Wireless Medium Access for Concurrent Communication

USC/ISI Technical Report ISI-TR-652, May 2008

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

Abstract

Most wireless medium access control (MAC) protocols today prevent nearby concurrent communication due to concern that it would corrupt ongoing communication. While recent work has demonstrated that channel capture allows *successful* concurrent communication, MAC protocols to date have not exploited this approach to improve performance. In this paper we conduct experiments with 802.15.4 radios to model concurrent communication, showing that concurrent communication is possible for nearly all topologies with appropriate power selection. From these experiments, we define a new MAC protocol, *gain-adaptive power control* (GAPC), with the goal of enabling concurrent communication when possible. Unlike prior power-adaptive MAC protocols, GAPC keeps a small power reserve and uses channel capture to support concurrent communication. We develop a fully distributed algorithm to allocate this reserve to boost minimum needed transmit powers, adapting the gain to overcome noise due to other transmissions or a varying environment. We show that this power reserve is *essential* to achieve significant performance advantage from concurrent communication. Finally, we quantify the benefits of concurrent communication, comparing GAPC performance to MACs which channel access that is optimal (Oracle-based); uses non-adaptive, RTS-CTS-based CSMA; or uses simple minimum-transmit power. The GAPC heuristic allows concurrent communication in 73% of cases of optimal in our evaluation. When communicating with neighbors, we show that GAPC can allow 2.6× more successful receptions than CSMA, and 3× more than CSMA with RTS/CTS. For multi-hop communication, GAPC completes a fixed-size transfer faster as well, with CSMA requiring 1.7× longer, or 1.2× longer with RTS/CTS, even ignoring control overhead. Our study shows embracing concurrent communication can significantly improve wireless throughput.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Wireless communication*; C.2.2 [**Computer-Communication Networks**]: Network Protocols

General Terms

Design, Experimentation, Measurement, Performance

Keywords

wireless MAC protocols, power control, channel capture

1 Introduction

Efficient communication is the ultimate goal of a medium-access (MAC) protocol. Preemptively avoiding collisions has been the means to reach this end for many years. Research in MACAW [2] 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 (CC)*—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 the efficiency of streaming data. In networks such as 802.15.4, where payloads are small and so RTS/CTS is not appropriate, applications often take steps to avoid concurrent communication.

Recently, researchers have begun to explore opportunities for concurrent communication in real-world wireless networks. ExOR and network coding exploits partial reception for bulk data transfer [3]. Experiments have shown that MAC protocols can exploit channel capture, either by retraining mid-reception [9, 20, 24], or using more aggressive carrier sense [8]. Other work has shown that power control can allow transmission "over the heads" of intermediate nodes [13]. Experiments have evaluated power control to maximize spatial reuse [11, 16], and have helped develop better models of wireless propagation [12, 15, 17]. 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 [13, 17]. This range of work

This research is supported by the National Science Foundation (NSF) through the following grants: CNS-0347621 (CAREEER), CNS-0325875 (NSF-ITR), CNS-0627028 (NeTs-NOSS), CCF-0430061 (Design Automation of Compute-Intensive Networked Embedded Systems), NeTS-NOSS-045517 (Sensor Networks for Undersea Seismic Experimentation), CNS-9626702 (Sensor-Internet Sharing and Search).

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 several of these efforts have been evaluated with modifications to existing protocols (Section 2), to date there has been no ground-up reevaluation of medium access.

This paper seeks to do just that: we propose *gain-adaptive power control* or GAPC, a MAC protocol designed to exploit concurrent communication in radios that provide power control and channel capture. Evaluating this MAC protocol requires three advances. First, we conduct extensive experiments with 802.15.4 radios (Section 3) to develop a new model of concurrent communication (Section 4). These experiments shows that concurrent communication is possible for nearly all topologies given radio power control, and with the model they suggest that *flexibility in power control* is needed to exploit concurrent communication.

Second, the insights from these experiments inform our design of our MAC protocol, gain-adaptive-power-control (GAPC), Section 5. Unlike prior power-adaptive MAC protocols, GAPC keeps a small power reserve and uses channel capture to support concurrent communication We develop a fully distributed algorithm to allocate this reserve to boost minimum needed transmit powers, adapting the gain to overcome noise from other, possibly distant, transmissions, or variations in channel noise. We show that this power reserve is *essential* to achieve significant performance advantage from concurrent communication.

Finally, this paper quantifies the benefits of exploiting concurrent communication. We seek to answer two fundamental open questions: *how significant are the potential benefits* of concurrent communication, and *to what extent can a practical MAC* realize this potential?

These questions appear deceptively straightforward. When possible, concurrent communication will obviously improve spatial reuse, and its existence has been shown experimentally [13]. 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 [13, 20]; our contribution is to show through testbed experiments how concurrent communication is affected by location and transmit power (Section 3) and to calibrate our SINR-based channel model [17]. We show that although it is a heuristic, GAPC allows concurrent communication in 73% of cases of optimal.

We then use simulations based on these experimentally-derived models to compare GAPC performance to a non-adaptive, RTS-CTS-based CSMA MAC, and both optimal power control and simple minimum-transmit power MACs. Our optimal power control requires perfect knowledge of all channel state, providing upper bound on performance. On the other hand, a very simple MAC might send at the lowest possible power to maximize channel reuse. Our proposed GAPC uses intelligent power control to enable more concurrent communication based on local knowledge; we compare it to these alternatives.

We compare the performance of GAPC with traditional MAC schemes with direct neighbor communication, and with multi-hop streams of data. When communicating with neighbors, we show that GAPC allows 2.6× more successful receptions than CSMA, and 3× more than CSMA with RTS/CTS (CS-RTS/CTS) Multi-hop communications better represent end-to-end performance. We show that for multi-hop communication, GAPC completes a fixed-size transfer faster as well, with CSMA requiring 1.7× longer, or 1.2× longer with CS-RTS/CTS. 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 Related Work

A great deal of prior work has empirically studied low-power wireless transmission, improving our understanding of the wireless communication and developing models and metrics of performance [1, 4, 6, 7, 10, 16, 19, 21, 22, 23, 25]. However, these studies do not consider the situation where multiple senders transmit packets simultaneously. In fact, a design goal of most current medium-access protocols is to avoid concurrent transmissions, often within a two-hop neighborhood of the sender.

Recent work has begun to relax this assumption and explore the implications of concurrent transmission [8, 9, 15, 17, 20, 24]. In densely deployed wireless sensor networks, concurrent packet transmission is endemic. Whitehouse *et al.* [20] and Son *et al.* [17] were the first to systematically explore the effects of interference and concurrent transmission, in Mica2 and MicaZ motes. However, unlike our focus here, they both study the case where multiple transmitters send to a common receiver. Thorough experimental study provides useful guidelines to interference-aware protocol design and demonstrates the feasibility of concurrent communication. Reis *et al.* introduce two physical layer models that provide effective prediction of the probability of packet delivery under interference from concurrent transmission [15]. These models are based on the RF measurements from real 802.11 testbed. Jamieson *et al.* [8] investigate MAC performance by allowing concurrent transmission in an 802.11 testbed. Recently, Moscibroda *et al.* analytically and empirically study the inaccuracy and inefficiency of protocol design based on graph-based model [13], and analyze the capacity of wireless network with a physical model allowing concurrent communications [12].

Transmission power control plays a key role in interference-aware protocol design by controlling the intensity of the signal and interference strength. Even though there have been extensive research efforts with transmission power control in wireless communication, there are few *empirical* studies that consider transmission power control. Son *et al.* [16] study the effects of transmission power control on

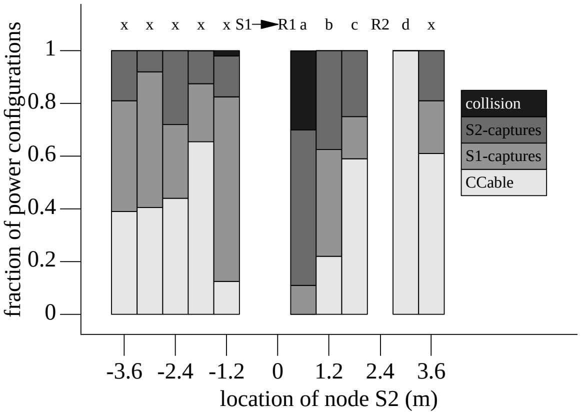

Figure 1. MicaZ experiments measuring CCability with two sender-receiver pairs. S1, R1, and R2 are placed at -0.6m, 0m, and 2.4m as shown above; S2 is measured at each of the 10 locations marked with lowercase letters.

wireless link quality on a real sensor network testbed with Mica2 motes. They propose a power control scheme with link blacklisting to improve link reliability and energy efficiency. Recently Lin *et al.* [11] proposed an adaptive transmission power control (ATPC) protocol based on the empirical measurements from the MicaZ motes with 802.15.4 radios [5] which reacts to the temporal change of the link quality with explicit on-demand feedback packets. However, both these works do not explicitly study the benefits from power control for concurrent communication or propose a MAC protocol.

3 Experimental Evaluation

We first present experimental results with real hardware to show how node location and transmit power allow concurrent communication, capture, or result in collision.

3.1 Methodology

Our testbed experiments follow the methodology of recent studies of concurrent communication [17]. 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 [5] 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.

We vary transmit power and test for success of 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.

3.2 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 5.

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 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 vividly show that control of transmit power is a necessary enabler for successful concurrent communication. It also shows that flexibility in power selection is a function of distance. 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 [17], 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. These results are consistent with previously reported demonstrations of capture [13, 20].

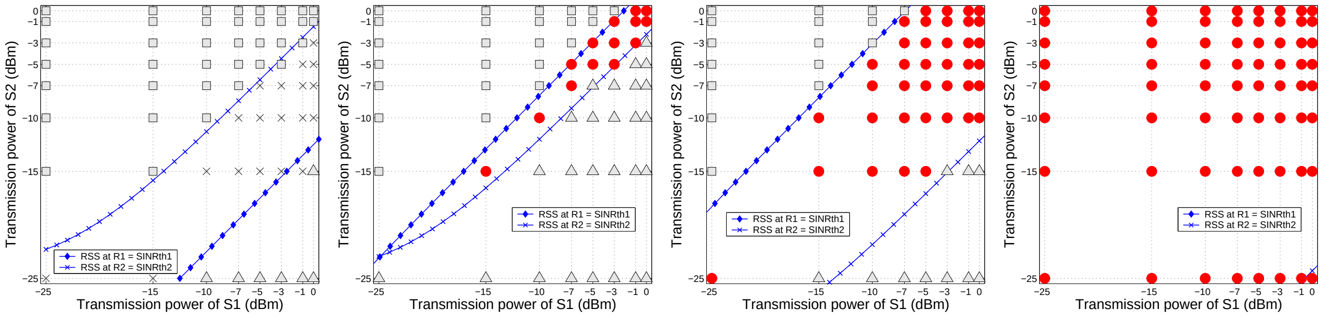

(a) S2 at 0.6m (b) S2 at 1.2m (c) S2 at 1.8m (d) S2 at 3.0m

Figure 2. CCability in the testbed experiment as S2 is moved (presented together with the expectation from simulation with our proposed formula in Section 4). 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.

3.3 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. We next use these experiments to parameterize our SINR model to predict CCability with two concurrent transmitters.

4 Modeling Concurrent Communication

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 [17]. We next evaluate the data from the experiments by comparing them to the mathematical estimation.

We begin by modeling mathematically when concurrent transmissions can occur for the case of two senders and two receivers. For our modeling, we use the *exponential path loss model with log-normal fading* [14, 25]:

$$\begin{aligned} PL(d)_{dB} &= PL(d_0)_{dBm} + 10n\log(d/d_0) + X_{\sigma_{dB}} \qquad (1) \\ P_r(d)_{dBm} &= P_{t_{dBm}} - PL(d)_{dB} \end{aligned}$$

Here P_t and P_r are the transmission and reception power in dBm. The sender-receiver distance is d, and d_0 is the reference distance for path loss *(PL)*. X_σ is the variance in path loss due to multipath fading, modeled as Gaussian random variable with zero mean and standard deviation σ_{dB}. This model defines the path loss and the received signal strength (RSS) at the receiver for a given transmission power level.

4.1 Power Setting for CCability

For concurrent transmission to be possible, the received SINR must be above the threshold for each receiver ($SINR_{\theta r}$ for receiver r):

$$\begin{aligned} S_{11_{dBm}} - 10\log(10^{I_{21_{dBm}}/10} + 10^{N_{1_{dBm}}/10}) &\geq SINR_{\theta 1_{dB}} \qquad (2) \\ S_{22_{dBm}} - 10\log(10^{I_{12_{dBm}}/10} + 10^{N_{2_{dBm}}/10}) &\geq SINR_{\theta 2_{dB}} \end{aligned}$$

For a given distance-based path loss model, such as the one we described in Equation 1, we get the following non-linear inequalities relating the transmission powers of both senders, given a sender x to receiver y distance of d_{xy} and transmission power of $P_t(s)$ for sender s:

$$\begin{aligned} P_t(S1) &\geq PL(d_{11}) + SINR_{\theta 1} + \qquad (3) \\ &\quad 10\log(10^{(P_t(S2)-PL(d_{21}))/10} + 10^{N_1/10}) \\ P_t(S2) &\geq PL(d_{22}) + SINR_{\theta 2} + \\ &\quad 10\log(10^{(P_t(S1)-PL(d_{12}))/10} + 10^{N_2/10}) \end{aligned}$$

We can visualize these non-linear inequalities as regions in a plot where the axes represent the transmission powers $P_t(S1)$ and $P_t(S2)$. The intersection of regions would then indicate when both conditions are satisfied simultaneously, i.e. when concurrent transmissions are possible. From the above equation, we see that the shape of these regions would be primarily determined by the path loss model and the inter-node distances. Figure 3 shows these regions for three different node topologies.

Figure 3(a) shows regions corresponding to the non-linear inequalities for the scenario shown in Figure 4 at 10 m of R1-S2 distance. Each line indicates the sender's optimal transmission power which meets the SINR threshold requirement at its intended receiver with equality. The line with circles shows calculated S1's optimal transmission powers if the S2's transmission power varies between -25 and 10 dBm as shown in the Y-axis. The region to the bottom-right of this curve represents all combinations of transmission pow-

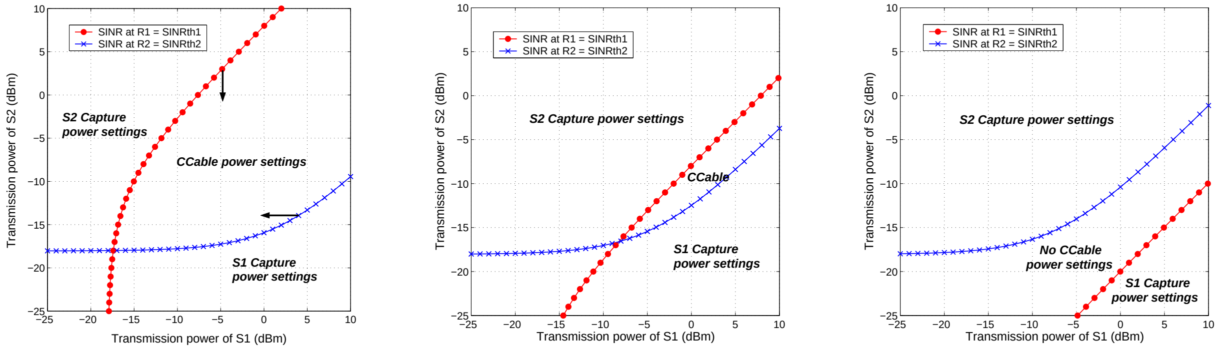

(a) $d_{R1-S2} = 10m$, $lf_1 = 4$, $lf_2 = 2$ (b) $d_{R1-S2} = 4m$, $lf_1 = 1.8$, $lf_2 = 0.8$ (c) $d_{R1-S2} = 2m$, $lf_1 = 1.4$, $lf_2 = 0.4$

Figure 3. The CCable transmission power relationship between S1 and S2. $PL_0 = 45$, $n = 4$, $SINR_\theta = 4$, $X_\sigma = 0$

ers that allow receiver R1 to capture the message. The line with crosses similarly shows S2's calculated optimal powers for different S1's transmission power selections. The region to the top-left of this curve shows all combinations of transmit powers that allow receiver R2 to capture the message. The overlapping region, therefore, shows the combination of transmission powers that allow for concurrent transmission (i.e., these are the CCable power settings).

As shown in the plots in Figure 3, the extent and the existence of the overlapping CCable region depends upon the inter-node distances. In particular, compared to (a), (b) shows a smaller CCable region requiring higher transmit powers as the R1-S2 distance becomes smaller; when the R1-S2 distance is reduced even further in (c), we find that the two regions no longer overlap.

The crossing point of the two lines in Figure 3 provides the optimal S1 and S2's transmission power combination, which is the minimum transmission power setting for CC. We can actually solve analytically for this crossing point (when it does exist) by treating the inequalities from Equation 3 as simultaneous non-linear equations. This yields the following expressions for the *optimal transmission power settings* for S1 and S2:

$$\begin{aligned} P_t(S1) =& PL(d_{11}) + SINR_{\theta 1} + \\ & 10\log(10^{(P_t(S2)-PL(d_{21}))/10} + 10^{N_1/10}) \\ P_t(S2) =& 10\log(10^{(PL(d_{11})-PL(d_{12})+SINR_{\theta 1}+N_1)/10} + 10^{N_2/10}) \\ & - 10\log(10^{-(SINR_{\theta 2}+PL(d_{22}))/10} \\ & - 10^{(PL(d_{11})-PL(d_{12})-PL(d_{21})+SINR_{\theta 1})/10}) \end{aligned} \quad (4)$$

Equation 4 provides the optimal transmission power to use for each sender S1 and S2 without exhaustive trial and error. Optimal power setting consumes minimum energy for concurrent communication causing minimal interference to the network.

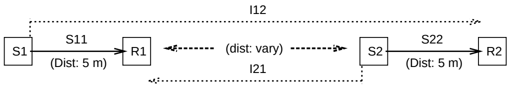

Figure 4. Example scenario with two concurrent packet sender-receiver pairs varying R1-S2 distance.

4.2 Topology Condition for CC

We can get some analytical insight into the impact of topology by deriving a necessary and sufficient condition for CCability. In order to ensure that the simultaneous non-linear equations have a bounded solution, it can be shown that the following topology condition is necessary and sufficient (this condition ensures that the logarithm in equation 4 has a positive argument):

$$\begin{aligned} SINR_{\theta 1} + SINR_{\theta 2} <& \\ & PL(d_{12}) - PL(d_{11}) + PL(d_{21}) - PL(d_{22}) \end{aligned} \quad (5)$$

Adopting the exponential path loss model from Equation 1, this can be written as:

$$SINR_{\theta 1} + SINR_{\theta 2} < 10n(\log(\frac{d_{12}}{d_{11}}) + \log(\frac{d_{21}}{d_{22}})) \quad (6)$$

Let us define the *location flexibility* lf_i for each sender i as the ratio of the distance between a sender and its intended receiver to the distance between a sender and its unintended receiver (i.e., interfered node). Thus $lf_1 = \frac{d_{12}}{d_{11}}$ and $lf_2 = \frac{d_{21}}{d_{22}}$. The lf value indicates the endurance level to the additional interference and noise under concurrent transmission. Depending on the lf value of each sender (the higher the better), the possibility of CC and the area of the CCable second sender location changes. This can be seen in Figure 3.

4.3 CCability with Limited Power Range

We have now shown how to determine optimal transmission power for concurrent transmission: evaluate the topology condition to determine if CCability is possible (Equation 5, and if so, compute the optimal transmission powers with Equation 4). Real hardware, however, has limited control over transmission power in terms of supported range and granularity. If the optimal power computed above is supported, we are done. If not, we next consider *how to adapt to constrained choice of power settings*:

- If either optimal transmission power level is greater than that supported by the hardware, CC is not possible.
- If either one of the optimal transmission powers is lower than supported power range, we set the transmission power of this node to the minimum supported transmission power for that node and calculate the transmission power of the other sender with Equation 4. CC is possible only if the calculated transmission power is within the supported power range.
- If both selected transmission powers are below the supported minimum power levels, we select the one with higher difference between the optimal transmission power and its minimum supported power (let's call this the first sender). It attempts to send at its minimum supported power level, and we compute the other required transmission power accordingly. If this exceeds its range, CC is not possible. Otherwise we use the suggested power for the second sender, or bring it to the minimal supported range if it was lower than what is supported. This is because the increase of the second sender's transmission power level to its minimum supported power range is still less than the increase of the first node's transmission power. Therefore, the first sender can tolerate the increase of the second sender's transmission power level.

The basic rule is that the increase of the same amount of transmission power for both CCable senders from the CCable power level always allows CC at their new transmission power level, if new power levels are supported. This is because the effect from the noise decrease at higher transmission power or higher received power level. Therefore, the same amount of signal and interference increase always ends up with higher SINR at the receiver.

4.4 Validation and Implications of the Model

While this analysis provides a framework suggesting an understanding of CCability, that understanding extends only as far as the model matches the complexities of real world. We next confirm that our experiments (Section 3) match the model well and both model and analysis provide consistent guidance to MAC design.

To compare our experiments with simulation, we use the algorithm defined in Section 4.1 and Section 4.3 to predict the power settings that enable concurrent communication. 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.

More importantly, both the model and experiments demonstrate the two factors that control CCability. First, CCability depends on the *location flexibility*. Higher location flexibility increases the possibility of CC, represented by a greater gap between the two lines in S1 and S2's optimal power plot. Second, CCability depends on the transmission *power flexibility*, which means the range of controllable transmission power (i.e., the minimum and maximum transmission power level). Higher transmission power flexibility improves the CCability by increasing the chance of meeting the required CCable transmission power for each sender. Therefore, we can expect higher CCability due to higher *flexibility* from location and transmission power range. We next use these principles to guide MAC design.

5 MAC protocols for Concurrent Communication

We have established that for the vast majority of topologies where senders have reasonable separation, concurrent communication can be possible. 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 and our analysis and experiments depend on complete information and exploration of all possible power levels.

We next consider four different power control algorithms for MAC protocols. Carrier sense with RTS/CTS (CS-RTS/CTS) represents the current state-of-the-art. We present an *Oracle* algorithm (based on Section 4) 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 specify GAPC here in Section 5.4.

We evaluate these protocols through simulation using the SINR-based concurrent communication model that we validated in Section 4.4. 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 5 shows CCability for each protocol, where black indicates the CCable region, gray shows where one communication or the other is capturable. 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. The x and y coordinates show the locations of each node, in meters. 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.

5.1 Today's practice: CSMA/CA and CS-RTS/CTS

We begin by evaluating a traditional control method. The carrier sense multiple access with collision avoid-

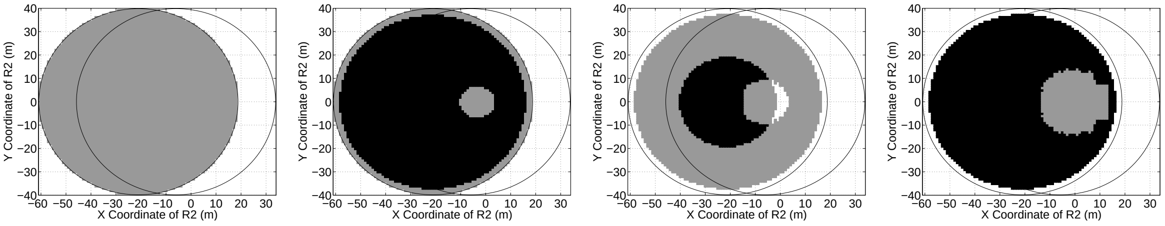

(a) CS-RTS/CTS(Max) (b) Oracle (c) MinPC (d) GAPC

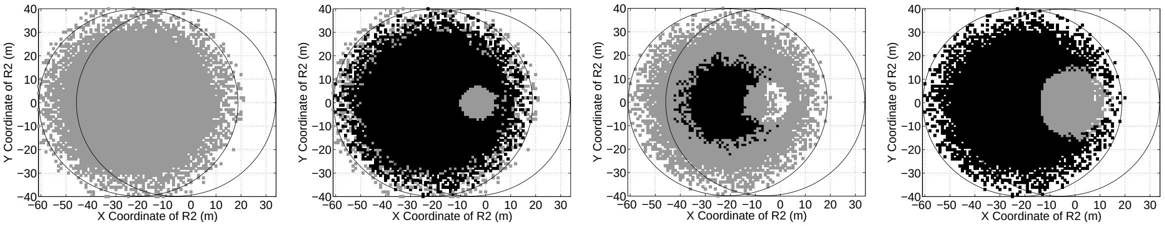

(e) CS-RTS/CTS(Max) (f) Oracle (g) MinPC (h) GAPC

Figure 5. 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.

ance (CSMA/CA) and carrier sense with RTS/CTS (CS-RTS/CTS). Figure 5(a) shows the behavior of a traditional CS-RTS/CTS MAC for concurrent communication. We use the static transmission power for data and control packet transmissions. It is simplest to always transmit RTS/CTS at maximum power to block any potential receivers. As shown in Figure 5(a), this always allows one sender, but never allows concurrent communication. The CSMA/CA has the same performance as CS-RTS/CTS at the given example. However, CS-RTS/CTS is a more conservative medium access control scheme that blocks more space using control packets to avoid packet collisions at the receiver.

5.2 A Upper Bound on Performance: 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 (Section 4.1). While gain can be observed in a network and its variance estimated, this estimate is never perfect and so this algorithm is not achievable in practice. However, it establishes 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 5(b) shows sample results with the Oracle. Compared to CS-RTS/CTS (Figure 5(a)) 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*.

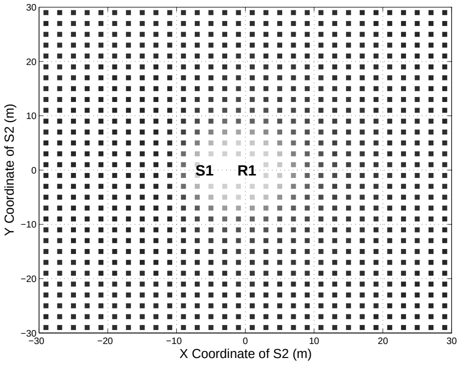

Figure 6. 2D simulation results with optimal transmission power settings. $S1 = (-6,0)$, $R1 = (0,0)$, $S2$ location varies within the presented coordinate space

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 (Figure 8).

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.

Greater degrees of freedom: The above simulations fix

S2's location and move R2 to all possible locations. To confirm that this simplification does not bias our results, we also carried out simulations where we move both S2 and R2 using the Oracle algorithm for power control.

Figure 6 shows the average CCability we observed for each S2 location. Each S2 point in this figure presents the mean probability of successful concurrent communication to all possible other R2 locations. The darkness of each block represents the mean probability, with black means 100% CCability, white meaning no concurrent communication is possible. Over the whole field, 89% of R2 locations are CCable when S2 is placed in the field of -30 to 30m. This simulation confirms that results in Figure 5 are representative of other topologies: CCability is almost always possible except when sources and destinations are nearly completely co-located. These simulations confirm the potential for much better spatial channel reuse, provided MAC protocols can make good power control decisions.

5.3 Exploiting Power Control and Channel Capture: MinPC

Ignoring interference, we maximize spatial reuse by always sending at the minimum power that will reach the intended receiver. Our *minimum power and capture* protocol (*MinPC* for short) selects the minimum power and employs MAC-level channel capture [20]. We assume nodes can track the gain needed to reach their neighbors. 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 valid assumption for slow-fading environments with little mobility as are typical in today's sensornets [16].

Figure 5(c) shows a moderate size region where concurrent communication is possible, 20% of the total area in this case. In addition, the large gray 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. Comparing this to CS-RTS/CTS demonstrates the advantage of channel capture over communications prohibition.

The penalty of allowing concurrent communication is the small white crescent region where transmission powers are evenly matched at the receivers, resulting in collisions without capture. With CS-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.

However, it is important to understand *why* traffic is CCable in this scenario. In theory, sending at the minimum transmission power cannot tolerate any level of interference. That means we should not find any CCable locations with MinPC when any other concurrent transmitter is within range. However, in practice, real hardware can be set only at discrete power levels within the supported transmission power ranges, resulting in some extra which provides some level of protection to noise. (We use discrete levels at 1 dBm increments in this simulation; the CC2420 radio provides slightly coarser power control capability.)

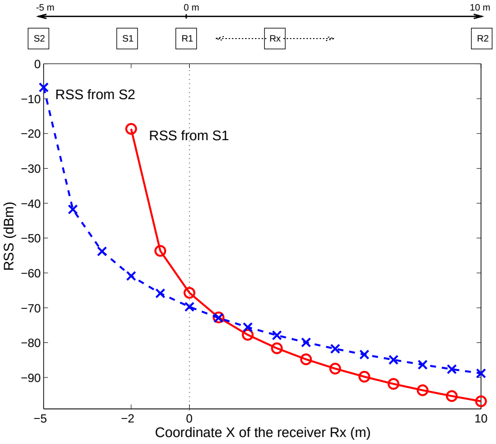

Figure 7. Comparison of RSS from S1 and S2 at hypothetical receiver Rx at location $(x,0)$. $S1 = (-2,0)$, $R1 = (0,0)$, $S2 = (-5,0)$, $R2 = (10,0)$, $n = 4$, $SINR_\theta = 4$, $X_\sigma = 0$

Evaluation of individual node power settings (omitted due to space) shows that the CCable region here is possible only because of our simulation of a limited number of discrete power levels and a fixed minimum transmission power (-25 dBm). CCability within the communication range of S1 comes from this *unintended* extra power allocation due to these hardware limitations.

Thus we conclude that, while minimum power transmission can improve overall spatial reuse by giving each transmitter a smaller "footprint", it is insufficient to promote concurrent communication. While the Oracle algorithm has perfect global knowledge to jointly adjust both concurrent transmission powers, it is not clear how to provide a distributed algorithm to do the same in practical situations. Next we propose GAPC, a simple scheme that gets significant benefit from only minimal information.

5.4 Gain-Adaptive Power Control: GAPC

While unintentional extra power in MinPC provides some buffer against noise, concurrent transmission provides strong sources of interference that limit the chance of concurrent communication. This problem is particularly noticeable in Figure 5(c) as the large gray torus surrounding the black CCable crescent. Compared to the Oracle algorithm, MinPC is succeeding at a few of the many opportunities for viable concurrent communication. The challenge here is the coupling between the two sender-receiver pairs, since the increase in transmission power at one sender also increases the noise at the unintended receiver.

To get around this coupling, we next show that a small *boost* in transmission power that is *in inverse proportion to the minimum transmission gain* allows senders to typically overcome this noise. We call this algorithm *Gain-Adaptive Power Control (GAPC)*. Since the gain is roughly proportional to distance, this means that short-distance transmissions get large boosts while longer transmission gets relatively less gain. We describe our intuition and the specific algorithm next.

5.4.1 Intuition for GAPC

Our intuition for GAPC comes from observations in prior theoretical and simulation-based studies [13, 18], which found that shorter distance communication is often overwhelmed by interference from longer transmissions.

Figure 7 illustrates the intuition for inverse-gain boosts with an example where S2-R2 communicate over-the-heads of S1-R1. This figure compares the RSS from each sender at all possible locations on the line topology with best possible (Oracle-determined) transmission powers ($-18.67\ dBm$ for S1 and $-6.77\ dBm$ for S2). In this case, both R1 (at 0 m) and R2 (at 10 m) receive stronger signal strength from its intended sender (S1 or S2), allowing them to meet their SINR threshold and capture their packets. We can also have the following two observations: First, sender-receiver pair at a shorter distance requires a larger boost in power to overcome interference because it is transmitting at higher power and has a much larger RSS. By contrast, longer-distance exchanges benefit from smaller changes. In addition, this example shows why it is surprisingly possible to transmit over-across another sender-receiver pair and have both communications received: it is because the different decay rates in logarithmic path loss of wireless propagation allows concurrent capture at different spatial locations.

We next map this intuition into a practical heuristic for distributed control.

5.4.2 Details of MAC design with GAPC

We showed that the Oracle can almost always accomplish concurrent communication (Section 5.2), and that this is because of logarithmic path loss (Section 5.4.1). A full GAPC MAC protocol requires three algorithms: channel capture, to allow the MAC to switch to stronger packets; power selection, to match a good power to the receiver; and transmission control, to coordinate one or more senders. MAC modifications to support channel capture have been described previously [20], so we next describe our two new algorithms for power selection and power-aware transmission coordination.

Power selection: In GAPC we allocate extra power for each communication to improve resilience to the interference from simultaneous transmissions, hoping to increase the possibility of concurrent communication. If we define P_{max} as the hardware-defined maximum possible transmission power, and $P_{S,R}$ as minimum power needed for source S to reach receiver R, then GAPC selects extra power boost ε as:

$$\varepsilon = \phi(P_{max} - P_{S,R}) \qquad (7)$$

Where ϕ defines the fraction of remaining (spare) power to allocate to a transmission. Large values of ϕ will quickly assign all headroom to transmissions and will increase the bonus given to shorter links. We varied ϕ and found that moderate values (0.3 to 0.7) provided the best levels of CCability (values that near 0 provide no boost and approach MinPC performance, while values approaching 1 always operate at maximum power). We adopt $\phi = 0.5$ as a reasonable, robust choice. In general, higher ϕ improves the CCability for concurrent transmissions over short distances. But too large a ϕ decreases the CCability by reducing the number of transmission attempts because each takes a larger footprint.

Figure 5(d) evaluates gain-adaptive power control with $\phi = 0.5$. We see that this approach comes very close to optimal: concurrent communication is possible with the receiver in nine-tenths of CCable area to the Oracle algorithm, much closer than MinPC, which has one-fourth of CCable area. These results suggest that gain-adaptive power control is a practical scheme that gets a significant fraction of optimal performance.

The cost of gain-adaptive control relative to the Oracle can be seen in two locations. The moderate-size gray 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 torus around the edge of the Oracle cannot be reached with gain-adaptive control because of slightly higher interference from the S1-R1 pair.

Transmission control: Although our goal is to allow concurrent transmission, we still must control access to the medium to avoid unsynchronized access or overwhelming interference from many concurrent transmitters.

Our medium access control consists of two rules. First, potential senders monitor the channel to observe ongoing transmissions. Each node records the source and destination of any transmissions that begin. If we define an ongoing exchange as from S1 to R1, a new sender S2 considering transmission to R2 can evaluate if concurrent transmission is possible. If the S2 has R1 as its neighbor, it can consider the relative gains required for it to reach R1 and R2. If the S2-R1 gain is smaller than S2-R2, then it knows no concurrent transmission is possible (since it will drown out the ongoing S1-R1 transmission) and defers channel access. On the other hand, if the S2-R1 gain is smaller it can proceed because of location flexibility (Section 4.2).

We use the following specific rule to decide when to attempt a second transmission over an ongoing one:

$$PL(d_{S2-R1}) \geq PL(d_{S2-R2}) + 10n\log C$$

Although nodes know only path loss, not actual distances, this rule is equivalent to $d_{S2-R1} \geq Cd_{S2-R2}$. The constant factor C controls the amount of headroom required. We use $C = 1.5$ in our simulations, based on tests with candidate values. Comparisons of C are omitted due to space.

Second, we allow at most two concurrent transmissions. As described above, senders can track one current transmission directly. They may estimate multiple on-going transmissions by tracking sustained changes of channel energy above that of the first transmission. Unlike 802.11 with RTS/CTS, we only observe transmissions at the sender. Note that estimating level of concurrent transmissions from channel energy in inaccurate with distant receivers; our goal is simply to ensure a relatively empty channel, so an exact count of transmissions is not required, and little is lost if multiple distant transmissions are assumed to be one nearer transmission.

6 Quantifying the Advantages of Concurrent Communication

While we have just shown the *potential* opportunities for concurrent communication and how it lead us to GAPC, we

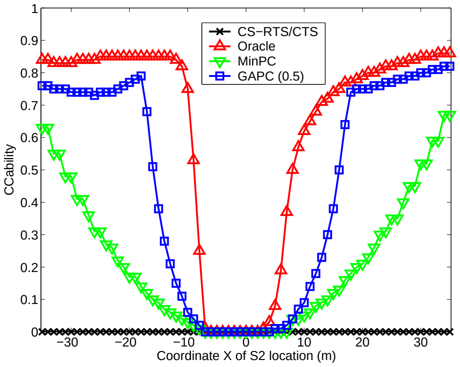

Figure 8. Comparison of CCability with limited power levels (-25 dBm to 0 dBm) for five MACs. $S1 = (-6,0)$, $R1 = (0,0)$**,** $S2 = (x,0)$**,** $R2$ **at all possible locations.**

next quantify the performance advantages of GAPC for communication compared to other MAC protocols. We do this in four steps, first more rigorously quantifying opportunities for concurrent transmission with two senders. We then consider communication between random neighbors in a large, dense, multi-hop network. Next we study the effects of power flexibility required for GAPC on multi-hop communications. Finally, we compare dense networks with a large number of random multi-hop data flows to evaluate end-to-end performance in a complex network.

6.1 Two Concurrent Senders

We begin by quantifying the opportunities for concurrent communication with two concurrent senders. Here we quantify the level of spatial reuse for the four MAC protocols we consider (RTS/CTS, Oracle, MinPC, and GAPC). We consider a two sender case to allow comparison of all possible topologies.

CCability: In Section 5, we compared the four MAC protocols we fixed three of the four nodes (there S1, R1, and S2) and varied the location of R2. From Figure 5 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 5(a)). We define this ratio as the *CCability*.

We next take this evaluation a step further by allowing the location of *both* S2 and R2 to vary. Figure 8 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 in the 2-D plane. Each point on the figure represents the fraction of R2 locations that allow concurrent communication for a given S2 x position.

Figure 8 shows that CS-RTS/CTS provides no spatial reuse of the channel through concurrent communication; this result is consistent with its design goal of preventing any possible interference. By contrast, the top line shows the Oracle algorithm, with ideal power settings: when S2 just a bit distant from R1 (say, outside the range of -8 to 6m), 70–85% of R2 locations allow concurrent transmission. This ideal is what we strive for.

Shifting to a channel-capture-based MAC (MinPC) provides considerably greater opportunity for concurrent communication because with lower transmit powers, distant nodes can send concurrently. We see this as CCability in MinPC rises to about 65% when S2 is at ±36m—as far from R1 as possible. Transmitting at lower power therefore improves spatial reuse by reducing co-channel interference, but it cannot allow concurrent communication when senders are near each other.

Finally, GAPC comes relatively close to the best possible Oracle result. We find that it allows CCability in about three-quarters of the optimal (GAPC vs. Oracle) averaged over all S2 locations, and and it reaches 95% of in regions with sufficient source separation (when S2 is left of -15m or right of 15m).

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.

Possible communication: We can use this same methodology to evaluate not just opportunities for concurrent communication, but some kind of successful communication by either one or both parties. In Figure 5, successful communication by exactly one party occurs in the light grey regions and CCability in black regions, so we next compare that area against the total transmission range.

When we compare the probability of having one or two successful packet receptions the oracle always gets one *or two* packets through. Figure 8 shows where two are possible, while in region of impossible concurrency it gets one packet through. By comparison MinPC and GAPC show similar, performance, getting one or two packets through 77–88% of the time (mean: 85%).

This high success rate is because even when concurrency is impossible, most often one sender captures the channel. The 12–23% gap between practical and Oracle represents lost capacity due to collisions. For GAPC, most of this loss occurs at near the edge of maximum communication range where hardware limitations prevent GAPC from boosting power enough to exceed interference.

Conclusions: From our evaluations of two concurrent senders we have shown that GAPC gets the about three-quarters of the opportunities for concurrent transmission, significantly better than MinPC. We also have shown that even when concurrent communication is impossible, most of the time at least one concurrent transmitter is successful. We conclude that GAPC has significant opportunity to improve performance provided its overhead is minimal.

6.2 Neighbor Communication in a Dense Network

We next turn to simulating a large, densely deployed multi-hop network. Section 6.1 evaluated opportunities for concurrent communication by exhaustively exploring many

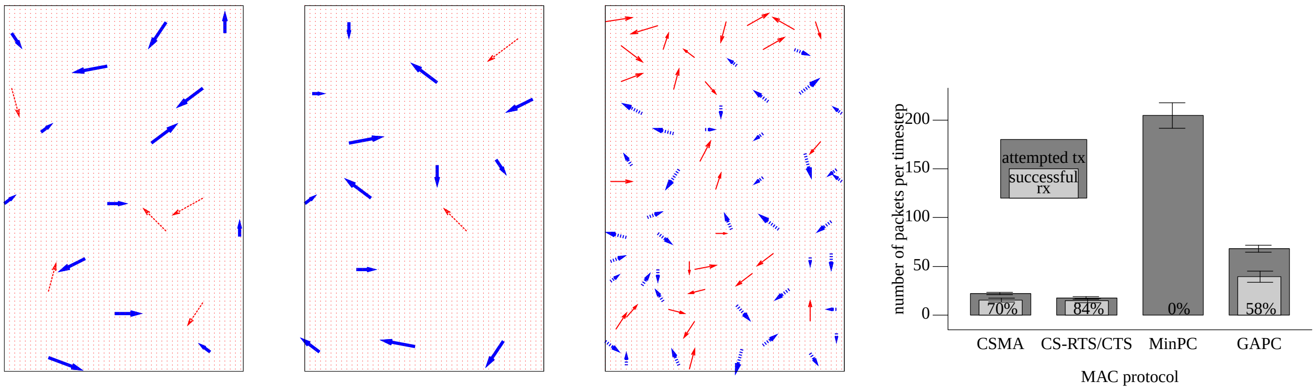

(a) CSMA (b) CS-RTS/CTS (c) GAPC (d) Performance comparison

Figure 9. Successful (thick solid arrows) and failed (thin dotted arrows) communication attempts for a representative timestep of neighbor communication, for three protocols (a, b, and c), and statistics for all steps of all simulations (d).

two-pair topologies. Here we instead consider random communications in a field of nodes to show that these opportunities persist even in the face of cumulative effects of noise from many neighbors.

Methodology: We consider the following topology. We deploy 4232 nodes in a uniform grid at a 26 m spacing over a 2400×1200m area. We select radios that allow transmission powers to vary between −25 and 25 dBm in 1 dBm steps. This range is 25 dBm greater than the 802.15.4 radios on MicaZ motes, and the degree of control is somewhat greater. With no interference, and the propagation parameters we describe next, this power range allows transmit distances of 8 to 207 m. In the grid, each node can cover 7 hops at maximum power, giving nodes up to 192 neighbors, making this a very dense network. Our path loss parameters (equation 1) are $n = 3.5$, $PL_{(}d_0) = 35$, $X_\sigma = 0$. Based on the observation with MicaZ motes, we adopt $SINR_\theta = 2dB$ for a channel capture and ambient noise of $N = -95$ dBm. We use $\phi = 0.5$ for extra power allocation for GAPC, and −93 dBm for carrier sensing threshold (CS_θ).

We perform simplified version of time-slot based simulations. For each time slot, all nodes transmit a packet with probability $PR_{tx} = 0.05$ to a randomly chosen neighbor. (At maximum power this probability implies an offered load of 9.6 at each node.) In each slot we randomly order senders to simulate randomized access through a slotted contention window. Each sender calculates the interference level based on the selected transmission order. We simulate 100 timesteps, each a round of communication from many nodes. For these 100 simulations, each one has a different communications pattern (from a different random seed).

We compare four MAC protocols: CSMA, CS-RTS/CTS, MinPC, and GAPC. We consider both CSMA and CS-RTS/CTS to separate out the effects of 2-hop silence from RTS/CTS.

We make no claims that this topology and traffic model represent realistic traffic—they do not. The are designed instead to provide an easily understood abstract model of a multi-hop network, allowing relative comparisons of MAC protocols in the face of concurrent communication. In addition, we select the size to avoid edge effects, and density to capture the densest portions of potential deployments.

Evaluation: Figure 9 presents a representative example of one timestep of one simulation. We show all successful and unsuccessful transmissions as thick, solid and thin, dotted arrows. Figure 9(d) summarizes performance over all simulations, including percentage successful reception and standard deviations.

One limitation of this simulation is that GAPC senders choose only neighbors reachable within 20 dBm transmit power or less, while others consider the full 25 dBm transmit power. However, GAPC boost transmissions to full (25 dBm) power if necessary to overcome interference. We quantify the effects of power management in the following section (Section 6.3).

As Figure 9 shows, GAPC attempts many more communications, and produces more successful receptions, than existing MAC protocols that strive to prevent concurrent communication to provide very high success. On average GAPC transmits about 3.1 times more packets and allows about 2.6 times more concurrent communications than CSMA (GAPC: 68 attempts, 39 successful; CSMA: 22 attempts and 15 successful). MinPC, by contrast, makes many attempts but none get through. This result shows the importance of a GAPC's adaptive power boost, something we return to in the next section. CSMA attempts many fewer transmissions, but its success rate is higher for what it sends (70% compared to GAPC's 58%). CS-RTS/CTS attempts even fewer transmissions since a two-hop neighborhood is blocked, with a mean rate of 15 attempts per slot, but blocking transmissions around the destination raises the success rate to 84% (13 per slot), one third of what GAPC can accomplish.

Overall, GAPC provides 2.6 or 3 times the overall throughput because of greater spatial reuse, but this through-

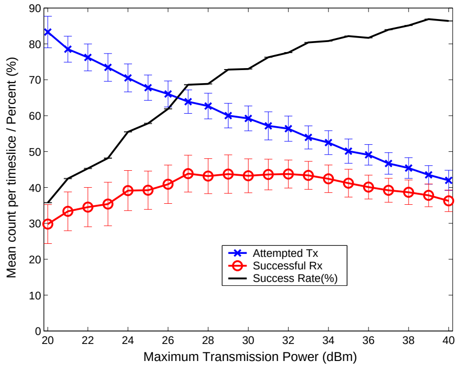

Figure 10. The effect of reserve power on GAPC

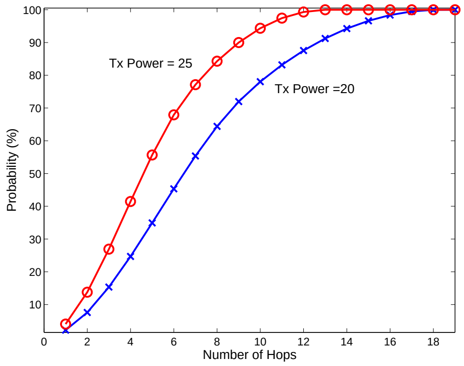

Figure 11. CDF of shortest-path hop-length for all pairs of nodes with two different neighbor selection powers, 20 and 25 dBm.

put comes at the cost of lower per-attempt reliability. We expect that a MAC-level ARQ scheme can improve reliability; we return to this question when we consider multi-hop routing in Section 6.4.

6.3 Effects of Power Reserve

GAPC requires a power reserve to allow it to boost power to overcome interference from a concurrent transmission. We must leave room for this reserve by doing routing at less than the maximum possible power. This spare transmission power is not wasted, but used only only to compensate for co-channel interference. We next evaluate how much power reserve is needed, and what effect his reserve has on multi-hop paths lengths.

To study the effect of different amounts of power reserve we use the same simulation methodology as our neighbor communications (Section 6.2), but we select neighbors with a 20 dBm power budget. We then vary the maximum supported transmission power from 20 to 40 dBm (in this experiment only, effectively improving our radio hardware). Higher maximum transmission powers provide a larger reserve to overcome interference.

Figure 10 shows the number of allowed concurrent transmissions, successful transmissions, and success rate as the amount of power reserve grows. and percentage for success rate of transmitted packet.

This experiment shows the importance of reserve power. First, we observe that additional reserve power consistently improves the rate of transmit success, from 35% with no reserve to 85% with 20 dBm reserve (the rising solid line with tiny error bars). This observation confirms the importance of power boost over minimum to overcome noise, confirming problems in deploying MinPC with more than two transmitters (Figure 9(d)). Second, we can observe the effect of greater power on reducing spatial reuse, as seen in the reduced number of transmission attempts from 84 packets/timestep at 20 dBm to 40 packets/timestep at 40 dBm (the falling line with X markers). This result quantifies the importance of power control to maximize spatial reuse. Finally, we see that these two effects interact, trading off in ultimate performance shown successful reception rate (the curve with circles). Successful reception varies over a range of from 30 to 40 packets/timesteps, slightly greater than the random effects of topology and transmitter variation (shown by standard deviations around ±5), with best success rate at 27 dBm maximum power (7 dBm reserve). Finally, this evaluation confirms that our choice of 5 dBm reserve power in other simulations is reasonable.

To provide a fair comparison with alternate MAC protocols, we cannot simply "create" extra power. GAPC reserve power by attempting communication only with closer neighbors, saving some power for boost. The effect of this is shorter possible transmission range. To quantify this, we compute shortest paths between all pairs of random topologies for our simulations, considering neighbors reachable with at most 20 or 25 dBm transmission power. The stronger power (25 dBm) corresponds to neighbors reachable with traditional MACs, while the weaker power (20 dBm) is the neighborhood GAPC considers for communication to allow 5 dBm of reserve power. From Figure 11, we can see that GAPC does require slightly longer paths; on average 2 extra hops in this network.

The prior section showed that GAPC has much more successful communication, but here we see that it requires slightly more hops. We therefore next consider multi-hop routing to evaluate this trade-off with end-to-end performance.

6.4 Concurrent Multi Hop Communications

We use the same dense network topology and parameters as the neighbor communication study (Section 6.2). However, we replace the traffic model with multi-hop communication. For each simulation, we randomly select 20 pairs of source and destination nodes within the network. Each source sends one packet to the destination. We route the packets using Dijkstra's shortest path algorithm, setting the cost of every available link within the communication range to 1 regardless of the distance.

First, to study MAC performance alone, we first simulate communication without collision recovery, modeling traffic without any link-layer ARQ. For each time slot, we schedule one packet transmission following the communication route for each data flow. Regardless of the success of communication, we schedule a packet transmission from the next hop for next time slot. While unrealistic, this model isolates MAC performance.

Figure 12 plots one example of the successful and failed

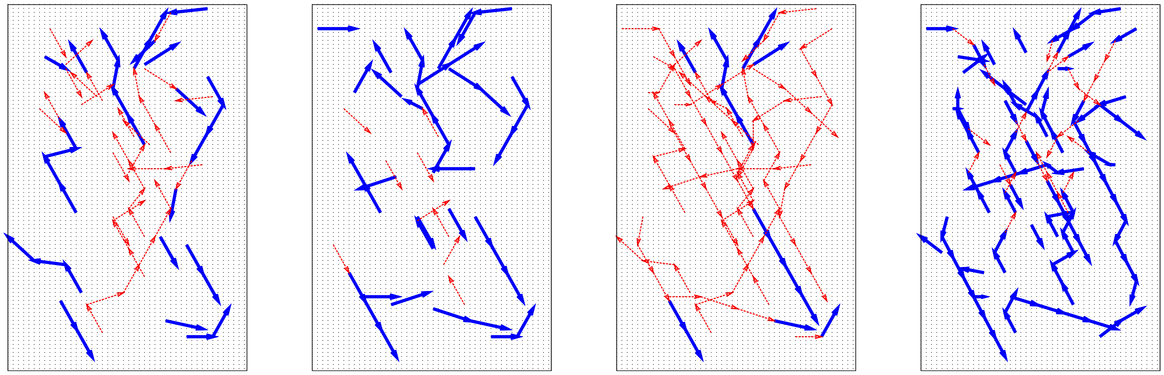

(a) CSMA (b) CS-RTS/CTS (c) MinPC (d) GAPC

Figure 12. Successful (thick solid arrows) and failed (thin dotted arrows) communication attempts for multi-hop concurrent communications in dense networks without a link level ARQ.

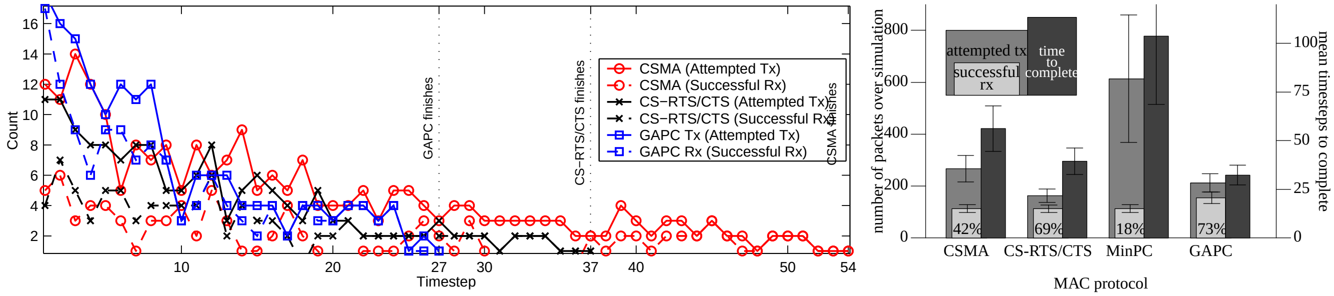

Figure 13. Packet transmission attempts and successful receptions per timestep for multi-hop communication with simple ARQ, with three MAC schemes (a), and performance comparison (b).

communication from using different MACs, out of 10 simulations with various random flows. On average, GAPC attempts 1.3× and 1.5× more transmissions than CSMA and CS-RTS/CTS. One the other hand, MinPC attempts 1.3× more than GAPC. For this effort GAPC gets 1.7×, 1.6×, and 3.2× more concurrent communications over CSMA, CS-RTS/CTS, and MinPC. These results generally confirm that our previous evaluations apply even with multi-hop flows: GAPC provides much better spatial reuse compared to more conservative CSMA and CS-RTS/CTS, and MinPC which allows concurrent transmissions without traffic control.

To compare end-to-end performance, we next add reliability by considering a link-level ARQ scheme in the MAC. Each MAC retries a failed transmission with a simple back-off algorithm, reducing PR_{tx} by half for the first collision and to a quarter its value for subsequent retries. Retransmission effort stops only when it successfully deliver a packet to the next hop. We carried out 10 simulations using the same random flow locations and topologies as the simulations without a link-level retransmissions.

Figure 13 shows the complete execution of a representative simulation (left), and statistics over all simulations (right). Adding retransmissions, we shift our evaluation from packets per timestep, to the end-to-end metric of packets to complete delivery of all data in the network.

First, we observe that GAPC completes sooner than alternatives. In the example, GAPC finishes after 27 timesteps, compared to 37 for CS-RTS/CTS and 54 for CSMA. These trends hold across all simulations, with CSMA, CS-RTS/CTS, and MinPC requiring 1.7×, 1.2×, and 3.2× longer to complete sending data compared to GAPC.

However, end-to-end scenario with ARQ changes some underlying results. First, GAPC's success rate is much better here (73%) compared to neighborhood communication (58%). This change is because the level of traffic in this network is lower, on average two packets per neighborhood per timestep, a good match for GAPC. The lower performance of CSMA and CS-RTS/CTS are because with shortest-path routing they will always strive to use the longest possible next hop, yet transmissions at extreme distances are very vulnerable to even slight interference. Also, sender-side carrier sensing only approximates interference at the receive. GAPC avoids this problem for two reasons. First, it selects hops with less than maximum power and then boosts power if necessary to overcome concurrent communication. Second, it transmits at lower power, thereby getting better spatial reuse and *causing* less interference. We can also observe GAPC's shorter-hop selection in the total number of successful packets sent. Each packet will be sent successfully once over each hop in the shortest-hop path (variance in this number is due to averaging over different topologies). GAPC requires a mean of 155 successful packets to deliver all data, while the other protocols each require 113. This difference reflects GAPC's use of more, shorter hops; it matches the two extra hops we reported in Section 6.3.

This experiment confirms the importance of RTS/CTS in multi-hop communication (compare time and attempts of CS-RTS/CTS compared to CSMA). Because we ignore the cost of the RTS/CTS control messages, these results overly favor RTS/CTS. GAPC, however, is more efficient, because of improved spatial reuse.

Finally, these results can suggest potential energy costs of these protocols. Ignoring control overhead, GAPC sends 30% more packets than CS-RTS/CTS. However, counting control overhead would triple CS-RTS/CTS' number of packets. In addition, GAPC transmits at lower powers, so range-dependent energy costs will be lower. MinPC also sends at lower power, but its low success rate makes its performance not competitive.

We conclude that GAPC outperforms traditional MAC schemes for end-to-end packet delivery from improved spatial reuse obtained from proper power control with channel capture in mind.

7 Conclusions

In this paper we have presented the first effort to quantify and utilize the the opportunity for concurrent communication in low-power wireless networks. Our experimental results show that complete packet collision is a rare event, and concurrent communication is almost always possible given appropriate selection of transmission power. We then model the conditions for concurrent communication that matches our empirical results, and provides a reliable background for simulation-based study.

Based on the insights obtained from experimental study, we propose a practically implementable MAC scheme, GAPC. We compare the performance of GAPC with traditional MAC schemes for single and multi-hop concurrent communication in dense networks. Improved spatial use from GAPC allows significant improvements to network throughput in dense wireless networks. Our study shows embracing concurrent communication through power control and channel capture can significantly improve wireless throughput.

We are currently working on a full implementation of GAPC for motes. Channel capture is well supported on CC1000 radios [20], but more challenging on 802.15.4 radios where only strongest-first capture is supported.

8 References

[1] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b mesh network. In *ACM SIGCOMM*, pages 121–132, Aug 2004.

[2] V. Bharghavan, A. Demers, S. Shenker, and L. Zhang. MACAW: A media access protocol for wireless lan's. In *ACM sigcomm*, pages 212–225, London, UK, Sep 1994.

[3] S. Biswas and R. Morris. Exor: Opportunistic multi-hop routing for wireless networks. *ACM SIGCOMM Computer Communication Review*, 35,4:133–144, Oct 2005.

[4] A. Cerpa, J. L. Wong, L. Kuang, M. Potkonjak, and D. Estrin. Statistical model of lossy links in wireless sensor networks. In *IPSN*, Apr 2005.

[5] Chipcon. CC2420 2.4 GHz ieee 802.15.4/ zigbee-ready RF transceiver. `http://www.chipcon.com/files/CC2420_Data_Sheet_1_4.pdf` .

[6] D. Couto, D. Aguayo, J. Bicket, and R. Morris. A high-throughput path metric for multi-hop wireless routing. In *IEEE Mobicom*, pages 134–146, Sep 2003.

[7] D. Ganesan, D. Estrin, A. Woo, D. Culler, B. Krishnamachari, and S. Wicker. Complex behavior at scale: An experimental study of low-power wireless sensor networks. Technical Report CS TR 02-0013, UCLA, 2002.

[8] K. Jamieson, B. Hull, A. Miu, and H. Balakrishnan. Understanding the real-world performance of carrier sense. In *SIGCOMM'05 Workshop*, pages 52–57, Aug 2005.

[9] A. Kochut, A. Vasan, A. Shankar, and A. Agrawala. Sniffing out the correct physical layer capture model in 802.11b. In *IEEE International Conference on Network Protocols (ICNP)*, pages 252–261, Oct 2004.

[10] D. Lal, A. Manjeshwar, F. Herrmann, E. Uysal-Biyikoglu, and A. Keshavarzian. Measurement and characterization of link quality metrics in energy constrained wireless sensor networks. In *IEEE Globecom*, pages 446–452, Dec 2003.

[11] S. Lin, J. Zhang, G. Zhou, L. Gu, T. He, and J. Stankovic. ATPC: Adaptive transmission power control for wireless sensor. In *Sensys*, pages 223–236, Nov 2006.

[12] T. Moscibroda. The worst-case capacity of wireless sensor networks. In *IPSN*, pages 1–10, Apr 2007.

[13] T. Moscibroda, R. Wattenhofer, and Y. Weber. Protocol design beyond graph-based models. In *HotNets-V*, pages 25–30, Nov 2006.

[14] T. Rappaport. *Wireless Communications: Principles and Practice*. Prentice Hall, 1996.

[15] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan. Measurement-based models of delivery and interference in static wireless networks. In *ACM Sigcomm*, pages 51–62, Sep 2006.

[16] D. Son, B. Krishnamachari, and J. Heidemann. Experimental study of the effects of transmission power control and blacklisting in wireless sensor networks. In *IEEE SECON*, pages 289–298, Oct 2004.

[17] D. Son, B. Krishnamachari, and J. Heidemann. Experimental study of concurrent transmission in wireless sensor networks. In *ACM Sensys*, pages 237–250, Nov 2006.

[18] 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.

[19] K. Srinivasan and P. Levis. RSSI is under appreciated. In *Emnets*, May 2006.

[20] K. Whitehouse, A. Woo, F. Jiang, J. Polastre, and D. Culler. Exploiting the capture effect for collision detection and recovery. In *IEEE Workshop on Embedded Networked Sensors (EmNetS-II)*, pages 45–52, May 2005.

[21] A. Woo, T. Tong, and D. Culler. Taming the underlying challenges of reliable multihop routing in sensor networks. In *ACM Sensys*, pages 14–27, Nov 2003.

[22] J. Zhao and R. Govindan. Understanding packet delivery performance in dense wireless sensor networks. In *ACM Sensys*, pages 1–13, Nov 2003.

[23] G. Zhou, T. He, S. Krishnamurthy, and J. Stankovic. Impact of radio irregularity on wireless sensor networks. In *ACM MobiSys*, pages 125–138, Jun 2004.

[24] G. Zhou, T. He, J. Stankovic, and T. Abdelzaher. RID: Radio interference detection in wireless sensor networks. In *IEEE Infocom*, pages 891–901, Mar 2005.

[25] M. Zuniga and B. Krishnamachari. Analyzing the transitional region in low power wireless links. In *IEEE SECON*, pages 517–526, Oct 2004.