Towards Concurrent Communication in Wireless Networks

USC/ISI Technical Report ISI-TR-648, July 2007

Dongjin Son[1,2] John Heidemann[2] Bhaskar Krishnamachari[1]
dongjins@usc.edu, johnh@isi.edu, bkrishna@usc.edu
[1]Department of Electrical Engineering-Systems, [2]Information Sciences Institute
Viterbi School of Engineering, University of Southern California

Abstract

Avoiding collisions is one of the key roles of media-access (MAC) protocols. Since MACAW and 802.11, carrier sense and exchange of request-to-send (RTS) and clear-to-send (CTS) packets have been used to prevent concurrent communication in wireless networks. Yet these approaches have significant cost: they prevent *all* concurrent communication, even exchanges that might not result in loss; they reduce end-to-end throughput in a multi-hop network; and control traffic imposes control overhead on networks with small data payloads such as 802.15.4. In this paper, we show that RTS/CTS is almost *never* desirable in modern wireless networks that support power control and channel capture. We use four-node experiments with 802.15.4 radios to show that *concurrent communication* is often possible, depending on node locations and transmit powers. We validate an SINR-based propagation model against these experiments and use simulation to systematically explore how node location affects the ability to communicate. Given optimal power settings, when two sources are outside some minimal distance, they can communicate concurrently with two receivers more than 80% of the time. An optimal algorithm requires perfect knowledge of the channel and transmission state, so we then sketch *Gain-Adaptive Power Control*, a MAC protocol that provides significant benefit with only local and prior information. Compared to optimal, we show that this practical MAC can transmit concurrently 75% of the time, but requires a larger minimum source separation. We also show that least one sender can capture the channel 77–88% of the time, regardless of source and receiver location, so the cost of failed concurrent communication is only slightly worse than RTS/CTS. These results provide compelling evidence that future MAC protocols should exploit power control and channel capture.

1 Introduction

Avoiding collisions is one of the key roles of media-access (MAC) protocols. Research in MACAW [1] and standards such as 802.11 employ carrier sense and exchange of request-to-send (RTS) and clear-to-send (CTS) packets to prevent concurrent communications and hidden terminal cases that might corrupt communication. Nevertheless, *concurrent communication*—allowing transmission by two senders at the same time over the same channel—can be beneficial, provided both receivers can successfully receive what is sent. The benefits of concurrent concurrent communication come because carrier sense and RTS/CTS greatly reduce opportunities for spatial reuse of the channel. In a multi-hop network, RTS/CTS-enforced-silence reduces end-to-end throughput. And for networks designed for relatively small data payloads, such as 802.15.4, the RTS/CTS exchange is avoided as control overhead.

Recent work has begun exploiting the richness of real-world wireless propagation and richer MAC protocols. Experiments have shown that MAC protocols can exploit channel capture, either by retraining mid-reception [4, 12, 13], or using more aggressive carrier sense [3]. Other work has shown that power control can allow transmission "over the heads" of intermediate nodes [7]. Experiments have evaluated power control to maximize spatial reuse [9, 5], and to develop better models of wireless propagation [6, 8, 10]. Experimental work has also suggested the importance of SINR-based channel models that represent the intermittent, power- and location-sensitive reception inherent in concurrent wireless communication [7, 10]. This range of work provides components for interference-aware protocol design and has shown the feasibility of concurrent communication with modern radios that provide per-packet power control and MAC-level channel capture. While very promising, this work has yet to suggest a specific new MAC protocol or quantify trade-offs.

This paper seeks to answer two open questions: *how significant are the benefits* of concurrent communication, and *how close can a practical MAC* approach optimal?

These questions appear deceptively straightforward. When possible, concurrent communication will obviously improve spatial reuse, and its existence has been shown experimentally [7]. However, it is not clear *how often* concurrent communication is possible, since if a sender transmits at minimum power needed to reach its receiver, any other concurrent communication raises the noise floor, forcing the original sender to further raise its power. Given this coupling and potential costs of coordinating multiple senders, quantifying the benefit of concurrent communication is essential.

To quantify the benefits of concurrent communication, we define *CCability*, the fraction of spatial locations where concurrent communication is possible with two concurrent senders. Prior work has shown instances where concurrent communication is possible [12, 7]; our contribution is to show through testbed experiments how concurrent communication is affected by location and transmit power (Section 2). We use these experiments to calibrate our SINR-based channel model [10], and then use simulation to quantify the idealized opportunities for concurrent communication as a function of location, transmission power limitation.

Our second contribution is to relate these bounds on concurrent communication to what can be accomplished in a real-world MAC protocol. Our optimal bounds require perfect knowledge of all channel state: all concurrent communication, node locations, and noise; information impossible to maintain in a realistic network, and complex and expensive to approximate. On the other hand, a very simple MAC might send at the lowest possible power to maximize channel reuse. We evaluate the benefits of designs that employ different amounts of information relative to our optimal performance bound (Section 3).

While this short paper does not advocate a complete new MAC, we show that two pairs of transmitters can communicate concurrently more than 80% of the time with sufficient source separation, given perfect channel knowledge. We also show that a practical gain adaptive power control-based MAC protocol can provide much of this benefit. These results suggest that future MAC designs should embrace concurrent communication through power control and channel capture and shift away from carrier sense and RTS/CTS.

2 Experimental Evaluation

Previous work has shown examples of channel capture for specific scenarios [12, 7]. We first present experimental results with real hardware to show how node location and transmit power allow concurrent communication, capture, or result in collision.

2.1 Methodology

Our testbed experiments follow the methodology of recent studies of concurrent communication [10]. We use two sender-receiver pairs of nodes, S1-R1 and S2-R2. A fifth node, the *synchronizer*, coordinates sender transmission with a trigger message. We use MicaZ motes with CC2420 radios [2] for our experiments because they provide power control, a completely programmable MAC, and low-level hardware access for accurate timing. We disable carrier sensing and randomized back-off from the MAC layer to allow concurrent packet transmission from multiple senders.

We consider both crossed and adjacent communication, as shown at the top of Figure 1. Sender 2 (S2) moves to each location indicated with a lowercase letter, while its receiver is positioned outside the S1-R1 pair. We vary the S2-R2 distance, considering ten different positions of S2, roughly every 60 cm. We skip positions where S2 would be in the same location as another node.

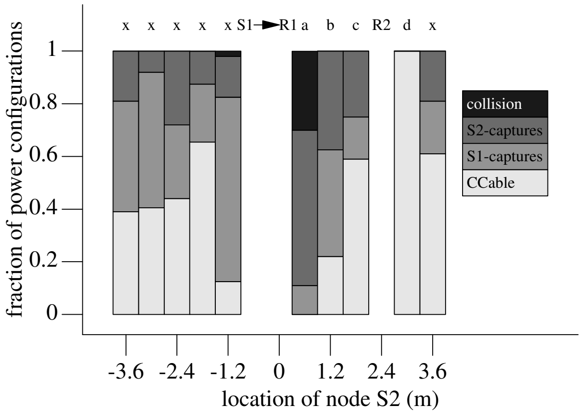

Figure 1. MicaZ experiments measuring CCability with two sender-receiver pairs. S1, R1, and R2 are placed at -0.6m, 0m, and 2.4m as shown above, while S2 is measured at each of the 10 locations marked with lowercase letters.

We vary transmit power to test communication. The MicaZ supports 8 different transmission power levels from -25 to 0 dBm. For each position experiment, we first measure the signal and interference strength with 10 packets and then test the CCability with 25 concurrent packet transmissions for each of all 64 different combinations of two senders' transmission power settings. We repeat the same experiment twice for each topology to verify that the results are consistent; the results were similar and we show only one representative experiment here.

2.2 Concurrent Communication, Capture, and Collisions

The graph in Figure 1 summarizes our experimental results. Considering each of 10 positions, we can see that in nine of the ten cases there is *some* power configuration that supports concurrent communication, and often there is considerable flexibility in the exact power settings. Only when R2 is at 0.6 m, close to R1, is concurrent communication impossible.

This experiment demonstrates the large opportunity for concurrent communication if MAC support for packet capture and appropriate power selection was available and RTS/CTS was revised. Nevertheless, current MAC protocols would prohibit many of these opportunities to transmit due to carrier sense detecting a busy channel, or RTS/CTS forbidding communication.

This experiment also shows that even sub-optimal power settings often allow at least one sender or the other to capture the channel, at least when nodes are not directly on top of each other. The fraction of CCable power combinations by itself is not a useful metric, since an intelligent MAC would not select transmission power randomly, but this level of *flexibility* in power selection is important to implementing a MAC with imperfect information, that is tolerant to environment noise and interference, as we show in Section 3.

Figure 2 shows a more detailed view of reception for the four positions of S2 marked (a) through (d) in Figure 1. We consider cases where S2 is located to the right of R1; we observe a similar trend and implications when S2 is to the left

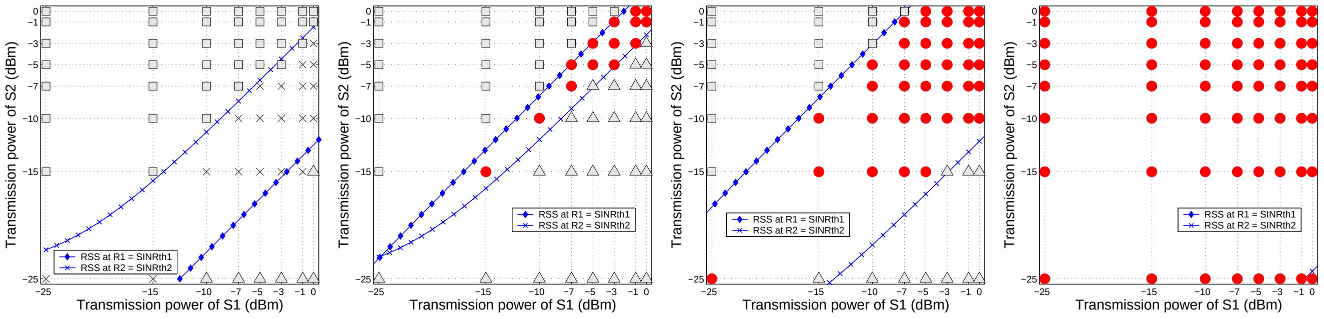

(a) S2 at 0.6m (b) S2 at 1.2m (c) S2 at 1.8m (d) S2 at 3.0m

Figure 2. CCability in the outside testbed experiment as S2 is moved (presented together with the expectation from simulation with our proposed formula) . Circles are *CCable*, triangles and squares are *S1* or *S2 capturable*, and Xs indicate a collision. Simulation results at the same topology are presented together with two dotted lines.

of S1 (sending over the S1-R1 pair) as well, but omit those results due to space. Figure 2 shows the result of each transmission for all possible power levels for these cases. Results of each test are shown by different symbols: filled circles are CCable, while empty triangles or squares indicate capture by S1 or S2, and Xs indicate collisions where neither receiver can capture data.

These plots show vividly that the key determinant of concurrent communication is power control by the two transmitters. As the source separation increases from (a) to (d), so do the number of power level settings that allow for concurrent communication, as shown by the greater number of circles.

We can also see from Figures 1 and 2 that even if two transmissions are not CCable, almost always one or the other can be delivered with the capture effect. The SINR threshold of the MicaZ is around 2 dB [10], and the low number of collisions in this experiment shows that it is rare for RSSs from both senders to fall within this 2 dB range. In our experiments, only 3% of power configurations resulted in collisions. This is consistent with previously reported experiments pertaining to capture [12].

2.3 Validating Our SINR Model

While experiments are the ultimate test of behavior, it is not practical to explore the entire parameter space experimentally, and impossible to predict the behavior of potential future hardware. Our prior experimental studies have evaluated how hardware, location, and power affect concurrent communication and suggested an SINR-based model, whereby each receiver receives a packet successfully if and only if the ratio of the signal power from its intended sender to the sum of interference power from the other sender and the noise power exceeds a given threshold [11]. We next evaluate the data from the experiments and compare them to our SINR model.

To compare our experiments with simulation, we use the algorithm defined in our prior work to predict the power settings that enable concurrent communication [11]. We compute these values plot them as two lines in the plots of Figure 2. The simulations require parameters for the channel propagation model that we do not know, so we use the measured path loss at each location. We also used observed values for SINR threshold (2 dB for MicaZ) and ambient noise level for each node (−96.3 dBm for R1 and −96 dBm for R2). We can see that our simulation results match closely the experimentally observed CCability at different transmit power settings.

2.4 Conclusions from Experiments

These testbed experiments provide the following three main observations: First, concurrent communication is highly probable in many previously restricted cases with traditional 802.11-like medium access control. Second, complete collisions and full corruption of both packets is rare and often at least one sender can capture a packet. Finally, controllable transmission power significantly improves CCability. In addition, they provide evidence that our SINR model can predict power settings for CCability with two concurrent transmitters.

3 MAC Protocol Designs

We have established that for the vast majority of topologies where senders have reasonable separation, concurrent communication is possible given complete information. Yet how close practice can come to this bound is not clear, since a practical MAC protocol must make control decisions based only on prior knowledge and local information.

We next consider five different power control algorithms for MAC protocols. Carrier sense with RTS/CTS at maximum and minimum power represents the current state-of-the-art. We present an *oracle* algorithm, to provide an upper bound on performance given unachievably perfect information. We then introduce two simple MAC protocols that use only local and prior information. *MinPC* sends at minimum power with channel capture; a very simple way to improve spatial reuse given prior knowledge of node locations. Finally, *gain-adaptive power control* (GAPC) adds a transmit-power-dependent boost to MinPC to overcome some potential interference.

We evaluate these protocols through simulation using the SINR-based model that we validated in Section 2.3. We use an exponential path-loss model with the option of realistic log-normal multipath fading in our simulations to obtain the pair-wise link gains. In each simulation we consider two sender-receiver pairs. We fix the location of one pair and the second sender, move the second receiver over all possible locations with possible reception, and measure which receivers can capture concurrently sent packets.

Figure 3 shows CCability for each protocol, in this figure black indicates the CCable region, gray shows where one communication or the other is capturable, and white shows inability to communicate, either due to power limitations (outside the circle) or due to collisions (inside). The two rows of this figure show results without (top) and with (bottom) variance in link gain due to model multi-path fading. As can be seen, while fading effects do introduce a degree of noise, they do not fundamentally change the results. For ease of exposition, therefore, we ignore fading in subsequent discussion as we consider each design alternative.

3.1 Today's practice: CS-RTS/CTS with Simple Power Control

We begin by evaluating traditional control methods. Figures 3(a) and 3(b) show the behavior of a traditional carrier-sense with RTS/CTS MAC. Simplest is to always transmit RTS/CTS at maximum power to block any potential receivers. As shown in Figure 3(a), this case always allows one sender, but never concurrent communication.

Slightly better is to send at minimum possible power needed to reach the indented receiver (assuming unicast communication). Taking this step requires that each node maintain a list of neighbors and estimates of the transmit power needed to exceed their SNR threshold. We assume this information is collected and reasonably stable, a valide assumption for slow-fading environments with little mobility [9]. In this case, the small black crescent in Figure 3(b) shows that even with CS-RTS/CTS we can get some concurrent communication when R2 is located near S2—approximately 2% of the area. While better than maximum power, we suggest that this gain is too modest relative to the measurement overhead to motivate use of power control with CS-RTS/CTS.

3.2 A Upper Bound on Performance with an Oracle

Given perfect knowledge of the gain (or path loss) between the nodes in the network, any concurrent communication, and noise, one can compute the optimal (minimal) transmit power for concurrent communication. We describe this approach in our prior work [11]. While gain can be observe in the network and its variance estimated, one cannot have perfect knowledge of all concurrent communications. We next describe an oracle algorithm that uses this perfect knowledge establish an upper bound on the benefit we can expect from concurrent communication. While this oracle uses perfect knowledge, it is still subject to hardware limitations of discrete power levels and minimum and maximum transmit power.

Figure 3(c) shows sample results with the oracle. Compared to CS-RTS/CTS at different power levels (Figures 3(a) and 3(b)), we can see that there is considerable room for concurrent communication. There is a small hole near the center, occupying about 3% of the possible area, where only one communication can be allowed. In this region receiver R2 is too close to sender S1 or receiver R1 for both to communicate given a maximum transmit power. We call this region the *region of impossible concurrency*.

Of course, this scenario represents just one topology. However, we observe similar results provided the two sources are separated by some minimum distance (outside the region of impossible concurrency). We examine different source and receiver placement in more detail below (Figures 4 and 5).

For this configuration, the oracle allows concurrent communication over 97% of the R2 locations. This evaluation demonstrates the potential of channel capture and power-control, if we define a MAC algorithm that uses practical information.

3.3 Exploiting Power Control and Channel Capture

Ignoring interference, we maximize spatial reuse by always sending at the minimum power that will reach the intended receiver. Our *capture(MinPC)* algorithm takes this approach to power control, and employs MAC-level channel capture [12]. It is therefore equivalent to CS-RTS/CTS(MinPC), but replacing CS-RTS/CTS with channel capture. As with CS-RTS/CTS(MinPC), we assume nodes maintain a list of neighbors and can track the gain needed to reach them.

In theory, sending at the minimum transmit power cannot tolerate any level of interference. However, in practice, real hardware can be set only at discrete power levels, providing some level of protection to noise. (We use discrete levels at 1 dBm increments in this simulation; the CC2420 radio provides slightly coarser levels.)

Figure 3(d) shows a moderate size region where concurrent communication is possible, 20% of the total area in this case. Comparing this to CS-RTS/CTS(MinPC) demonstrates the advantage of channel capture over communications prohibition. In addition, the large grey region shows that, even when concurrent communication is not possible, at least one receiver or the other will get their data through. In this case, CC or capture is 87% of the total area.

The penalty of allowing concurrent communication is the small white crescent region where transmit powers are evenly matched at the receivers, resulting in collisions without capture. With RTS/CTS, one sender or the other would win the contention and send, but with capture we depend on random backoff and retry when nodes are at this range.

3.4 Gain-Adaptive Power Control and Capture

While discrete power levels provide some buffer against noise with the capture(MinPC) algorithm, concurrent communication provides strong sources of interference that limit the ability of min-power to approach oracle. This problem is particularly noticeable at the edges of S2's range, where interference for S1 prevents S2 reception. In Figure 3(d) this case appears as the large grey doughnut surrounding the black CCable region.

This conditions can be overcome by systematically adding a boost of power in inverse proportion to the gain needed to reach the intended receiver. We call this algorithm

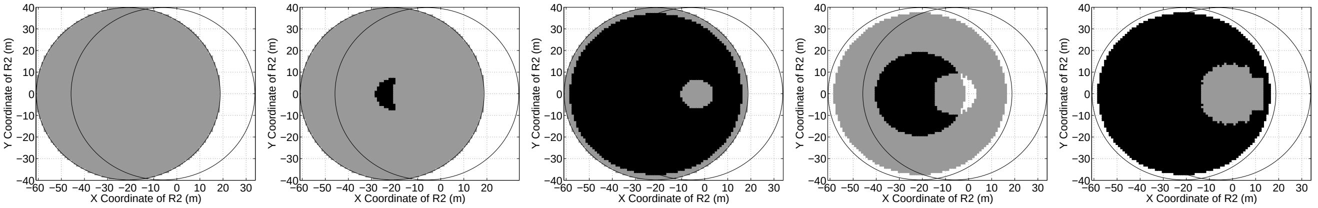

(a) CS-RTS/CTS(Max) (b) CS-RTS/CTS(MinPC) (c) Oracle (d) MinPC (e) GAPC

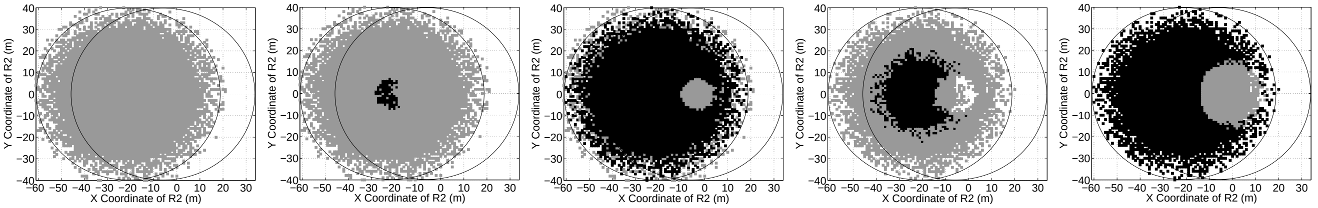

(f) CS-RTS/CTS(Max) (g) CS-RTS/CTS(MinPC) (h) Oracle (i) MinPC (j) GAPC

Figure 3. MAC power control comparison of CCability. The bottom row shows the case with fading variance. $S1 = (-6,0)$**,** $R1 = (0,0)$**,** $S2 = (-21,0)$**, and vary** $R2$ **between -35 m and 35 m both in X and Y directions,** $n = 3.5$**,** $SINR_\theta = 2$**,** $X_\sigma = 0$ **for the top row,** $X_\sigma = 3$ **for the bottom row. Black = CCable, Gray = Capture, White = No communication.**

Gain-Adaptive Power Control. Since gain is roughly proportional to distance, this means that short-distance transmissions get large boosts while longer transmission gets relatively less gain. Our intuition for this scheme comes from observations in a detailed simulation study [11]: we found that short distance communication is often overwhelmed by interference from longer transmissions.

If we define P_{max} and $P_{S,R}$ as the maximum possible transmit power and the power needed for source S to reach receiver R, then we can define the power boost ε as:

$$\varepsilon = (P_{max} - P_{S,R})\varepsilon_{ratio} \qquad (1)$$

In this equation, ε_{ratio} represents the fraction of remaining power to allocate to a transmission. Large values of ε_{ratio} will quickly assign all headroom to transmissions and will increase the bonus given to shorter links. We varied ε_{ratio} and found that moderate values (0.3 to 0.7) provided the best levels of CCability (values that near 0 provide no boost, while values approaching 1 always operate at maximum power). We adopt $\varepsilon_{ratio} = 0.5$ as a reasonable, robust choice.

Figure 3(e) evaluates gain-adaptive power control with $\varepsilon_{ratio} = 0.5$. We see that this approach comes very close to optimal: concurrent communication is possible with the receiver in 76% of the area compared to the oracle algorithm, much closer than capture(MinPC).

The cost of gain-adaptive control relative to the oracle can be seen in two locations. The moderate-size grey area when R2 is placed near $(0,0)$ is larger than optimal. This area corresponds to cases where R1 and R2 are competing and the power boost prevents concurrent communication. In this region it is best if only one sender transmits. Second, communication in the narrow gray ring around the edge of the oracle cannot be reached with gain-adaptive control because of slightly higher interference from the S1-R1 pair.

These results suggest that gain-adaptive power control is a practical scheme that gets a significant fraction of optimal.

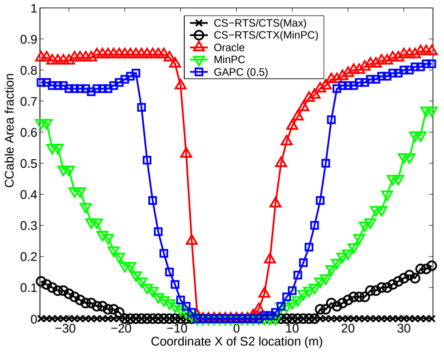

Figure 4. Comparison of CCable area with limited power levels (-25 dBm to 0 dBm) for five MACs.

3.5 Comparing MAC Protocols

Figure 3 compares the five MAC algorithms for a particular topology. We can quantify the benefit of concurrent communication by observing the ratio of area of concurrent communication (anything black) to the total reachable area when there is no interference (indicated by thin black circles, also equal to the the gray area with CS-RTS/CTS at maximum power in Figure 3(a)). We define this ratio as the *CCable area*.

Figure 4 compares this CCable area for each of the MAC schemes we consider. This graph provides a single slice through the 2-D simulation with nodes placed at $S1 = (-6,0)$, $R1 = (0,0)$, S2 at coordinate (x, 0), with the x-coordinate indicated on the horizontal axis of the graph, evaluated for all R2 locations over all potentially receivable locations. Each point on the figure represents the fraction of R2 locations that allow concurrent communication for a given S2 x position.

Even with power control, Figure 4 shows that CS-RTS/CTS provides little spatial reuse of the channel through concurrent communication; this result is consistent with its design goal of preventing any possible interference. Shifting

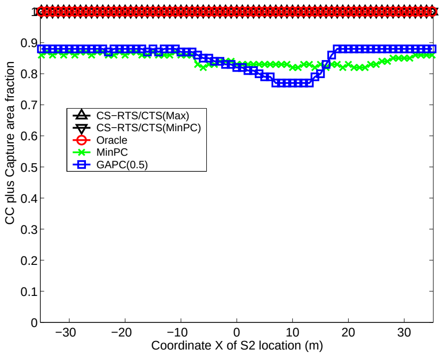

Figure 5. CC plus Capture rate comparison with limited power levels (-25 dBm – 0 dBm)

to a channel-capture-based MAC (min-power) provides considerably greater opportunity for concurrent communication. Finally, GAPC comes relatively close to the best possible oracle result. We find that its CCable area is within 73% of the optimal, averaged over all S2 locations, and is as high as 95% in regions with sufficient source separation.

We can also observe *where* concurrent communication is possible. All algorithms, including the oracle, fail when both senders (or receivers) are nearly in the same place. In this region of impossible concurrency, no level of power is sufficient to capture the channel. The algorithms differ mainly in the width of this region—more sophisticated algorithms are closer to the oracle's best-possible result.

We can use this same methodology to evaluate not just opportunities for concurrent communication, but for CC *or* channel capture. Figure 5 shows that the oracle performance almost strictly dominates CS-RTS/CTS by this metric—CS-RTS/CTS always gets exactly one packet through, while the oracle always gets one *or two* packets through. Figure 4 shows where two are possible, while in region of impossible concurrency it gets one packet through. By comparison, the realizable algorithms capture(MinPC) and GAPC can get one or two packets through 77–88% of the time. Even when concurrency is impossible, most often one sender captures the channel. The 12–23% gap represents lost capacity in conditions where packet collision allows neither sender to communicate. Most of this loss occurs at near the edge of maximum communication range where GAPC cannot boost power adequately to exceed interference because of hardware limitations.

4 Future Work and Conclusions

It is now widely understood that wireless propagation is much more than receive/no-receive links. Prior work has demonstrated that channel capture and power control are possible in individual cases.

This paper has established that the benefit of exploiting these characteristics is significant. We provide a theoretical upper bound on performance given realistic hardware and perfect knowledge. We then showed that a practically implementable power control algorithm, the GAPC scheme, can get near-optimal performance, averaging 73% of area of concurrent communication obtained the oracle, with successful capture in 77–88% of the cases.

While promising, our work is still a preliminary step. We focused on two pairs of concurrent communications; we believe the results generalize to *n*-node communication (through preliminary simulations not shown here), but through evaluation is future work. More importantly, full implementations of the MAC algorithms that we propose are necessary to provide full experimental validation of our conclusions.

We believe this work establishes an essential direction for future MAC research, away from the use of carrier sense and RTS/CTS to avoid concurrent communication, instead embracing concurrency through power control and channel capture.

5 References

[1] V. Bharghavan, A. Demers, S. Shenker, and L. Zhang. MACAW: A media access protocol for wireless LAN's. In *Proc. of the ACM SIGCOMM Conference*, pages 212–225, London, UK, Sept. 1994. ACM.

[2] Chipcon. CC2420 2.4GHz IEEE 802.15.4/ZigBee-ready RF transceiver.

[3] K. Jamieson, B. Hull, A. Miu, and H. Balakrishnan. Understanding the real-world performance of carrier sense. In *Proc. of the ACM SIGCOMM Workshop on Experimental Approaches to Wireless Network Design and Analysis*, pages 52–57, Philadelphia, PA, USA, Aug. 2005. ACM.

[4] A. Kochut, A. Vasan, A. U. Shankar, and A. Agrawala. Sniffing out the correct physical layer capture model in 802.11b. In *Proc. of the 12th IEEE International Conference on Network Protocols*, pages 252–261, Berlin, Germany, Oct. 2004. IEEE.

[5] S. Lin, T. He, J. Zhang, G. Zhou, L. Gu, and J. A. Stankovic. ATPC: Adaptive transmission power control for wireless sensor. In *Proc. of the Fourth ACM SenSys Conference*, pages 223–236, Boulder, Colorado, USA, Nov. 2006. ACM.

[6] T. Moscibroda. The worst-case capacity of wireless sensor networks. In *Proc. of the 6th ACM International Workshop on Information Processing in Sensor Networks*, pages 1–10, Cambridge, Massachusetts, USA, Apr. 2007. ACM.

[7] T. Moscibroda, R. Wattenhofer, and Y. Weber. Protocol design beyond graph-based models. In *Proc. of the ACM Workshop on Hot Topics in Networks (HotNets-V)*, pages 25–30, Irvine, California, Nov. 2006.

[8] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan. Measurement-based models of delivery and interference in static wireless networks. In *Proc. of the ACM SIGCOMM Conference*, pages 51–62, Pisa, Italy, Sept. 2006. ACM.

[9] D. Son, B. Krishnamachari, and J. Heidemann. Experimental study of the effects of transmission power control and blacklisting in wireless sensor networks. In *Proc. of the First IEEE Conference on Sensor and Adhoc Communication and Networks*, pages 289–298, Santa Clara, California, USA, Oct. 2004. IEEE.

[10] D. Son, B. Krishnamachari, and J. Heidemann. Experimental analysis of concurrent packet transmissions in wireless sensor networks. In *Proc. of the Fourth ACM SenSys Conference*, pages 237–249, Boulder, Colorado, USA, Nov. 2006. ACM.

[11] D. Son, B. Krishnamachari, and J. Heidemann. Evaluating the importance of concurrent packet communication in wireless networks. Technical Report ISI-TR-639, USC/ISI, Apr. 2007.

[12] K. Whitehouse, A. Woo, F. Jiang, J. Polastre, and D. Culler. Exploiting the capture effect for collision detection and recovery. In *Proc. of the IEEE Workshop on Embedded Networked Sensors*, pages 45–52, Sydney, Australia, May 2005. IEEE.

[13] G. Zhou, T. He, J. A. Stankovic, and T. Abdelzaher. RID: Radio interference detection in wireless sensor networks. In *Proc. of the 24th IEEE Infocom*, pages 891–901, Miami, FL, USA, Mar. 2005. IEEE.