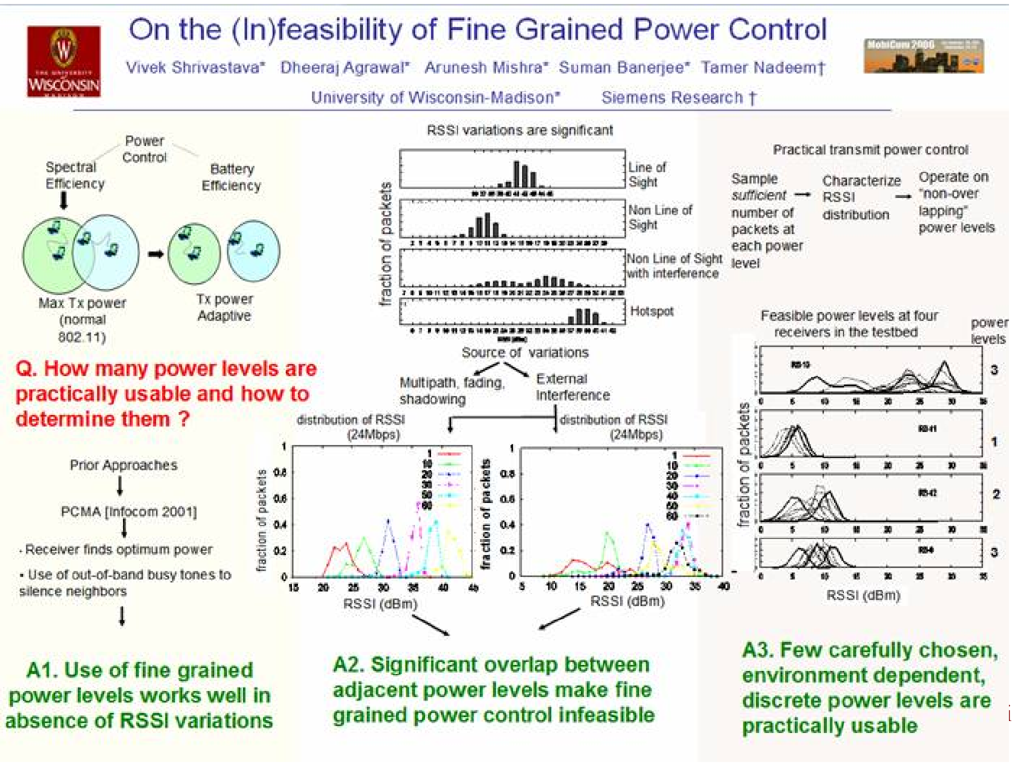

On the (In)feasibility of Fine Grained Power Control

Vivek Shrivastava* Dheeraj Agrawal* Arunesh Mishra* Suman Banerjee* Tamer Nadeem†

University of Wisconsin-Madison* Siemens Research †

Q. How many power levels are practically usable and how to determine them ?

Prior Approaches → PCMA [Infocom 2001]

- Receiver finds optimum power
- Use of out-of-band busy tones to silence neighbors

A1. Use of fine grained power levels works well in absence of RSSI variations

RSSI variations are significant

Source of variations: Multipath, fading, shadowing; External Interference

A2. Significant overlap between adjacent power levels make fine grained power control infeasible

Practical transmit power control

Sample *sufficient* number of packets at each power level → Characterize RSSI distribution → Operate on "non-over lapping" power levels

Feasible power levels at four receivers in the testbed

A3. Few carefully chosen, environment dependent, discrete power levels are practically usable

Linl interested
area

I. 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 work, our focus is to study practical power control as applicable to 802.11 based networks. For interference mitigation, power control mechanisms like PCMA[3] 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. On the other hand, energy conservation mechanisms try to minimize the transmission power to support the required quality of service.

A number of prior research efforts have studied power control based on the theoretical abstraction of wireless signal propagation in free space. They generally consider transmit power as a continuous variable (i.e., a fine grained parameter), that can be set on a per packet basis. However lack of vendor support acts as a deterrent to a practical realization of such prior mechanisms [1]. Hence many few power control mechanisms have made it to practice. In this work we ask the following questions: ***Is fine-grained power control really useful and would lead to a better realization of power-control algorithms? If not, what is the minimum granularity of power control that is useful in different wireless environments?*** . We answer the first question in the negative. As we discuss in detail, 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 power control mechanisms. We demonstrate this through detailed experimentation in different indoor wireless network environments. Our answer to the second question is that a power control algorithm can make practical use of only a **few** 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. In particular, through this work we build an empirical model that allows us to characterize the specific set of power levels that is useful for a given environment

"Is fine-grained power control really useful? If not, what granularity of power control is useful in practical settings ?"

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 hotspots to characterize **Receive Signal Strength Indicator (RSSI)** variations in different practical environments. We observe that the number of feasible power levels that can be plugged in a transmit power control mechanism is few and discrete and is a function of the environment.

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. We believe that our experiments highlight some fundamental issues with transmit power control, that can help in design of future wireless cards.

The remainder of the article is organized as follows. Section II discusses various transmit power mechanism proposed in literature and their respective evaluation in context of our practical models for transmit power control. Section III describes our experimental methodology and section IV explains the fundamental limitations of fine-grained power control. In section V, we provide basic guidelines for developing a practical transmit power control mechanism. We discuss some applications in section VI and conclude in section VII.

Next: Experimental Methodology Previous: Introduction Index

II. Motivation : Power Control Approaches and Limitations

Through controlled experiments, we first question the usefulness of fine-grained transmit power control in practical wireless environments. We show that significant variations in the received signal strength indicator (RSSI) render such control practically ineffective. Next, we summarize prior approaches proposed in the literature that take advantage of such 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.

Infeasibility of Fine Grained Power Control

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 shown in [1], most of the wireless cards support only 4 to 5 power levels at the hardware, which is in stark contrast to the fine grained power control preferred by most power control schemes. 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 realistic wireless environments, which limits the number of feasible power levels that is useful in such mechanisms. We present preliminary results from our detailed set of experiments explained later to illustrate the fundamental constraints of fine-grained power control.

“There are fundamental constraints to power control in realistic environments, which limits the number of feasible power levels”

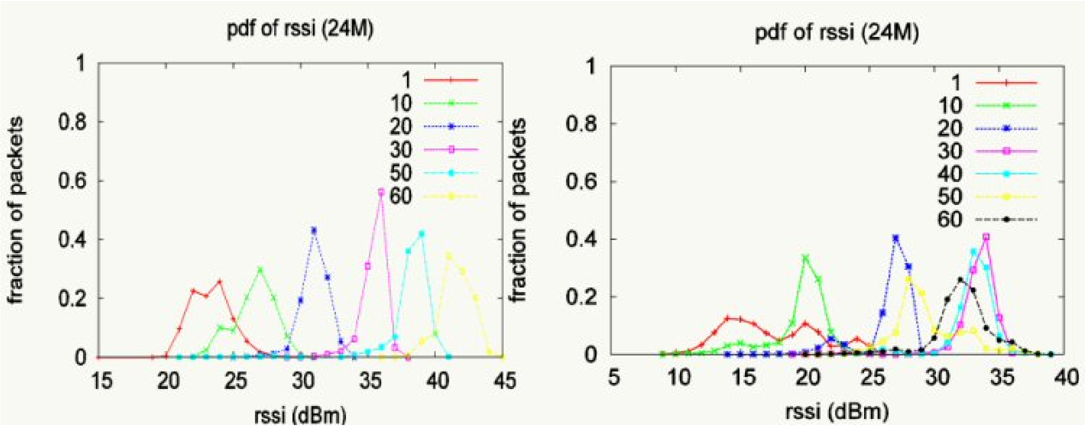

Figure 1: Probability distribution of RSSI for varying power levels is shown. The left plot corresponds to no interference with 6 distinguishable power levels while right plot shows the interference effects on the same scenario with the number of power levels reduced from 6 to 3. Transmit powers are varied from 1mW to 60mW in step of 10mW.

This sample experiment consists of a transmitter-receiver pair (T2-R2) shown in figure 5 operating in 802.11g. At R2 we capture the packets transmitted by T2 for different power levels available at T2's atheros based wireless chipset. Figure 1 shows the probability density function (PDF) of RSSI distribution for various power levels at the transmitter. The power levels are increased from 0dBm (1mW) to 18 dBm (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 60mW, 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 overlap 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. 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 1 (confirmed through extensive empirical analysis in later sections).

“Significant overlap between the RSSI distributions of two power levels make them practically indistinguishable at the receiver”

No interference scenario: It is evident from figure 1 (left plot) 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 (1,10,20,30,50,60)mW 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.

Effects of interference: The above experiments were conducted in an environment devoid of any interference from other wireless transmitters. We repeat the experiments for transmitter-receiver pair T2-R2 (figure 5 - floor plan) with external sources of interference from 802.11 devices operating on the same channel. The resulting distribution of RSSI values is shown in figure 1(right plot). As expected the RSSI variations increase, thereby increasing the overlap between RSSI of neighboring power levels. This observation indicates that in the presence of interference, the number of power levels having non-overlapping RSSI distributions are further reduced, thereby making fine-grained transmit power control even less effective. These experiments further reinforce our claim that transmit power control mechanism may not be feasible in practical deployments where external interference is substantial.

Effect on Proposed Power Control Approaches We categorize some of the prior power control methods applicable to WLANs into two : fine-grained power control mechanisms and per-packet power control mechanisms.

Fine Grained Power Control: Monks et al. proposed a power controlled medium access control (PCMA [3]), which optimizes the `floor space' used by the transmitter-receiver pair by reducing transmit power. The mechanism allows the flows to be more tightly packed and provides better spatial reuse of spectrum resources as compared to naïve 802.11. Their mechanism avoids collisions by sending ``busy tone signal" on an auxiliary channel, the transmit power of the tone determining the floor space acquired by the flow. The authors treat transmit power as a continuous parameter and do not consider the effect of RSSI variations on the mechanism.

“Use of fine grained power levels works well in the absence of RSSI variations ”

Packet-based Power Control: Akella et al. [4] discuss some power control mechanisms in their paper 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. Using our empirically derived power control model, the above mechanisms could dynamically determine an exact set of discrete power values to be used.

Cellular Networks: In cellular networks, a centralized base station just needs to maintain a predetermined QoS level (low data rate for voice) at the clients. Hence there is not much incentive in performing very fine grained power control. However, in WLANs, where data rates can take a wide range of values, varying power at very fine granularity may provide good gains by achieving higher data rates [5].

Next: Characterizing RSSI Distributions **Previous:** Motivation Index

III Experimental Methodology

In order to characterize RSSI variations and its effect on power control schemes we collected extensive traces in different wireless environments. This serves 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

Testbed Setup

The primary vehicle for our empirical studies is an indoor tested of fourteen stationary router boards shown in Figure 5. The testbed is located on the third floor of an enterprise, mimicking a deployment of wireless nodes in an office scenario. Each testbed node has an (ar-5212) 802.11 a/b/g Atheros card that we operate using the Madwifi driver. In all our experiments, we use a fixed data rate of 1Mbps and fixed packet size of 1KB. For our experiments 1 sec of receiver time window ~125 packets. We repeated the same experiments with other wireless cards and found the results were consistent with the ones reported here.

RSSI measurements

We use received signal strength indicator (RSSI) values that are reported by all commodity wireless cards. RSSIs 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 10log(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 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.

Validating Hardware Supported Power Levels We use the current drawn by the wireless card circuit to determine the number of power levels the card supports. This is based on the fact that the peak current drawn by a wireless card circuit (set up shown in leftmost diagram of Figure 2) increases with the transmit power used. It is a more accurate measure of determining the transmit power levels used by a wireless card as compared to inferring it using the RSSI at the receiver.

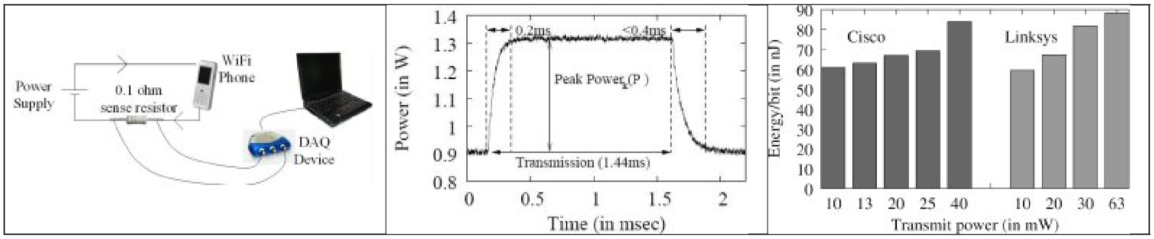

Figure 2: Left figure shows setup for measuring current. Middle figure shows the current drawn using our setup for a packet. Right figure shows energy consumed per bit of data transferred for some power levels in Cisco & Linksys cards.

Next: Empirical Model **Previous:** Experimental Methodology Index

IV Characterizing RSSI distribution

In this section, we characterize the distribution of RSSI in the various scenarios mentioned above. 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.

On the basis of our experiments, we observed that due to multipath, fading and other propagation effects, the RSSI values at the receiver show significant variation (also corroborated by Figure 1). Depending on the exact environment, two RSSI distributions for successive values of transmit power can have substantial overlap, making them practically indistinguishable at the receiver. Therefore, the number of distinguishable power levels vary depending on the RSSI distribution of the individual power levels. 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 ~10 minutes 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 ?*

"What is the minimum number of packets we should sample to get a ``good" approximation of RSSI distribution ? "

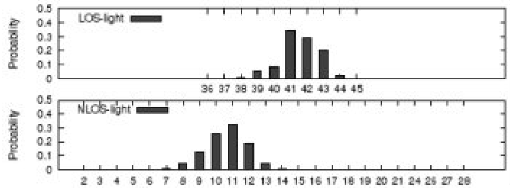

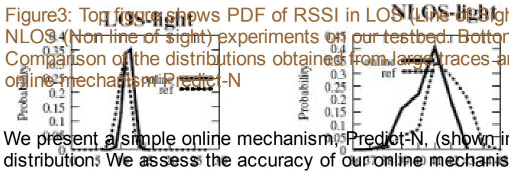

Figure3: Top figure shows PDF of RSSI in LOS (Line-of-Sight) and NLOS (Non line of sight) experiments on our testbed. Bottom Comparison of the distributions obtained from large traces and online mechanism Predict-N

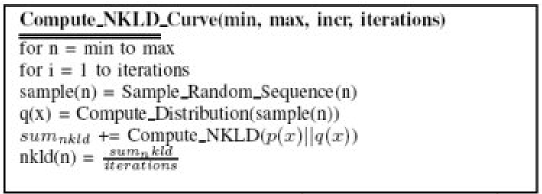

```
Compute_NKLD_Curve(min, max, incr, iterations)
for n = min to max
    for i = 1 to iterations
        sample(n) = Sample_Random_Sequence(n)
        q(x) = Compute_Distribution(sample(n))
        $sum_{nkld}$ += Compute_NKLD($p(x)||q(x)$)
    nkld(n) = $\frac{sum_n kld}{iterations}$
```

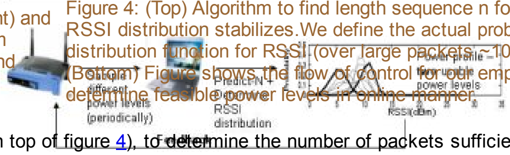

Figure 4: (Top) Algorithm to find length sequence n for which the RSSI distribution stabilizes.We define the actual probability distribution function for RSSI (over large packets ~100,000) as p(x). (Bottom) Figure shows the flow of control for our empirical model to determine feasible power levels in online manner.

We present a simple online mechanism, Predict-N, (shown in top of figure 4), to determine the number of packets sufficient to get a good distribution. We assess the accuracy of our online mechanism for two scenarios: (1) LOS (Line of Sight) - Transmitter and receiver are in line of sight, like T2 and R2 for our testbed. (2) NLOS (Non Line of Sight) - Transmitter is T1 and receivers can be RB-1 to RB-12. RSSI variations for the two scenarios is shown in top part of figure 3. The approximate distribution obtained by our online mechanism, Predict-N is shown in figure 4. We use NKLD as a measure for overlap between two distribution functions. The detailed analysis can be found in [2]. The accuracy of our mechanism is shown in figure 3 (bottom) for different scenarios.

Next: Applications Previous: Characterizing RSSI Distribution Index

V Towards an Empirical Model for Power Control

Based on the insights gained from studying the RSSI distributions for various environments, we present an empirical model for power control. Figure 4 (bottom) shows the flow control for our empirical model, where the AP transmits at some sample power levels (as advertised by hardware vendor) and the client determines the set of feasible power levels, which is provided as a feedback to AP, that may use it for performing power control.

Some indoor environments may be 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). Figure 5 shows the distribution of RSSI for various receivers in the test bed, when T1 is used as a transmitter and power level is varied at the granularity of 2 dBm. 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. 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.

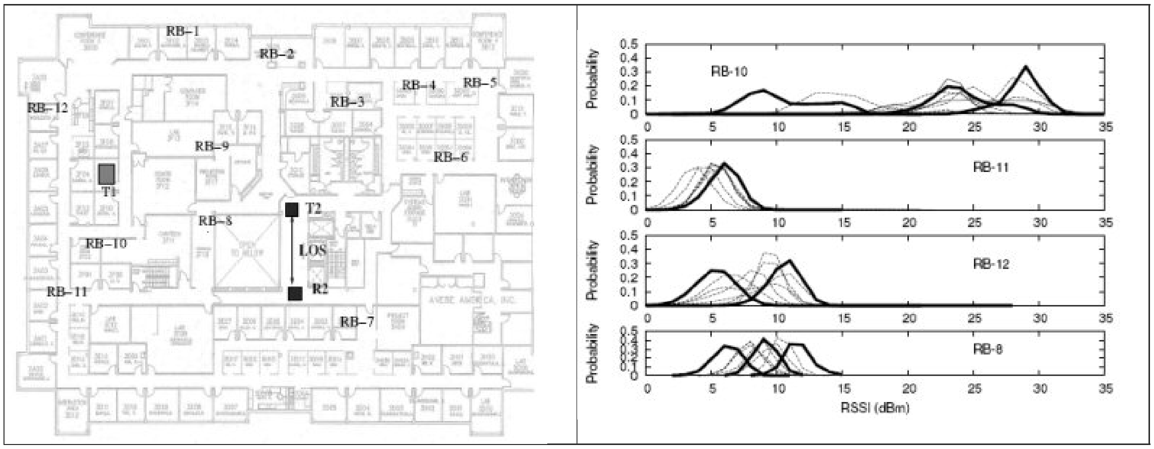

Figure 5: The left figure shows the wireless testbed, consisting of seven 802.11 a/b/g nodes (transmitters marked by T1, T2 and receivers marked by RB-1 to RB-12. Right figure shows the probability distribution of RSSI for varying power levels at the transmitter for receivers RB-10, RB-11, RB-12 and RB-8 (Top to bottom).. Number of feasible power levels vary with location of the receiver.

Next: Conclusions Previous: Empirical Index

VI Applications

We present two scenarios where information about the usable power levels can provide good gains to existing power control mechanisms.

Adding a profiling plug-in to current power control approaches

As an example application, we have augmented an existing power control protocol, PCMA, to use the right set of power levels. Our model shown in Figure 5, computes the right set of power levels for each receiver, thereby eliminating the need for PCMA to scan a large number of potentially available power levels (which is again difficult to infer). Our initial results demonstrate rapid convergence for PCMA by operating on few feasible power levels.

Energy Savings using empirical model

If two power levels are indistinguishable at the receiver, it might be advantageous from an energy perspective to use the lower power level. Figure 2 shows that the energy consumed (in nJ per bit) by the linksys and Cisco cards vary substantially with power levels, with 40% more energy consumed if transmit power of 63mW then 10mW. For e.g. in Figure 1 (right figure), there is significant overlap between

"Reduces search space for power levels in practical settings and can provide solid energy savings if used in conjunction with existing battery conserving mechanisms."

"Very few carefully chosen, environment

30mW and 40mW transmit power. Therefore, the receiver should use 30mW instead of 40mW. So, it is possible to leverage our model to intelligently select power levels for energy saving mechanisms as well.

dependent power levels are actually usable”

Next: References Previous: Applications Index

VII Conclusions

Multipath, fading, shadowing and external interference from wireless devices, make the implementation of power 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 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 guides these appropriate number and choices of power values that is adequate. We believe our work provides an important framework that can be used by researchers to develop robust and practical power control mechanisms.

Previous: Conclusions Index

VIII References

[1] Abdesslem et al. On the feasibility of power control in current IEEE 802.11 devices. In PERCOMW ’06
[2] Shrivastava et al. Are fine grained power levels really useful? Tech report, UW-Madison, www.cs.wisc.edu/~viveks/fine-grained.pdf
[3] V. Bhargavan, J.Monks and W.Hwu. A power controlled multiple access protocol for wireless packet networks,. In INFOCOM ’01
[4] A. Akella, G. Judd, S.Seshan and P.Steenkiste. Self Management of chaotic wireless networks,. In Mobicom ’05
[5] Price Performance: 3G and Tropos MetroMesh, Technology Whitepaper, Tropos Networks, April 2004