarios.

Line of Sight - light interference(LOS-light)

These experiments represent a scenario where the transmitter-receiver pair are in direct line-of-sight and have minimal to zero external interference. Figure 2 shows the placement of transmitter-receiver pair T2 and R2 respectively for LOS-light experiment. The experiment used 2 IBM Thinkpad laptops running Linux kernel 2.6. Each of the laptops housed an Atheros chipset based 802.11a/g Linksys wireless card and used Madwifi drivers. We used *Netperf 2.2* to generate UDP flows between the two laptops and collected MAC-level traces for the packets received at the receiver using the *pcap* standard library. We vary the power of the transmitter to understand their corresponding effects on RSSI.

Non Line of Sight - light interference(NLOS-light)

The experiment comprises of a single transmitter T1 and 5 receivers (RB-1, RB-8, RB-10, RB-11 and RB-12) as shown in Figure 2 placed at various locations in the building and used *netperf* and *pcap* library to generate flows and collect traces respectively. None of the receivers were in direct line-of-sight of T1 and this setup too had minimal to zero external interference.

Line of Sight - heavy interference(LOS-heavy)

We investigate the effect of controlled interference on RSSI. We use our experimental testbed shown in figure 2 for line of sight experiments to evaluate the effect of heavy interference (like bulk data transfers) on RSSI variations. Nodes RB-12, RB-11 and RB-2 act as separate APs and perform bulk data transfers with their respective clients (3 IBM laptops). Nodes T2 and R2 form a transmitter-receiver pair.

Non Line of Sight - heavy interference(NLOS-heavy)

We use our experimental testbed shown in figure 2 for non-line of sight experiments to evaluate the effect of heavy interference (like bulk data transfers) on RSSI variations. Nodes RB-12, RB-11 and RB-2 act as separate APs and perform bulk data transfers with their respective clients (3 IBM laptops). Nodes T1 and RB-8 form a transmitter-receiver pair.

3.4 Analyzing WLAN Traces

Figure 5 shows the smoothed moving average of RSSI per packet for the four categories of traces described in previous section. Although we collect many traces from each category (namely LOS-light, NLOS-light, LOS-Heavy and NLOS-heavy), we present only one representative trace from each category. The representative trace is chosen such that it manifests the basic characteristic of traces from that particular category. All these traces are collected at 1Mbps of data rate with packet size of 1000 bytes.

As clear from Figure 5, the variations in RSSI is minimum for LOS-light trace and is maximum for the NLOS-heavy trace. This behavior is expected because the factors contributing to RSSI variations increase in both number and magnitude from the topmost plot to the bottom. Figure 6 show the probability distribution of RSSI values at the receiver for the four scenarios. Clearly, the distribution of RSSI becomes flatter (larger variation) with the increase in interference and multipath effects, with the distribution of LOS-light and NLOS-light resembling a Gaussian distribution. Next we analyze these trace in detail to understand

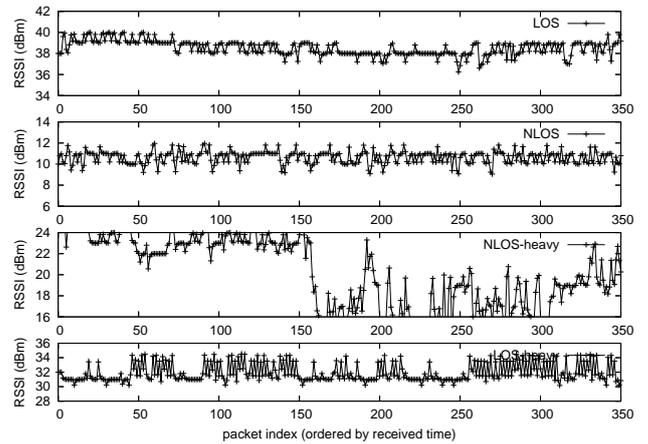


Figure 5: Exponentially weighted moving average of RSSI over time for four traces collected under various practical scenarios, with varying degree of external interference, multipath, shadowing and fading effects. The packets are sorted in order of received time. The traces from topmost plot to the bottom belong to LOS-light, NLOS-light, NLOS-heavy and LOS-heavy. The high variation of RSSI for NLOS-heavy can be observed.

temporal variations in RSSI and propose an algorithm to dynamically characterize the distribution of RSSI in any environment.

Stationarity

Figure 5 shows the variation of RSSI on a per packet basis, but it would also be useful to observe the amount of fluctuation over a set of packets (or a burst). Such an analysis would reveal any characteristic burst intervals where RSSI values vary largely over different bursts but deviate minimally within a burst. Also note that since our experiments are conducted with the traffic sent at uniform rates packet intervals directly correspond to time intervals (modulo 802.11 DCF effects). One way to summarize changes at different time scale is to plot the Allan deviation [3] at each packet interval. Allan deviation is the square root of the two sample variance formed by the average of the squared differences between successive values of a regularly measured quantity taken from sampling periods of the measurement interval. Allan deviation differs from standard deviation in that it uses differences between successive samples, rather than the difference between each sample and long term mean. In this case, the samples are the fraction of packets delivered in successive intervals of a particular length. The Allan deviation is appropriate for data sets where data has persistent fluctuations away from the mean. The formula for the Allan deviation for N measurements of T_i and sampling period τ_0 is:

$$\sigma_y(\tau_0) = \sqrt{\frac{\sum_{i=1}^{N-1} (T_{i+1} - T_i)^2}{2(N-1)}} \quad (1)$$

The sampling period is varied by averaging n adjacent values of T_i so that $\tau = n\tau_0$. Now the Allan deviation for different

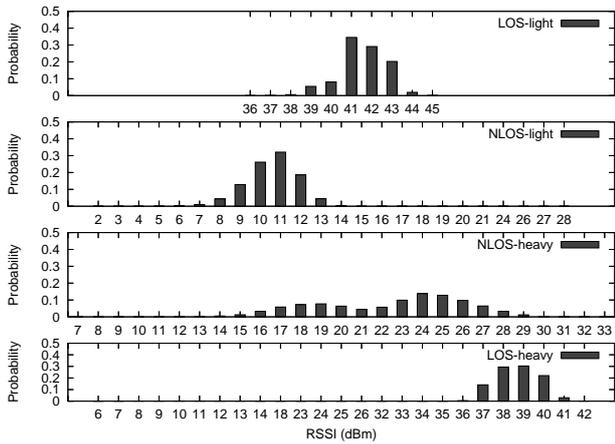


Figure 6: Probability distribution of RSSI for the four traces shown in Figure 5. The spread in RSSI distribution is noticeable in all the traces, with the NLOS-heavy trace having the maximum deviation. In the NLOS-heavy scenario, the RSSI values show persistent fluctuations about two different RSSI values (bimodal distribution).

values of n can be given by:

$$\sigma_y(\tau) = \sqrt{\frac{\sum_{i=1}^{N-2n+1} \left[\frac{1}{n} \left(\sum_{j=i+n}^{i+2n-1} T_j - \sum_{j=i}^{i+n-1} T_j \right) \right]^2}{2(N-2n+1)}} \quad (2)$$

The Allan deviation inherently provides a measure of the behavior of the variability of a quantity as it is averaged over different measurement time periods, which allows it to directly quantify and distinguish between different types of RSSI variations. The Allan deviation will be high for interval lengths near the characteristic burst length. At smaller intervals, adjacent recent samples will change slowly, and the Allan deviation will be low. At longer intervals, each sample will tend towards the long term average, and the Allan deviation will again be small.

Figure 7 shows the Allan deviation of RSSI over large scale packet intervals (thousands of packets). We can observe that although there are no prominent peaks for the RSSI bursts for any scenario, but Allan Deviation becomes quite stable (between 0.2 and 0.5) for LOS-light, NLOS-light and LOS-heavy scenarios. The NLOS-heavy has relatively higher deviation and shows significant fluctuations in the range of (1.6-1.8). In Figure 8, we show the zoomed version for Allan deviation for intervals less than 100 packets. This figure shows the short term characteristic of RSSI variations. As clear from the figure, Allan deviation for LOS-light, NLOS-light and LOS-heavy is maximum at 1 packet, then decreases sharply because averaging over longer intervals rapidly smoothes out fluctuations. This means that the RSSI variations for the aforementioned three categories are independent for intervals less than 100 packets. On the other hand, NLOS-heavy shows sharp increase in Allan Deviation from 0.6 to 1.4. This indicates that in NLOS-heavy trace, the RSSI averaged over small sample sizes (τ in Equation 2), change quickly leading to a sharp increase in Allan Deviation at such small scales. On further analysis, we found that deviation for NLOS-heavy reaches 1.7 for about 400-

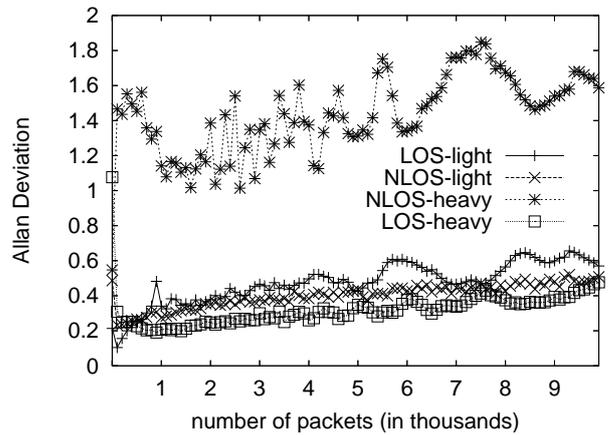


Figure 7: Allan deviation for the four representative traces shown in figure 5. The y axis shows the Allan deviation ($\sigma(\tau)$), while the value of n (sampling period in Equation 2) is varied on the x axis. It shows that there are no clear peaks for the RSSI bursts for any scenario, however it is clear that Allan Deviation becomes quite stable (between 0.2 and 0.5) for LOS-light, NLOS-light and LOS-heavy scenarios. The NLOS-heavy has relatively higher deviation and shows significant fluctuations but remains in the range of (1.6-1.8).

500 packets and as shown in Figure 7, fluctuates around that value for larger packet intervals as well. We agree that there is no clear decrease in the Allan deviation for any scenario, so we approximate the value of burst size at the point when the deviation becomes quite stable (or the rate of increase in deviation becomes very low). Hence we choose ≈ 400 packets for NLOS-heavy and on the order of thousand packets for LOS-heavy, LOS-light and NLOS-light.

We report these burst size for various LOS and NLOS scenarios in Table 1. The burst size information is used by our algorithm Online-RSSI (explained in Section 3.5), that samples the packets in multiples of these burst sizes for determining the signal strength distribution for a given transmit power level. As RSSI varies significantly across bursts, the online mechanism needs to consider at least an increment of burst size in its sampling process to determine if the online distribution being computed has stabilized. This allows us to quickly converge on an accurate RSSI distribution as explained in Section 3.5.

Summary: RSSI variations are bursty for intervals of the order of ≈ 1000 packets for LOS-light, NLOS-light and LOS-heavy scenarios. But for NLOS-heavy traces, the Allan deviation increases even in the small interval of 100 packets, depicting bursts even in short packet intervals. This can be explained because the interference coupled with multipath effects make the wireless channel highly variable and leads to bursts even in very short time intervals. This behavior was observed in all our NLOS-heavy traces (for various receivers) and indicates high variability in wireless environment. Allan deviation provides an estimate of burst length of a trace and could be interpreted as an effect of temporal variations in wireless channel. So if Allan deviation shows that a trace

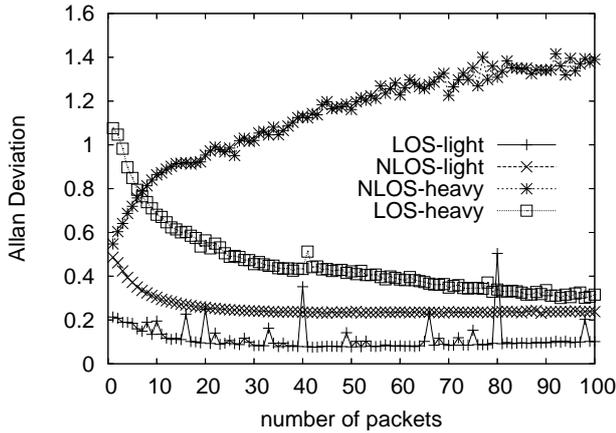


Figure 8: Zoomed version of Allan deviation for short interval of time (≈ 100 packets). Allan deviation decreases sharply for LOS-light, NLOS-light and LOS-heavy traces, indicating independent packet losses. But Allan deviation for NLOS-heavy increases, indicating very small bursts and highly variable wireless channel. This is a strong indication that fine grained power control becomes even more difficult when multipath effects are coupled with 802.11 interference.

has very small burst periods (as in the case of NLOS-heavy), it can be used as an indication that per-packet power control will be highly unpredictable. Finally we observe that all the scenarios show substantial non-stationarity in RSSI variations, which will further impede fine grained mechanisms for power control.

Entropy

Through the empirical analysis presented in Section 2.1, we observed that due to multipath, fading and other propagation effects, the RSSI values at the receiver show significant variation (also corroborated by Figure 6). Depending on the exact environment, RSSI distributions for close transmit power levels can have substantial overlap, making them practically indistinguishable at the receiver. For a power control scheme to be effective, it needs to determine the set of useful power levels i.e. power levels with minimum overlap. In order to estimate the number of power levels in any setting, we need to estimate the corresponding RSSI distribution for various power levels. Ideally, we can sample the RSSI values for a very long period of time ($\approx 10mins$) to obtain the true behavior of the RSSI distribution. But, as we show next, sampling very large number of packets may not be necessary (or practical, due to computation and storage limitation on the clients) in most settings. This observation leads us to the following question: *What is the minimum number of packets we should sample to get a “good” approximation of RSSI distribution ?*

We first describe an offline mechanism to determine the number of samples that are required to generate a distribution close to the one computed over large number of packets, as shown in Figure 6. On the basis of insights obtained from the offline analysis, we then present a simple online mecha-

nism to dynamically determine the number of packets sufficient to characterize RSSI distribution in any environment.

Let us define the actual probability distribution function for RSSI (over large packets $\approx 100,000$) as $p(x)$. The approximate distribution obtained by our mechanism is denoted by $q(x)$. We now describe the statistical measure that we use to quantify the performance of the model.

Let $p(x)$ and $q(x)$ be two probability distribution functions defined over a common set χ . We describe a commonly used statistical measure *Kullback-Leibler Divergence* (KLD) that quantifies the ‘distance’ or the relative entropy between two probability distributions. This comprises a general measure and allows us to compare the statistics of all the orders for the two distributions. The *Kullback-Leibler Divergence* (KLD) [6] is defined as

$$D(p(x)||q(x)) = \sum_{x \in \chi} p(x) \left| \log \frac{p(x)}{q(x)} \right| \quad (3)$$

The KLD is zero when the two distributions are identical and increases as the distance between two distributions increase. The KLD is a measure used in information theory to calculate the ‘distance’ between two distributions $p(x)$ and $q(x)$. The definition of the KLD carries a bias for random variables with higher entropy. Hence to evaluate the relative distance accurately for our purposes, it is important to weigh in the entropy of the original distribution which can be large. The entropy $H(p(x))$ of the random variable x with distribution $p(x)$ is the average length of the shortest description of the random variable given by:

$$H(p(x)) = \sum_{x \in \chi} p(x) \log \frac{1}{p(x)} \quad (4)$$

Hence we use the normalized Kullback-Leibler divergence NKLD [13] defined below as a measure of distance between two distributions

$$\text{NKLD}(p(x)||q(x)) = \frac{D(p(x)||q(x))}{H(p(x))} \quad (5)$$

However the above metric is asymmetric and we make it symmetric by taking an average of $\text{NKLD}(p(x)||q(x))$ and $\text{NKLD}(q(x)||p(x))$. The symmetric average distance between two distributions is given by

$$\text{NKLD}(p(x), q(x)) = \frac{1}{2} \left(\frac{D(p(x)||q(x))}{H(p(x))} + \frac{D(q(x)||p(x))}{H(q(x))} \right) \quad (6)$$

Ideally we could have characterized the distance between two probability distributions by calculating the area of their intersection on some data set X . However this will require calculating their points of intersections and some numerical integration techniques, which may be cumbersome depending on the exact shape of the distribution. Hence we use NKLD as it compares the statistics of all orders for two distributions and is very simple to compute in real time. Further NKLD works efficiently for our experimental scenarios.

We consider the long term probability distribution as $p(x)$ and those derived from our offline mechanism as $q(x)$. Let n be the length of the packet sequence that is used for computing the distribution $q(x)$. The value of n is varied and we measure the corresponding NKLD for each $q(x)$ (with $p(x)$ as the reference long term distribution).

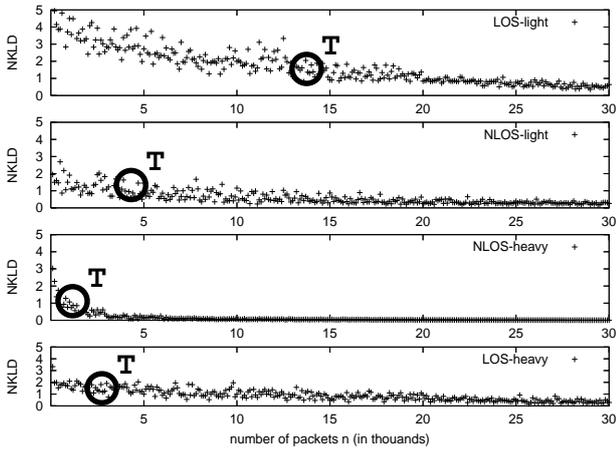


Figure 9: Normalized Kullback-Leibler Divergence (NKLD) for the four representative traces. Clearly for NLOS-heavy trace, NKLD decreases sharply with the increase in number of packets, reaching a value of 1 for a sample size of the order of 5000 packets. For LOS-light however, this value is around 30,000 packets.

Figure 9 shows the NKLD curve obtained for the representative traces from the four categories discussed before. NKLD is a decreasing function of n , although the exact shape of the curve varies as per the environment. We assume without the loss of generality, the tolerable error or relative distance between actual distribution and distribution obtained by sampling n packets be 10%. Figure 9 can be used to calculate the length of packet sequence required to achieve the error bound under varying scenarios. While LOS-light and NLOS-light require about 20,000 packets each, LOS-heavy and NLOS-heavy scenarios require less than 10,000 packets as shown in Table 1.

Summary: The number of packets required to determine a close approximation for RSSI distribution is especially high for the LOS-light scenario while for a NLOS-heavy scenario the number is relatively lower. The accuracy of an RSSI distribution varies directly with the number of bursts captured. Since, the NLOS-trace trace seen has short burst sizes we can obtain large number of bursts using a smaller trace to accurately model the RSSI distribution while the trace required for LOS-light scenario is larger owing to longer burst sizes. This analysis shows that sampling very large number of packets (≈ 100000) to obtain RSSI distribution is not required in majority of traces, with the notable exception of LOS-light scenario.

3.5 Algorithm Online-RSSI

Based on the above analysis, we describe an online algorithm to compute the RSSI distribution in an online fashion by predicting the number of packets needed in order to accurately characterize the distribution in any environment. As shown in Figure 9, initially NKLD (or error) decreases rapidly with the increase in n , but stabilizes after a threshold T , slowly tending to zero. It implies, that beyond a certain length of packet sequence, the decrease in NKLD(or error) is minimal and hence there is not much gain in sampling

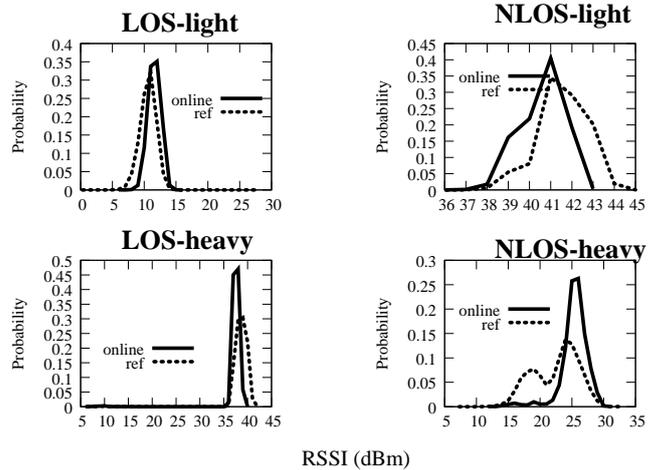


Figure 10: Comparison between distributions obtained from n packets (as determined by the online algorithm) and the true distributions obtained from long term traces. The two distributions are remarkably similar thereby indicating the efficacy of our online mechanism

Online-RSSI(burst_size,tolerance)

```

initialize n to 1
sample(n) = Sample_Random_Sequence(n)
q(x) = Compute_RSSI_Distribution(sample(n))
do
n' = n + k * burst_size
sample(n') = sample(n) + Sample_Random_Sequence(k *
burst_size)
q'(x) = Compute_RSSI_Distribution(sample(n'))
if Compute_NKLD(q'(x)||q(x)) ≤ tolerance
return q(x)
update n = n', q(x) = q'(x) and continue

```

Figure 11: Algorithm to find length sequence n for which the RSSI distribution stabilizes

longer packet sequences. The online algorithm is shown in Figure 11. The enabling observation for the above algorithm is that after the NKLD curve stabilizes, increasing the length of packet sequence does not change the distribution substantially. So we compute the RSSI distribution for n and $n + burst_size$ for varying values of n and return the value for which both the distributions have relative distance less than the tolerance level. We use $burst_size$ as an increment, as RSSI varies significantly across bursts and we need to consider at least a gap of more than $burst_size$ to conclude that the RSSI distribution has stabilized. For our experiments we find that typically an increment of one $burst_size$ is sufficient to yield correct results using the online mechanism. Table 1 shows the values of n obtained for the four representative traces shown in figure 5. The value of n obtained using an online mechanism is close to the value obtained using offline analysis of the traces. In order to evaluate the efficacy of our online mechanism (to determine n) we compare the distribution obtained using a packet sequence of length n with the distribution obtained using large

traces ($\approx 100,000$). Figure 10 shows that the distribution obtained using n as determined by the online mechanism closely approximates the true distribution for all the traces.

Trace	Burst Size # of pkts	Offline		Online	
		# of pkts	NKLD	# of pkts	NKLD
LOS-light	≈ 1000	30,000	0.5	22,000	0.8
NLOS-light	≈ 2500	20,000	0.5	20,000	0.8
LOS-heavy	≈ 3000	16,000	0.5	9000	0.8
NLOS-heavy	≈ 400	3000	0.5	5000	0.05

Table 1: Minimum packet length sequence for capturing the distribution of RSSI, as calculated by offline and online mechanisms. Corresponding NKLD distance with the long term "true" distribution is also given. NKLD of 0.5 is chosen as the threshold for determining the packet length sequence in the offline mechanism. Burst sizes corresponding to first noticeable peak in Allan deviation is shown.

Validating Efficiency of Online-RSSI : We validate the efficiency of Online-RSSI by using the traces collected in our indoor WLAN deployment as described in Section 3.1. Using those traces, we first build an accurate estimate of the signal strength distribution for each scenario for different power levels. These distributions are computed over large traces (comprising of $\approx 100,000$ packets) and act as a baseline against which we compare the distribution generated by Online-RSSI. Figure 10 shows the accuracy of Online-RSSI for a given power level in different scenarios. The results for different power levels are similar in nature to the ones presented here. The base line distributions for different scenarios are shown in dotted lines and the real time distribution generated by Online-RSSI is shown in solid lines. As shown in the figure, Online-RSSI is able to accurately estimate signal strength distribution and the errors (NKLD distance between baseline and estimated) are found to be within **5%** for LOS-light, NLOS-light and NLOS-light, while for NLOS heavy it was found to be with **20 %**. This indicates that the algorithm has reasonable accuracy in estimating the RSSI distribution in an online fashion for different scenarios.

4. EMPIRICAL MODEL FOR POWER CONTROL

As discussed in Section 2.1, RSSI values of neighboring power levels tend to overlap significantly in indoor scenarios, with some indoor settings more prone to multipath effects (like cubicles) than others (like large conference halls). Similarly the interference and other factors that determine the extent of RSSI variations will be different for different indoor environments. Hence, it is possible that some indoor environments may allow more power levels to be distinguishable (where RSSI variations are low) as compared to others (where RSSI variation is high). Based on our online mechanism to dynamically determine the number of packets required to characterize RSSI distribution in any environment, we present an empirical model for transmit power control, *Model-TPC*, that outputs the set of feasible (non-overlapping distribution) power levels for a given indoor setting.

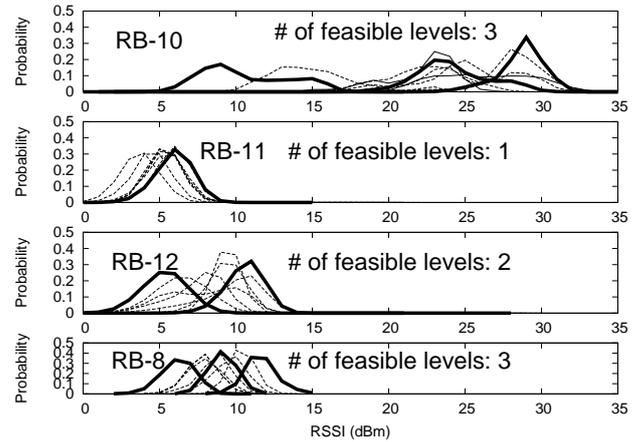


Figure 12: Probability distribution function for RSSI values received at varying power levels at the transmitter. The plots represent the distributions at receiver RB-10, RB-11, RB-12 and RB-8, in order from top to bottom. The exact positions of these receivers with respect to the transmitter can be seen in figure 2. The amount of overlap varies with the location and only 2-3 power levels are distinguishable at most of the receivers.

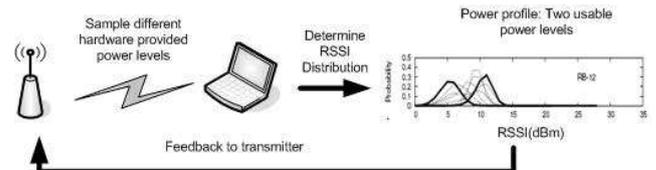


Figure 13: Steps involved in construction of Model-TPC. The receiver estimates the RSSI distribution using our Online-RSSI and computes set of feasible power levels as applicable to itself. This information is then sent to the transmitter to be used in power control

4.1 Model-TPC

Construction of our model proceeds through the following important steps, also shown in Figure 13. Assume we are operating in the context of a wireless node X.

- 1. Estimating RSSI distribution:** The RSSI distribution for any given power level is estimated using the Online-RSSI algorithm described in Section 11. Note that the RSSI distribution is captured at the receiver and communicated back to the sender as a feedback, as shown in Figure 13. Many proposed approaches (such as [2]) already incorporate protocol-level constructs to implement such functionality. Ongoing data communication between the participating nodes can be leveraged to gather this information. This process is repeated for different power levels available in the hardware. Note that for our experiments, this procedure is repeated for different hardware available power levels (6 for Cisco Aironet). In future, if the wireless hardware supports a large number of power levels, the

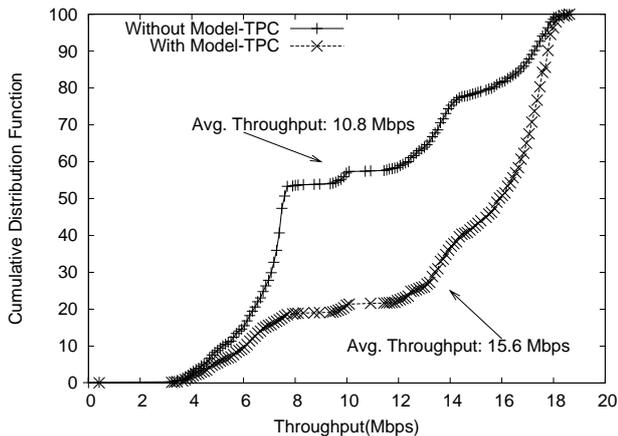


Figure 14: Cumulative distribution of throughput achieved by the wireless clients with/without the empirical model for adaptation at location T1. The average throughput for the adaptation process is also shown in the figure

cost for this step can be limited through a combination of sampling and simple approximation techniques to determine the RSSI distribution of power levels. We leave such extensions as directions for future work.

- Deciding the feasible power levels:** At completion of Step 1, the wireless node X would have built an empirically tuned model for the different power levels, much like Figure 12. At this point, if the NKLD of distributions of any two power levels is greater than a threshold $NKLD_{thresh}$, then the two power levels are considered to be distinct and can be used simultaneously. In theory, dynamic programming can be used to determine the largest set of feasible power levels satisfying above condition. For simplicity, we scan from maximum power level to lowest power level, picking all the power levels that satisfy the $NKLD_{thresh}$ criteria.

Figure 12 shows the distribution of RSSI for various receivers in our indoor WLAN deployment (Figure 2), when T1 is used as a transmitter and power level is varied at the granularity of 10mW. The power levels are not shown in the graph for the sake of clarity. The top most plot is for receiver RB-10, followed by RB-11, RB-12 and RB-8 in that particular order. We use the steps outlined above to determine the feasible power levels for the aforementioned receivers. The distributions corresponding to these feasible power levels are marked in black in Figure 12. As can be seen, the selected power levels overlap minimally ($NKLD \geq 4$). We also computed the error (captured by the NKLD function) between the accurate distributions and the distributions estimated by Online-RSSI. For each of these power levels, we found the error to be within 10 % of the desired maximum error. Clearly the amount of overlap (and hence the number of distinguishable power levels) depends on the location of the receiver, with RB-10 observing less overlap as compared to RB-11, which practically observes only a single power level. These results clearly indicate that the set of feasible power levels is highly correlated with the location of the receiver and motivates the case for location-based power

control, where the transmitter uses a different set of power levels for each client depending on client's location.

4.2 Summary

For a given a wireless environment, our proposed model and its associated algorithms were able to accurately determine a good and useful set of power levels. The set of useful power levels as computed by Model-TPC are valid till traffic characteristics (other interference source) and wireless environments (physical obstacles etc) remain similar. Using our Online-RSSI algorithm, we already sample sufficient packets to reflect small scale changes in the wireless environments in our model. However the set of power levels must be re-computed against large scale changes in the wireless wireless environment like transmitter mobility, introduction of a new physical obstacle or a new interference source. We are investigating various triggering mechanisms to refresh the Model-TPC, although a simple strategy to refresh the model every 10 minutes seems to work fine for our indoor experiments.

5. EXPERIMENTAL EVALUATION OF MODEL-TPC

To validate our model, we pick an existing algorithm [15] that uses transmit power control for improving client throughput and spatial re-use. The algorithm proposed increases transmit power in steps and measures signal quality to ascertain the optimal power setting for a given client.

At a high level, the algorithm operates as follows. It starts with the lowest power level and performs normal data rate adaptation using Onoe [1](a standard data rate adaptation mechanism). Once the data rate stabilizes around a value, the power level is increased and the rate adaptation process is continued. This process is repeated until the transmitter reach the maximum rate available or reaches the highest power level.

To demonstrate the benefits of our proposed model, we create a set of useful power levels through Model-TPC and restrict the above algorithm to use only this set of power levels in its adaptations. We then compare the adaptation performance of the algorithm under two different scenarios – (i) which uses all possible power levels as available from the wireless interface, and does not use our model-TPC, and (ii) which uses the power levels provided by Model-TPC.

There are two benefits of Model-TPC: First, it allows for significantly faster convergence for the transmitters to the best suited power level in their operating environments. Second, by eliminating the need to explore many redundant power levels with corresponding poor throughput performance, the transmitters achieve higher throughput over the entire adaptation duration. This is particularly important for clients that are mobile in nature and hence, need to adapt their transmission parameters, including power levels, quite frequently. We illustrate these gains through our reference implementation of the algorithm in [15], both with and without Model-TPC.

5.1 Setup

For the experiment described, the setup is identical to NLOS scenario, with the transmitter using an Atheros card having five power levels as validated by our power level validation setup in Figure 4. The mobile client continuously

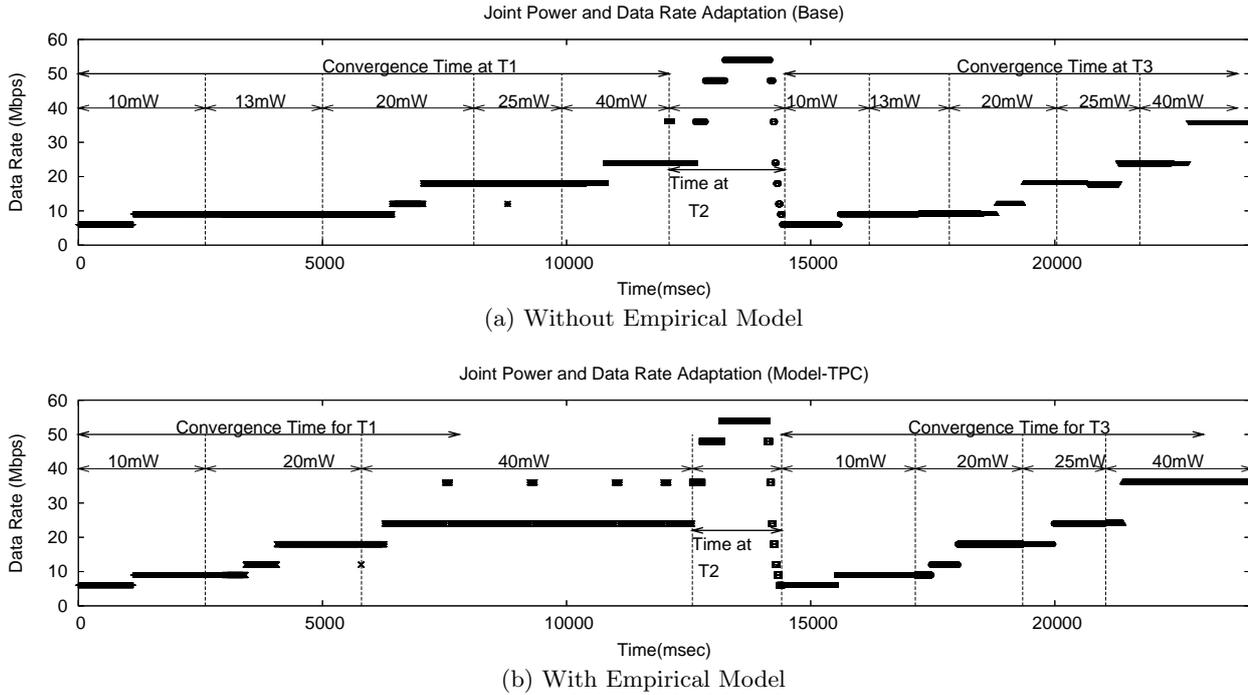


Figure 15: Joint power and data rate adaptation mechanism with/without the empirical model. Convergence is much faster with the empirical model.

transmits data from itself to a departmental server located at the position of receiver R2, shown in Figure 2. The client roams from locations T1 to T2 to T3, which are annotated in the Figure 2 of our indoor WLAN deployment. Initially the client is at T1, which has 3 feasible power levels of 10mW, 20mW and 40mW, as per Model-TPC. After 12 seconds, the client goes to location T2, which is very close (LOS) to the server R2 and hence the client decreases its power level and is able to use the lowest power level of 10mW to achieve a data rate of 54Mbps. After 2 seconds, the client again moves to location T3, which has four feasible power levels as per our empirical model. We show the data rate and power adaptation process at T1 and T3 (The adaptation at T2 is obvious, with the client simply reducing power levels as it is very close to the server).

5.2 Results

We present the cumulative distribution function of the instantaneous throughput (measured every 100 ms) of the two variants of the transmit power control algorithm in Figure 14. The figure shows that using Model-TPC to restrict power levels lead to higher instantaneous throughput for a significant part of the experiment as shown in Figure 16. We explain this difference by examining the adaptation mechanisms in the two cases in Figure 15.

Figure 15(a) shows the adaptation behavior when all five power levels are used by the algorithm. We can see that over time, the algorithm attempts to identify signal quality at each different data rate and power level, spending a significant amount of time testing parameter values which are redundant for a give environment, thus impacting performance. In contrast, Figure 15(b) shows that adaptation with our Model-TPC. Clearly adaptation is much faster with

our model, with more pronounced gains at T1 (as difference between hardware and feasible power levels is more) than T3.

Note that here we only show the throughput gains arising from quicker convergence from a small power level to the right (greater) power level for locations T1 and T3. Model-TPC also provides much better convergence when adapting from a high power level to lower (right) power level as for T2, by skipping all the redundant high power levels in between. A faster convergence reduces the energy consumed in scanning high power levels and leads to energy savings, which is an important consideration for mobile clients. Due to lack of space, we do not present our energy results in this section.

5.3 Summary

Our gains of in the above Internet-oriented wireless experiments stem from faster adaptation achievable when using the Model-TPC as an input to power control. Note that in our experiments, we compared benefits when only five power levels are available from the wireless interface. The performance gains of Model-TPC will only be greater if the wireless interface makes more power levels available to the system software, that will clearly increase the number of redundant channels that transmitter will scan in a typical power control algorithm, while our model will facilitate much faster convergence and performance.

6. DISCUSSION

While our work in this paper is targeted towards indoor WLANs, we discuss the relevance of our work in context of cellular networks, where power control is again an impor-

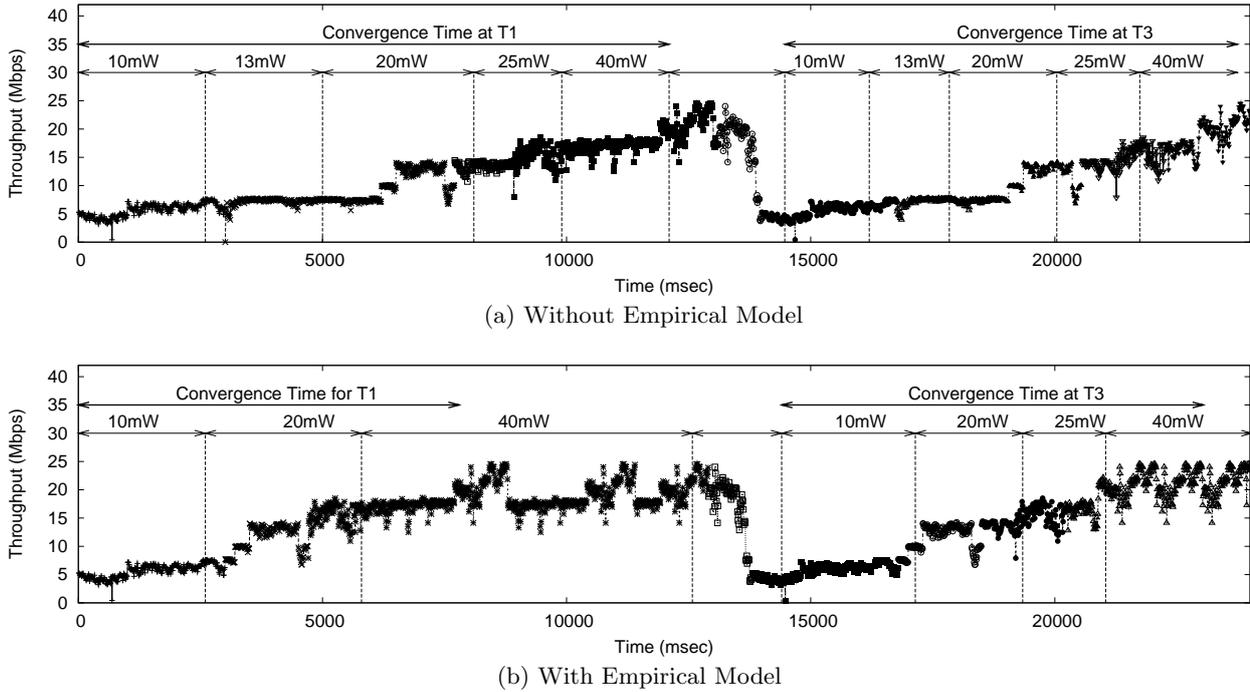


Figure 16: Goodput of the end wireless clients in joint power and data rate adaptation mechanism with/without the empirical model.

tant design parameter. Power control in cellular networks is used for reducing co-channel interference, managing voice quality, dealing with fast fading and near-far problem [10, 17]. However cellular networks primarily operate in outdoor environments, where the effects of multipath are much less pronounced as shown in Figure 3. Moreover, cellular networks do not perform rate adaptation in the inner loop (real time or per packet basis) of power control, whereas data rate adaptation is an integral component of 802.11 based WLAN systems. Thus the SNR threshold for cellular networks is varied slowly in the outer loop of power control, whereas in WLANs, data rate adaptation is performed on very small time scales, thereby making RSSI variations even more critical for system performance.

7. RELATED WORK

We discussed prior work on power control mechanisms in Section 2. In this section, we present other previous work that deals with signal strength measurement and characterization for wireless networks.

- **Measurement based modeling :** Some recent efforts have been made to use empirical observations to improve wireless protocols. Reis et al.[7] propose a measurement based model for delivery and interference in static wireless networks. Their work takes RSSI values of wireless packets as an input to predict the delivery rate and interference in the system. Divert [4] attempts to reduce packet loss rates in WLAN systems by rapidly switching between APs to tolerate bursty losses. ExOR [5] 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 such tool.

- **RF-based location determination :** RF-based location determination mechanisms [21, 14] use signal strength values for fingerprinting different locations in a WLAN. Kaemarungsi et al. [12] study the properties of indoor received signal strength for location fingerprinting. They propose an analytical model for indoor positioning system by modeling the RSSI variations as a Gaussian distribution. While such work in location determination had to examine RSSI (and hence, power) variations between a transmitter and a receiver, the focus of such RF-based localization technique did not require a careful exploration of various power level choices, and their implications on power control mechanisms.
- **Feasibility Analysis :** Abdesslem et al [8] describe the hardware and software limitations, like limited power levels in the wireless chipsets and lack of suitable device drivers, that hinder the implementation of transmit power control mechanisms. While their work is an important step towards determining feasibility of power control mechanisms (due to hardware/software limitations), they do not explore the more fundamental issues with fine grained power control that arise due to the inherent nature of wireless medium.

8. CONCLUSIONS AND FUTURE WORK

Multipath, fading, shadowing and external interference from wireless devices, make the implantation of power con-

control mechanism challenging in practical settings. The focus of this paper has been in understanding what the right set of power control mechanisms are useful to design efficient power control algorithms. More specifically, we show that fine-grained power control cannot be effectively used by such algorithms in a systematic manner. In fact, our work suggests that a *few 3-5* discrete power level choices is sufficient to implement any robust power control mechanism in typical indoor WLAN environments. Through our work, we also build an empirical model that guides these appropriate number and choices of power values that is adequate. Our model can be used as a plug-in to previously proposed power control mechanisms, to make them implementable in real settings. We believe our work provides an important framework that can be used by researchers to develop robust and practical power control mechanisms.

We have used NKLD as a statistical tool to measure the distance between two RSSI distributions. Although it works well for our environments and is easy to compute in a real time fashion, there are other statistical tools like moment based estimators, that capture the spread of the two distributions better and may be more effective in distinguishing between two probability distributions. Comparing the performance of NKLD with moment based estimators is an avenue of future work for us.

We are also investigating various triggering mechanisms to refresh Model-TPC. Currently we refresh it periodically every 10 minutes, which not be optimal for every scenario. The refresh period is tightly coupled with large scale variations in the wireless environments, like addition of an interfering source or a new obstacle and we are performing more experiments to detect such changes that can trigger the update of Model-TPC.

Although our main objective is to evaluate the feasibility of using fine grained power control in indoor environments, the model developed in this section can be readily used by access points (APs) to profile various locations in the environment and perform location-based power control. Through collaborative measurements made by different clients over time, the APs can create a location-dependent model for power control which can be downloaded to clients during association. Also the network as a whole could aggregate such models to create a single model which is network dependent but location independent that trades-off complexity and accuracy for simplicity of use. We are currently looking such location-specific power control approaches that can be implemented using a centralized controller that manages the wireless APs in enterprise WLANs. The flexibility in estimating and building the empirical model allows for its applicability to a wide range of power control algorithms and might find interest outside the scope of 802.11 networks.

Acknowledgment: We would like to thank all the reviewers for their comments. Special thanks to our shepherd Dr. Darryl Veitch, whose guidance bought this paper together into its final form. Vivek Shrivastava, Arunesh Mishra, and Suman Banerjee were supported in part by NSF awards CNS-0639434, CNS-0627589, CNS-0627102, and CNS-0520152.

9. REFERENCES

- [1] ONOE rate control.
http://madwifi.org/browser/trubk/ath_rate/onoe.
- [2] P. Steenkiste A. Akella, G. Judd and S. Seshan. Self management in chaotic wireless deployments. In *ACM MobiCom 2005*.
- [3] D.W. Allan. Time and frequency (time domain) characterization, estimation and prediction of precision clocks and oscillators. In *IEEE Transactions Nov. '87*.
- [4] Hari Balakrishnan Allen Miu, Godfrey Tan and John Apostolopoulos. Divert: Fine-grained path selection for wireless lans. In *ACM Mobisys '04*.
- [5] Sanjit Biswas and Robert Morris. ExOR: opportunistic multi-hop routing for wireless networks. In *SIGCOMM 2005*.
- [6] T.M. Cover and J.A. Thomas. *Elements of Information Theory*, 1991.
- [7] Charles Reis et al. Measurement-based models of delivery and interference in static wireless networks. In *SIGCOMM '06*.
- [8] Fehmi Ben Abdesslem et al. On the feasibility of power control in current IEEE 802.11 devices. In *PERCOMW 2006*.
- [9] P. Gupta and P. Kumar. Capacity of wireless networks. In *IEEE Transactions on Information Theory*.
- [10] Prashanth Hande, Sundeep Rangan, and Mung Chiang. Distributed uplink power control for optimal SIR assignment in cellular data networks. In *INFOCOM*, 2006.
- [11] V. Bharghavan J. Monks and W. Hwu. A power controlled multiple access protocol for wireless packet networks. In *INFOCOM '01*.
- [12] Kamol Kaemarungsi and Prashant Krishnamurthy. Modeling of Indoor Positioning Systems Based on Location Fingerprinting. In *IEEE INFOCOM '04*.
- [13] S. Kullback. *Information theory and statistics*, 1959.
- [14] J. Lansford and P. Bahl. The Design and Implementation of Home-RF: A Radio-Frequency Wireless Networking Standard for the Connected Home. In *IEEE '00: 1662-1676*.
- [15] K. Leung and L. Wang. Controlling QoS by Integrated Power Control and Link Adaptation in Broadband Wireless Networks.
- [16] Allen K. Miu, Hari Balakrishnan, and Can E. Koksal. Improving Loss Resilience with Multi-Radio Diversity in Wireless Networks. In *MOBICOM '05*.
- [17] Zvi Rosberg. Asymptotically optimal transmission power and rate control for CDMA channels with multiple user classes. In *INFOCOM*, 2005.
- [18] Anmol Sheth and Richard Han. SHUSH : Reactive Transmit Power Control for Wireless MAC Protocols. In *WICON '05*.
- [19] M. Subbarao. Dynamic power-conscious routing for manets: An initial approach. In *IEEE Vehicular Technology Conference '99*.
- [20] Chi-Hsiang Yeh. IPMA: An interference/power-aware mac scheme for heterogeneous wireless networks. *ISCC*, 2003.
- [21] M. Youssef and A. Agrawala. The Horus WLAN Location Determination System. In *Mobisys '05*.