On Estimating Joint Interference for Concurrent Packet Transmissions in Low Power Wireless Networks

Ritesh Maheshwari, Shweta Jain, Samir R. Das
Department of Computer Science, Stony Brook University, Stony Brook, NY 11794-4400, USA
ritesh,shweta,samir@cs.sunysb.edu

ABSTRACT

In a wireless network it is important to understand the nature of the joint interference generated at a receiver by multiple concurrent transmitters. This understanding helps developing packet scheduling algorithms. Prior experimental work using older generation mote-class radios (CC1000) have showed systematic deviations between estimation and direct measurement of the joint interference power, thus questioning whether the standard assumption that received signal powers are additive is applicable in practice. We, however, show via extensive experimentation that on newer generation radios (CC2420), the additive assumption is valid, particularly at the low power end.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Wireless communication

General Terms

Measurement.

Keywords

Interference, CC2420, Joint RSS, Additive RSS, SINR.

1. INTRODUCTION

There is a growing interest in understanding and modeling interference in wireless networks. This helps in developing efficient protocols for MAC-layer packet scheduling. Clearly, the packets must be scheduled in an interference-free fashion. While many algorithmic studies have used simple interference models based on network topology or physical node distance, focus has recently shifted towards more realistic models, based on *signal-to-interference-plus-noise ratio (SINR)*. Essentially, in SINR model, success of a packet reception depends on the ratio of the signal power on one hand and the aggregated interference and noise power on the other. The SINR model is driven by the fundamentals of radio receiver behavior at the physical layer.

While realistic, the SINR model is no longer 'pairwise' as the other simpler models, where the interference is modeled between the link in question and only one other interfering link. In SINR model, interference is considered in an aggregated form – *as the sum of all interference powers from all interfering links*. This subtlety not only makes algorithm design complex, but also makes instantiation of the model harder. Prior research [8] has done significant empirical work in understanding the interference properties using the SINR model for low power wireless links using the first generation motes (Mica2 motes with CC1000 radio [11]). However, they reported several interesting observations that cannot be easily supported by radio fundamentals. In particular, they reported that aggregated or joint signal power cannot be modeled by simply summing the individual signal powers. This obviously makes modeling aggregated interference much harder.

In our work, we perform a careful set of measurements on a later generation mote (TelosB [1] with CC2420 radio [10]) and report our findings. Our goal is to verify the classical additive model of signal powers on newer generation hardware. We find, in contrast to the observations in [8], additive model works quite well in practice barring measurement noises. To make this observation, we repeat similar experiments in [8], albeit on the newer mote platform, and then back up the observation with more elaborate experiments and measurements so that a fairly robust conclusion is possible. For brevity, we particularly emphasize on the low-power end. Much of our paper describes the details of the experimental setup, methodology and observations.

Our experience in this work points out that specific observations about radio behaviors in wireless networks are quite platform specific. Later generation hardware is likely to have lesser imperfections and are prone to lesser measurement errors and are thus likely to track theory more closely. The research community must do serious validation work in various platforms before embracing or questioning any model.

2. EXPERIMENTAL PLATFORM AND METHODOLOGY

Our experimental testbed consists of TelosB motes [6] using a CC2420 radio [10] which is compliant with the IEEE 802.15.4 PHY layer standard [4] in the 2.4 GHz ISM band and operates at the nominal bit rate of 250 Kbits/s. The radio provides some flexibility in terms of choice of transmit

WinTech'08, September 19, 2008, San Francisco, California, USA.

power that has been quite useful in our work. A custom MAC layer (described below) is implemented to enable concurrent packet transmissions without any carrier sensing. Concurrent packet transmissions are important as our goal is do power measurements when multiple concurrent transmitters are active.

Since we are interested in power measurements, unintended variations in transmit power cause significant modeling errors. We have noticed such variations when motes are driven by battery. Thus, we have always powered the motes using their USB interfaces to wall sockets via USB hubs.

2.1 Received and Transmit Powers

The CC2420 radio provides a measure of the received signal strength (RSS) in dBm, which is an estimate of signal strength averaged over 32 bit periods ($128\mu s$) and is continuously updated. This value can be either read directly from the RSS register or obtained from the metadata in the received packet. Since packet reception is not always possible for weak signals, we read the RSS from the register periodically to obtain signal strength even when the packet is not received correctly. The CC2420 datasheet [10] specifies that the transmit power can be programmed at 8 discrete levels between -25 to 0 dBm by setting the $TXCTRL.PA\_LEVEL$ register to values from 3 to 31 in steps of 4. But we verified experimentally that the power levels can be varied at a finer scale by setting the $TXCTRL.PA\_LEVEL$ register to values from 1 to 31 in steps of 1.[1] Transmit power is directly proportional to square of the signal amplitude. Assuming that the $TXCTRL.PA\_LEVEL$ register values and the signal amplitude are linearly related, the resulting extrapolated transmit powers vary from -32.5 to 0 dBm. This gives us a choice of a wide range of transmit powers.

2.2 MAC Layer and Measurement Process

Since concurrent packet scheduling is of essence, we have implemented a simple TDMA protocol in TinyOS-2.0 [2], where motes transmit at designated time instants without performing carrier sensing or backoff as in the default MAC implementation in TinyOS. We achieve time synchronization between nodes in the testbed as follows. One mote outside the testbed is directly connected to a laptop via USB. This mote and laptop combination is loosely referred to as the 'base station'. All network motes can directly talk to the base station using the maximum transmit power (0 dBm). This is the power the base station also uses. The base station periodically (500 ms intervals) transmits 'beacons' that the motes use to synchronize their clocks.

We have also implemented a 32 KHz precision timer to achieve low jitter between the actual and the scheduled transmission start times across motes. We have experimentally observed that the maximum jitter in transmission start times in our steup is less than the duration of the start frame delimiter (SFD) which is 128 μs. This guarantees almost concurrent transmissions.

The 'base station' (BS) also acts as a command and control center for the network for the measurement process. Any measurement activity in the testbed is initiated by broadcast 'command' message(s) from the BS. The command message contains specific instructions for each node and the nodes then start the necessary 'activity' (RSS measurements, packet transmissions etc., possibly at the scheduled time instants as indicated in the command). Similarly, when the 'activity' is over (the period of activity is predetermined), the BS mote sends 'poll' messages to motes to collect measurement data one at a time. Similar technique also has been in used in [8] for evaluations.

Care is taken so that all measurements are done within the

[1]This undocumented feature was confirmed by the mote manufacturer [1]

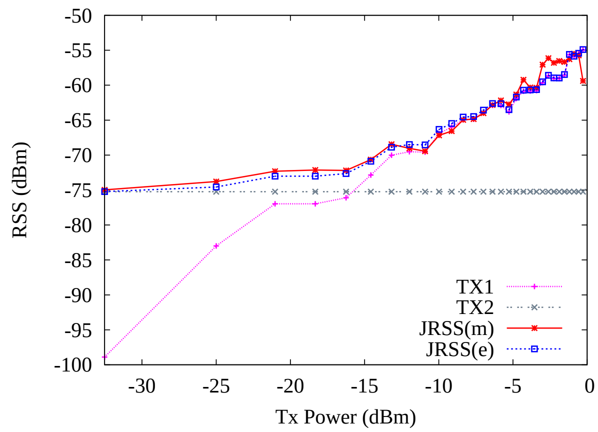

(a) Low transmit power for TX2.

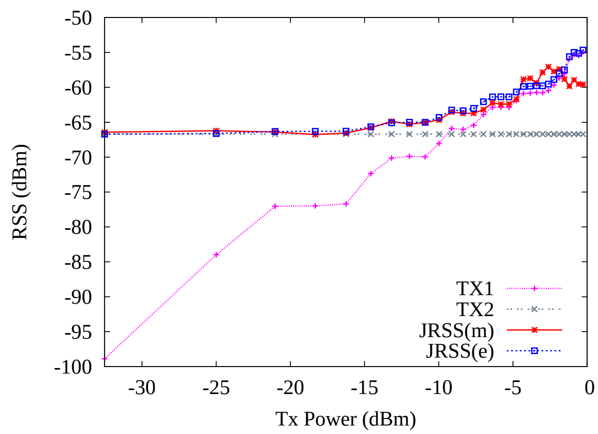

(b) Medium transmit power for TX2.

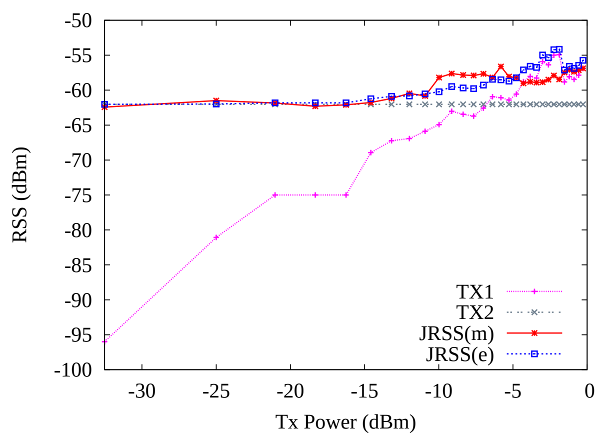

(c) High transmit power for TX2.

Figure 1: Results of two transmitter experiments comparing $JRSS(e)$ and $JRSS(m)$. Transmit power of one transmitter (TX2) is kept constant.

timing beacon interval so that the beacons do not interfere with the measurements. But they are repeated in different beacon intervals for obtaining desired confidence levels.

3. EXPERIMENTAL RESULTS

Given n concurrent transmitters, define joint *estimated* received signal strength (power) or RSS at a designated receiver as,

$$JRSS(e) = \sum_{i=1}^{n} RSS_i \tag{1}$$

Here, RSS_i refers to the individual RSS of transmitter i as measured at the receiver when no other transmitter is active. In contrast, the joint *measured* RSS or $JRSS(m)$ is an average of several RSS measurements at the receiver when the same n transmitters are *actually* transmitting concurrently. Given a set of concurrent transmitters, all experiments described here use the following method to measure individual RSS_i, $JRSS(e)$ and $JRSS(m)$.

1. *Individual RSS_i measurements and estimation of JRSS(e):* Each transmitter i in the given set takes turn to broadcast 1000 packets of size 128 bytes in succession. Each packet transmission time is approximately 4 ms. Any designated receiver samples the RSS register every 3 ms to obtain RSS on its link with the transmitter. (More frequent sampling did not change the measured RSS.) The average of these RSS values are recorded as the individual RSS_i and used for $JRSS(e)$ computation in equation 1.

2. *JRSS(m) measurement:* Here, the same set of transmitters concurrently transmit 1000 packets each. The designated receiver samples the RSS register every 3 ms to obtain the received signal strength. The average of this value over the given 1000 packet interval is taken as $JRSS(m)$.

In the following, we first study the relation between $JRSS(e)$ and $JRSS(m)$ for the scenarios similar to one described in [8] with two and more concurrent transmitters. Later, we describe our results of a more extensive study of the additive assumption for signal strengths with multiple concurrent transmitters particularly targeting low power scenarios.

3.1 Two Transmitters

Two concurrent transmitters (TX1 and TX2) are placed at equal distances from a receiver. Keeping the transmit power of TX2 constant, the transmit power of TX1 is varied from -32.5 dBm to 0 dBm. The resulting RSS at the receiver for TX1 varies from -100 dBm to -55 dBm. For three different experiments, three different transmit powers of TX2 were used, resulting in RSS of -75, -66 and -62 dBm at the receiver. Thus, the three experiments represent three distinct regions of transmit powers: low, medium and high. The results are shown in Figure 1. Here, TX1 and TX2 refer to the individual RSS's observed by the receiver from the transmitters TX1 and TX2 respectively, in absence of any other transmission. $JRSS(e)$ and $JRSS(m)$ refer to the joint estimated and measured RSS at the receiver respectively, when both TX1 and TX2 transmit concurrently.

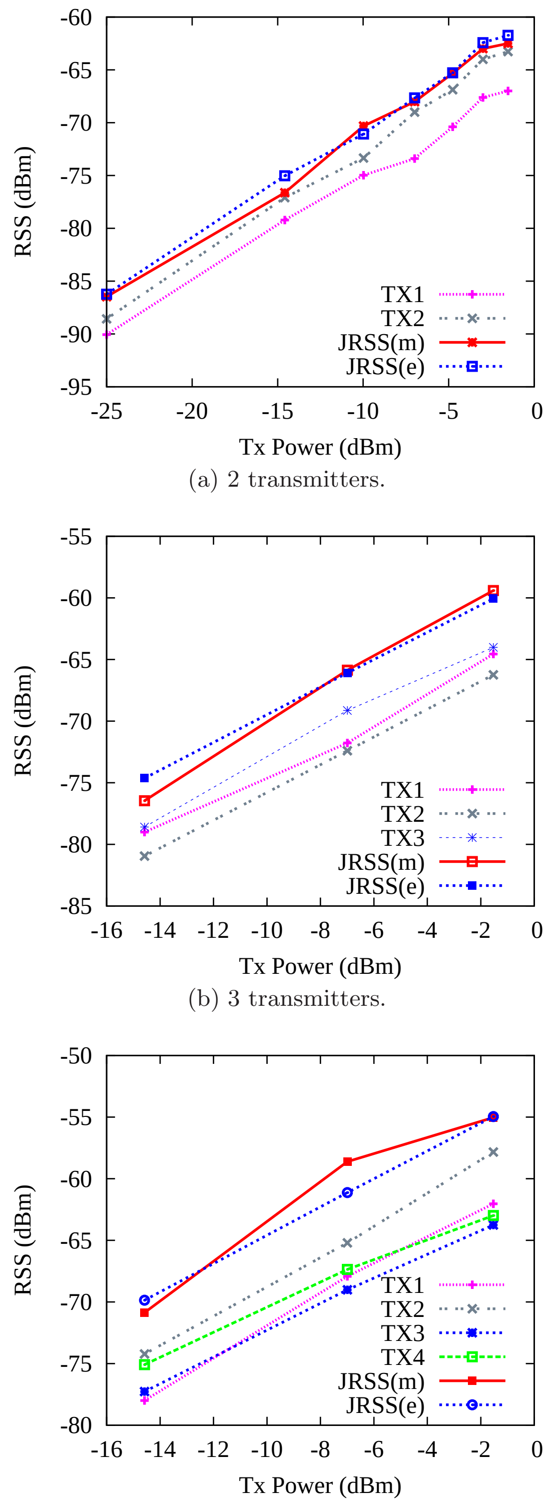

(a) 2 transmitters.

(b) 3 transmitters.

(c) 4 transmitters.

Figure 2: Results of 2–4 transmitter experiments comparing $JRSS(e)$ and $JRSS(m)$. Transmit powers of all transmitters are varied.

95% confidence interval of $JRSS(m)$ is very small, typically within 0.4% average, hence it is not shown in the plots. The same holds true for all other measurements.

With a very similar setup, the authors in [8] observed that $JRSS(e)$ always overestimates $JRSS(m)$ by about 3-5 dB. However, in Figure 1 they are very similar except at the high power end (high transmit power for TX1). This trend is consistent across all three plots. Also note that the additive interference assumption holds true for $JRSS(e)$ less than -60 dBm. When $JRSS(e)$ is less than -60 dBm, $JRSS(m)$ differs from $JRSS(e)$ on an average by about $+0.067$ dBm, -0.38 dBm, and -0.22 dBm in Figures 1(a), 1(b) and 1(c) respectively. Beyond that, $JRSS(m)$ and $JRSS(e)$ sometimes differ with difference usually within a few dBs. The maximum is about 5 dB. But unlike [8], there is no specific trend for over- or under-estimation.

To negate any effect on JRSS due to an overly dominant nature of TX1 signal over TX2 or vice versa, we also conducted experiments where both TX1 and TX2 vary their transmit powers similarly. The results are shown in Figure 2(a). $JRSS(m)$ and $JRSS(e)$ again are found to be quite similar.

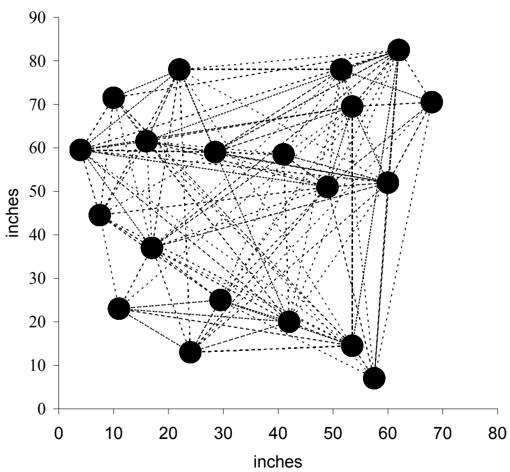

Figure 3: Topology of the 20 motes testbed for -32.5 dBm transmit power. Links shown have at least 99% packet reception rate.

3.2 Multiple Transmitters

To observe the relation between $JRSS(e)$ and $JRSS(m)$ when more than two transmitters are transmitting, we performed experiments with three and four concurrent transmitters. The transmitters were carefully placed at equal distances to the receiver such that the individual RSS's at the receiver do not differ substantially. Three experiments are performed corresponding to low (-14.5 dBm), medium (-7 dBm) and high (-1.5 dBm) transmit powers. The same transmit powers are chosen for all transmitters. The results are plotted similar to the case of two transmitters before. See Figure 2(b) and Figure 2(c).

Note the close match between $JRSS(e)$ and $JRSS(m)$. The minor differences observed in a couple of cases are likely to due to measurement noises. Again, there is no specific trend of over- or under-estimation. Note again that for low power power cases (JRSS below -60 dBm) the match is quite good. To strengthen this observation of close match between estimated and measured $JRSS$ at least for lower power links, we use a 20-node TelosB motes testbed for further experimentation. The nodes are randomly placed on a large table. All nodes use transmit power of -32.5 dBm to limit joint interference power under -60 dBm resulting in a topology as shown in Figure 3. In the figure, a link is shown if packet reception rate on the link is greater than 99%.

We conduct experiments with n concurrent transmitters, varying n from 2 to 6. Out of the 20 nodes, 10 nodes are randomly chosen to be transmitters and the remaining 10 as receivers. For each experiment, we enumerate all possible sets of n transmitters out of the chosen 10 transmitters. We then measure the $JRSS(e)$ and $JRSS(m)$ for all these sets in the same way as before.

The results are shown in Figure 4. Each subfigure in Figure 4 corresponds to n concurrent transmitters with n varying from 2 to 6. Each datapoint obtained in the above experiments is plotted as a scatterplot with $JRSS(m)$ along the X-axis and the corresponding $JRSS(e)$ along the Y-axis. If the additive interference assumption is true, then $JRSS(m)$ would be equal to $JRSS(e)$ and thus, all points in Figure 4 would lie on the $Y = X$ line (shown on the plots). All points indeed lie on or very close to the $Y = X$ line. The corresponding *coefficient of determination*, or R^2 (describing how well the data fits the model $JRSS(m) = JRSS(e)$), is also shown for each n. Note that the R^2 values for all the different number of transmitters is above 95%, signifying an excellent fit. No specific trend on the goodness of fit is observed with the number of concurrent transmitters n. This further strengthens the observation that the additive model is independent of the number of transmitters.

4. DISCUSSIONS

Our observations are in sharp contrast to those in [8] regarding additive signal strength assumption. Our experiments follow similar setup and methodology, but we see no evidence of systematic under- or over-estimation of the joint received signal power, when the joint power is modeled as sum of individual powers measured via individual experiments. We did observe small differences between estimated and measured powers when the joint power is high. However, still there is no evidence of any systematic under- or over-estimation. We hypothesize that these differences are due to measurement errors.

Our strongest observation is, however, at the low power ends (joint power below -60 dBm). Here, we have shown statistically significant evidence that received powers are indeed additive. This is true even when used in a network-context with many motes and with several interferers (upto 6).[2] Interestingly, results in [8] shows consistent over-estimation even at low power ends.

Note that the authors in [8] did not have a clear conclusion why the additive signal power assumption consistently overestimates the real joint signal power observed at the receiver when a set of transmitters transmit concurrently. They attributed this to a combination of poor radio hardware, hardware variations and measurement noises.

It is instructive to summarize here the differences in the hardware setup in the two works. First, the older generation

[2]In this small workshop paper, we are unable to include results from similar experiments at the high power end. However, initial results have shown higher levels of noise, but no systematic errors.

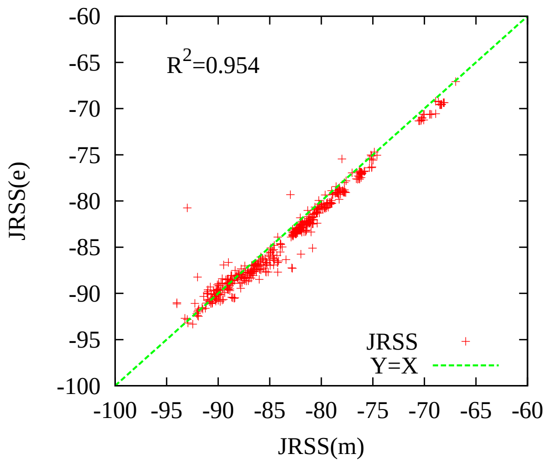

(a) 2 transmitters.

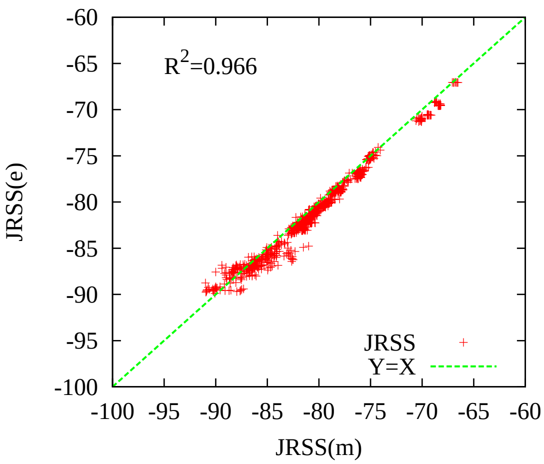

(b) 3 transmitters.

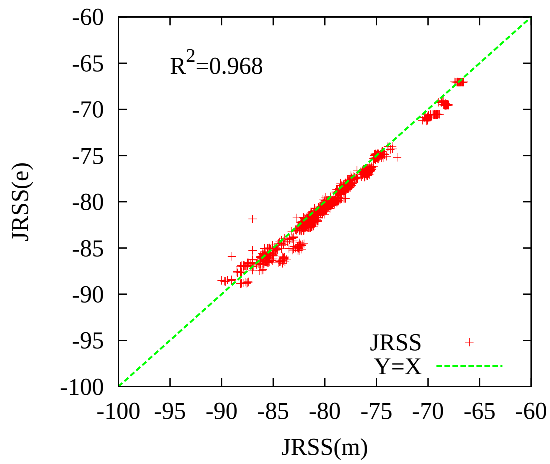

(c) 4 transmitters.

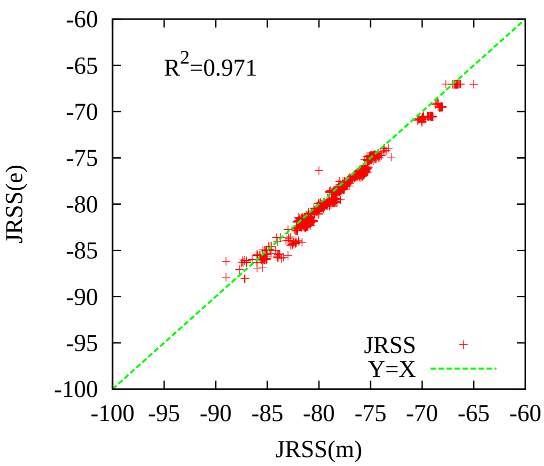

(d) 5 transmitters.

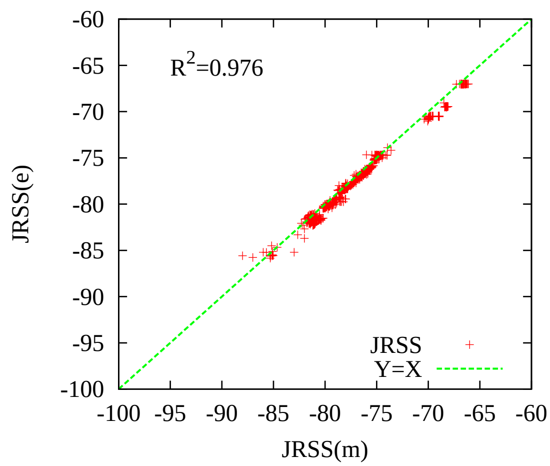

(e) 6 transmitters.

Figure 4: Results of experiments with low power links in a 20 node testbed comparing $JRSS(e)$ and $JRSS(m)$ for various number of concurrent transmitters.

CC1000 radio used in [8] was tuned to 433 MHz band while the CC2420 radio used in our work operates in the 2.4 GHz band. We do not anticipate this is a cause of the differences. We have taken care to use channels that are outside the operating regions of 802.11 devices, which are predominant in our laboratories. Second, RSS is obtained through different methods in the two radios. For CC1000, RSS measurements are taken by measuring voltage at an output pin of the radio by an ADC in the mote micro-controller and then converting the voltage to dBm according to the manufacturer datasheet [11]. This indirect way to measure RSS could be prone to systematic errors. For CC2420, RSS is directly obtained in dBm by reading a register. The RSS register always stores the average RSS over last 32 bit periods. Also, a status bit indicates whether the data available is valid. This direct means of measurements may be less error prone. Note that $JRSS(m)$ measurements in [8] showed large variations with 95% confidence interval sometimes 3-4 dBm wide. In contrast, in this work, the 95% confidence interval of $JRSS(m)$ is found to be almost negligible. This could be due to the differences in the RSS measurement methodologies. Third, we have noticed via independent experiments that use of battery power in motes causes significant power fluctuations, particularly if the battery is not fully charged. Thus, in our work, we have powered the motes always via their USB port, thus effectively eliminating any fluctuations due to battery effects. It is unclear how motes were powered in [8].

5. RELATED WORK

As should be obvious by now, the most closely related work is the study reported in [8]. Other than the talking point in the current paper, [8] made several other conclusions related to multiple concurrent transmissions in mote-based wireless networks. For example, they concluded that the SINR threshold for correct packet reception is dependent on number of interferers and that slightly different behaviors dependent on received power ranges.

Researchers have only begun to study effect of interference in mote-based wireless networks using experimental methods. Authors in [13] have studied the *transition region* and quantified its effects. The analysis in the paper is also supported by experimental validation using a motes testbed, though with the older generation (CC1000) radios.

In a different line of work [9], the authors have concluded from measurements on MicaZ motes with CC2420 radios, that RSSI is a good estimate of link quality.

Experimental work has also considered 802.11-based systems to study interference behavior. The difference here is that the sender-side (carrier-sense) behavior in the MAC protocol must also be modeled. Notable articles are as follows. Single and multiple interferer scenarios have been modeled in [7] and [5], respectively. The need for modeling multiple interferers has been motivated in [3].

6. CONCLUSIONS

We have demonstrated – using newer mote-class radios (CC2420) – that received signal power is indeed additive, at least for low-power links. This is in contrast with existing work [8] that has used older CC1000 radios and has concluded that the additive model overestimates the joint signal power. The significant deviation of our observation vis-a-vis prior work shows the exact nature of the hardware platform matters significantly in modeling radio behaviors.

In this regard, it is worth pointing out that literature has seen departure from the expected for the older radios in other contexts as well. For example, in [9] it has been shown that RSSI is indeed a reasonable indicator of link quality while prior research using older CC1000 and TR1000 radios concluded otherwise [9, 14]. We believe that wireless sensor network community needs to make serious validation in multiple platforms before building acceptable empirical models related to radio characteristics. Also, measurement noises are expected to be lower in newer platforms bringing empirical observations closer to theoretical predictions.

A direct impact of our work is in interference modeling for low power networks. Theoretical interference models like SINR-based models can now be used in practice as joint interference can simply be taken as a sum of individual interference powers. This reduces the complexity of realizing these models from exponential to polynomial. This, in turn, makes possible direct application of these models to practical TDMA scheduling.

7. REFERENCES

[1] "Moteiv," http://www.moteiv.com.

[2] "TinyOS community forum," http://www.tinyos.net.

[3] S. Das, D. Koutsonikolas, Y. Hu, and D. Peroulis, "Characterizing multi-way interference in wireless mesh networks." in *ACM Wintech*, 2005.

[4] IEEE Computer Society LAN/MAN Standards Committee, "802.15.4: Wireless medium access control (MAC) and physical layer (PHY) specifications for low-rate wireless personal area networks (LR-WPANS)," 2003.

[5] A. Kashyap, S. Ganguly, and S. R. Das, "A Measurement-Based Approach to Modeling Link Capacity in 802.11-based Wireless Networks," in *ACM MobiCom*, 2007.

[6] Moteiv Corporation, *tmote sky: Ultra low power IEEE 802.15.4 compliant wireless sensor module*, San Fransisco, CA, November 2006.

[7] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan, "Measurement-based models of delivery and interference in static wireless networks," in *SIGCOMM*, 2006.

[8] D. Son, B. Krishnamachari, and J. Heidemann, "Experimental study of concurrent transmission in wireless sensor networks," in *SenSys '06*, 2006.

[9] K. Srinivasan and P. Levis, "RSSI is Under Appreciated," in *EmNets 2006*, 2006.

[10] *CC2420 Radio Datasheet*, 1st ed., Texas Instruments, October 2005.

[11] *CC1000 Radio Datasheet*, 2nd ed., Texas Instruments, January 2007.

[12] K. Whitehouse, A. Woo, F. Jiang, J. Polastre, and D. Culler, "Exploiting the capture effect for collision detection and recovery," in *IEEE EmNetS-II.*

[13] M. Zuniga and B. Krishnamachari, "Analyzing the transitional region in low power wireless links," in *IEEE SECON 2004*, 2004, pp. 517–526.

[14] J. Zhao and R. Govindan. "Understanding packet delivery performance in dense wireless sensor networks," in *ACM SenSys*, 2003.