

The Shadowing Phenomenon: implications of receiving during a collision

Alec Woo Kamin Whitehouse Fred Jiang Joseph Polastre David Culler

{awoo,kamin,xjiang,polastre,culler}@cs.berkeley.edu

Department of Electrical Engineering and Computer Science

University of California, Berkeley

Berkeley, CA 94720-1776

UCB Technical Report: UCB//CSD-04-1313

Abstract

Many wireless networking protocols in use today assume that a connectivity implies interference. While this assumption is usually valid, it does not hold with many radios, possibly including those used in 802.11 and 802.15.4 systems. In systems where a node's interference range is less than its connectivity, receivers may recover the strongest-power packet in a collision, which we call the *shadowing phenomenon*. This phenomenon changes the semantics of many collision avoidance schemes such as the use of Acks and RTS/CTS. It also changes network dynamics and connectivity for higher-level protocols such as multihop routing, neighbor selection, and clustering. Techniques can exploit this phenomenon to detect collisions, recover packets, and provide new information to network stacks that can be used to design future protocols.

1 Introduction

The notion of a packet collision is a piece of conventional wisdom that pervades wireless networking: when two transmitters, A and B , both have connectivity to a receiver C and transmit simultaneously, a collision occurs, the data is corrupted and neither packet is received correctly. This notion that collisions result in transmission failure is used in many protocols to determine backoff times, adjust transmission rates, achieve fairness, or improve spatial reuse, power consumption, or bandwidth. That is, many protocols assume that *connectivity* implies *interference* in the case of collisions. We call this the *connectivity-interference assumption*, a principle that lies at the heart of many of the most important wireless networking trade-offs.

To our surprise, for a significant class of radios that are widely used, the connectivity-interference assumption does not hold. The receiver will sometimes correctly receive one of the two messages in a collision, in what we call the *shadowing phenomenon*. In fact, with certain modulation schemes, such as FSK and possibly PSK, the connectivity-interference assumption may only hold for as few as 10% of a node's connected neighbors. This phenomenon is not a newly observed phenomenon; it is in fact well-known among those familiar with the radio

physical layer. The contribution of this paper is to show that shadowing is very common, that many current networking protocols do not account for it, and that several new techniques can be used to exploit it.

One simple way to exploit shadowing is to enable the receiver to detect and recover from collisions. It is also possible to recover portions of the corrupted packet, such as the header fields. The basic approach borrows an idea from CDMA and continues scanning for start symbols while receiving an incoming packet. Thus, with a modest change to the link-level receiver processing, it is possible to detect and recover from collisions when the stronger packet arrives during reception of a weaker packet, called *stronger-last*. However, when the stronger message arrives first, called *stronger-first*, it will be received without any changes to receiver processing because the weaker packet will not even be noticed until the first message is complete.

These two types of shadowing both have far-reaching, yet different implications on networking protocols. The ability to recover useful data in *stronger-last* shadowing can be extremely valuable in multihop routing, network flooding, or any protocol with correlated transmissions in which collision would be extremely common. On the other hand, *stronger-first* shadowing can be devastating for protocols that rely on failure to detect collisions; collision do not result in failure nor are they detectable by the detection scheme described above.

Two experiments were performed to characterize the shadowing phenomenon and the collision detection and recovery techniques. In the first experiment, controlled collisions were created by generating two packets with a precise time difference Δt . This experiment confirmed that shadowing occurs even with a negligible difference in signal strength. It also confirmed that the detection and recovery techniques can succeed in up to 84% of *stronger-last* shadowing instances. In the second experiment, a 36-node broadcast storm was generated with multiple transmission power settings. This experiment confirmed that shadowing is very prevalent. It also characterized the relationship between shadowing and the distances between nodes in a collision.

This paper makes two different contributions corresponding to the two different types of shadowing. First,

by exploiting the physical layer properties at the networking stack, we enable receiver-side collision detection and data recovery in stronger-last shadowing. Second, we demonstrate how the common connectivity-interference assumption breaks down in stronger-first shadowing and analyze the potential implications on protocols that rely on it. By analyzing the cross-layer implications of shadowing from the physical layer to high-level protocols, this work suggests that closing the disconnect between radio manufacturers and protocol designers could enable new protocols as well as strengthen existing ones.

Section 2 describes the shadowing phenomenon and the physical layer characteristics that can cause it. Section 3 discusses how collisions can be detected and how complete packets and additional information can be recovered from collisions. Section 4 explores the effect of shadowing, collision detection and recovery on several existing networking protocols. Section 5.1 demonstrates shadowing in a controlled experiment, showing that the connectivity-interference assumption is indeed violated when there is even a trivial difference in signal strength. Section 5.2 shows the extent of the shadowing phenomenon in a multihop broadcast using a 36-node flooding experiment. In Section 6, we suggest future directions for protocol design in light of the shadowing phenomenon. Sections 7 and 8 discuss related work and conclusions.

2 The Shadowing Phenomenon

Shadowing is the phenomenon in which, during a collision, one message is received correctly, even though all messages would have been received were there no collision at all.

This phenomenon breaks the connectivity-interference assumption, which would require that no messages be received upon collision. We can differentiate the two cases with a simple example shown in Figure 1. Node A and node C are not in range with each other, but both are connected to node B. Under the connectivity-interference assumption, if both A and C send simultaneously, node B would receive nothing. If A and C are both strong enough to be reliably received at B, they are strong enough to cause interference at B. The shadowing phenomenon would be exhibited with two different scenarios: either node A's message is received or node C's message is received.

These latter two scenarios contradict common intuition and experience with radios, but are quite possible with various modulation schemes, such as FSK, GFSK and PSK. These modulation schemes are very commonly used in wireless networks such as 802.11, Bluetooth, and soon 802.15.4/ZigBee networks. As an example, FSK works by sending one frequency to indicate a logical 1 and another to indicate a logical 0. The receiver compares the amount of energy in each frequency to interpret the bit. When two incoming messages are superimposed, the bit value corresponding to the stronger signal will be received correctly unless they are both carrying almost exactly the same energy.

How much stronger must one signal be than the

other is determined by the SNR of the receiver. The SNR in FSK is determined by the fidelity of the integral over energy in each frequency and, in modern radios such as the CC1000 [1], the SNR is extremely low. Experiments from Section 5 indicate that, using FSK, up to 40% of all neighbors and 70% of all collisions exhibit the shadowing phenomenon. This extreme prominence of shadowing indicates that, with some modulation schemes, the connectivity-interference assumption that is relied upon by many network protocols is simply not true.

The extent of shadowing heavily depends on the choice of the modulation scheme at the physical layer. Amplitude-shift-keying (ASK), for example, has very little shadowing effect. ASK works by sending a high amplitude signal (relative to the noise floor) to indicate a logical 1 and a low amplitude signal to indicate a logical 0. The superposition of two signals would not be equal to either signal individually, unless one signal is so strong and the other so weak that one is considered noise. With ASK, therefore, shadowing only occurs in pathological cases, perhaps when one transmitter is extremely close to the receiver while the other is extremely far. Therefore, the likelihood of shadowing is much higher with a modulation scheme such as FSK than with ASK, though both radios are used in Berkeley's mica sensor platform [8].

3 Collision Detection and Recovery

In stronger-last shadowing, collisions can be *detected* by simultaneously searching for the beginning of a new message while receiving an incoming message. If the receiver can properly stop receiving and resynchronize, it can also *recover* the second packet. In Section 5 these detection and recovery techniques are shown to be highly effective.

A *collision detection* is defined to be the finding of a preamble during the reception of a packet. This detection method is different from most protocol-based collision detection schemes, such as the use of Acks and RTS/CTS, because the information is obtained by the receiver instead of the transmitter; the *availability* of this information to mediate future collisions is different. Since this scheme is relying on information at the data level, it is not as powerful or transparent as techniques that might be implemented at lower physical levels. For example, the transmitter must perform byte-stuffing to prevent false positives in case a message inadvertently contains a preamble. It also implies that a longer preamble may be required to reduce false negatives. In addition, collision can only be detected if the strong-last message arrives after the start symbol of the first message.

If the link-layer notices detects a preamble while receiving, it has the option of dropping the weaker message and resynchronizing to the second. This is called *collision recovery* and requires only modest modifications to link-layer processing at the receiver. While a conventional link layer would generate one completely corrupted packet, collision recovery has two advantages: the point at which the first packet became corrupted can be identified and the later incoming message can be completely recovered. This not only improves network throughput, but retrieving the destination and source ad-

dress fields of the colliding packets can be quite valuable in mediating future traffic.

One limitation of these detection and recovery techniques is that they only work in *stronger-last* shadowing, ie. when the stronger message arrives last. In the *stronger-first* shadowing case, where the stronger message is received first, it will be received normally by conventional link-layer processing and the weaker, later packet could never be noticed until the first message is complete. Stronger-first shadowing is likely to be occurring all the time but, being undetectable outside of a controlled experiment, has not been noticed by the networking community.

A straightforward extension of this work is to use a termination symbol to end each message as we have already paid the overhead of byte-stuffing. The receiver could search for extra termination symbols after receiving each message to see if other, weaker messages had been arriving. This extension, however, has not been implemented. As a result, only the strong-last cases can be detected, and thus, on average, at most 50% of all collisions are detectable.

There are also other ways to detect collisions using techniques at the physical layer. For example, the receiver could exploit a sudden change in the signal to noise ratio and use coding and signal processing techniques to detect collisions or recover the clobbered bits. However, just as network protocol designers have largely not accounted for the shadowing phenomenon, radio designers have not accounted for the needs of networking protocols. Baseband controllers almost never expose methods for higher level protocols to utilize collision detection and recovery. For example, while this section explains how to implement collision detection and recovery on the CC1000 [1], its next generation, the CC2420 [2], inadvertently *prevents* the higher-level protocols from doing this by not providing the ability to resynchronize during a packet reception.

It is possible that techniques such as these have not been used in conventional wireless MACs for exactly this reason: the low-level receiver processing is delegated to a separate controller or proprietary firmware, and is difficult to access. This could be especially true in 802.11 networks, which have a strong commercial influence. Because of this, even though shadowing may occur in 802.11, Bluetooth, and 802.15.4/ZigBee, these detection and recovery techniques were implemented on Berkeley's Mica platform [8], which permits very low-level access to the radio and networking stack. Nonetheless, the success of the techniques in this section suggest that, even if existing baseband controllers do not currently allow collision detection and recovery, future models could be redesigned with this in mind. Furthermore, since stronger-first shadowing occurs even *without* changes to radio receivers, protocol designers may need to account for shadowing regardless.

4 Implications of Shadowing

In either forms of shadowing, stronger-first and stronger-last, recovery of the stronger packet is almost always beneficial to a network protocol; it will receive one mes-

sage that it would not have otherwise. In the stronger-last case, it also knows that it was in a collision, and possibly which other transmitter was involved. The stronger-first case, which happens transparently with most wireless links in current use, can be quite problematic. Because current implementations of collision detection cannot identify this as a collision, protocols that rely on the connectivity-interference assumption will break down.

This section explores how collision recovery and the shadowing phenomenon can both help protocols such as flooding and cause unforeseen problems in others such as media access protocols.

4.1 Flooding

Flooding is one of the important mechanisms in ad hoc networks, for example, in both state dissemination and route discovery [9, 12]. Because flooding uses broadcast, it generally skips the RTS/CTS and Ack handshakes of the MAC protocol. Flooding also generates highly correlated traffic and many hidden terminals, making it likely to generate a lot of collisions; this phenomenon is often described as the broadcast storm problem [11].

Collisions not only make a flood slower, they reduce reliability and create complex dynamic behavior. As messages are lost, the flood can run depth first through the network, then backtrack to fill the gaps left by collisions. This is problematic for flooding based algorithms such as spanning tree. With shadowing, many of the collided packets can be recovered, completely changing the dynamics of the flood. Furthermore, because collision detection techniques are receiver based, collision information can be used to help make the decision of when to rebroadcast, the most critical decision a node must make during a flood.

4.2 Media Access Control

We usually assume that no successful packets can be received during a collision. Protocols in media access control or congestion control exploit this assumption and rely on a successful acknowledgment to mean that no collision occurred. With shadowing, the semantics of an Ack is changed; it can only mean that the packet is correctly received, but contains no affirmation that it was collision-free.

A similar situation arises in the RTS/CTS handshake [10, 4]. Traditionally, a CTS should be sent by the receiver only if it is idle for reception. However, shadowing allows a receiver to receive an RTS from the sender even though it is overhearing messages not destined for it. By replying with a CTS, the receiver would be committed to receive messages **during** a collision, which is very different from the original meaning of a CTS.

Fairness is a main concern for media access control, but, with shadowing, not all transmitters can contend fairly. In the extreme case for CSMA/Ack, an overpowering transmitter can keep sending over packets from others, and create the common channel capture problem. It might also always receive an Ack, so it would never receive any indication that other transmitters are contending for the channel. This problem arises even without the new recovery mechanism we propose since stronger-first

shadowing occurs regardless. Furthermore, it applies to all contention based MAC protocols such as RTS/CTS.

4.3 Spatial Reuse

Shadowing enables potential spatial reuse that is not possible with the traditional connectivity-interference assumption. Figure 2 shows an interesting scenario of spatial reuse of four nodes all having connectivity with each other. With shadowing, it is possible that node B and node C can recover messages from node A and node D even though collisions are constantly occurring. Another kind of spatial reuse is for multihop routing. In multihop routing, nodes are constantly trying to avoid suffering from the hidden terminal problem; most likely, their parents and children cannot hear each other and are often sending correlated transmission. Figure 3 shows a typical scenario where each node is only connected with its neighboring nodes. The data direction is flowing from node A to D. The common case is that node C will expose node B, making it impossible to receive A's messages. With careful use of shadowing, it is possible to reuse the channel, and allow node B to receive A's messages while C is transmitting to D.

4.4 Localization

Cricket [13] is a localization scheme that proposed embedding the ultrasound signal within a radio transmission. Many other radio/acoustic time-of-flight schemes also use a similar technique [14]. Since collisions cannot be detected in ultrasound, the radio transmission is used to guarantee a lack of collisions. If a radio message collides, nothing should be received, allowing both the ultrasound signals to be ignored. With shadowing, however, this is no longer true. A radio message could be received, and perhaps associated with an ultrasound message sent by the node that sent the packet that was collided with. This can cause errors unforeseen by the localization system.

In summary, we have seen new issues starting to arise at the protocol level when shadowing becomes possible. In Section 6, we will revisit these issues, discuss how they should be changed in the presence of shadowing, and layout the corresponding costs and benefits.

5 Empirical Observations

Two experiments were performed to characterize the nature of collisions, shadowing, and the collision detection and recovery techniques described in Section 3. The first experiment demonstrates shadowing and characterizes detection and recovery in the context of both stronger-first and stronger-last shadowing. The second experiment characterizes the frequency of shadowing and collisions in a 36-node flooding experiment and explores their relationship to network topology.

Both experiments were done using the Berkeley Mica2 [8] hardware platform. The Mica2 platform consists of a 7.38MHz Atmel 128 microprocessor, with 128kB program and 4kB data memory. The radio is a 433MHz Chipcon CC1000 [1] transceiver. Its transmit power can be changed from -20dBm to +10dBm. RSSI values can be obtained through the on-board 10-bit

ADC on the processor. Manchester encoding is used at the physical layer, and delivers 19.2kbps encoded bandwidth. No extra byte-level coding or forward error correction is used in the default networking stack in TinyOS. The corresponding byte time is about 450us. The default packet size is 36 bytes, with 5-byte preamble and 2-byte start symbol. The packet consists of a 5-byte header, 29-byte payload, and 2-byte CRC checksum. The destination address is the first 2 bytes in the header. In our experiments, we overload the source address on the first 2 bytes of the payload. The TinyOS default MAC layer is CSMA, with a 1-32 byte window of random delay before carrier sensing. The antenna used is a standard quarter wave-length whip antenna which, during these experiments, was in a random orientation.

5.1 Controlled Collisions

The first experiment was intended to both exhibit the shadowing phenomenon and also to characterize the collision detection and recovery techniques under controlled conditions. It involved two transmitters A and B and a receiver, all placed in a one-meter isosceles triangle. The receiver was using collision detection and recovery as described in Section 3. The transmitters were synchronized to transmit at times t_A and t_B , such that the time between them was $\Delta t = t_B - t_A$. Thus, when Δt is positive A transmits first and when it is negative B transmits first. Δt was varied from $-23ms$ to $23ms$ at 1ms intervals and 10 collisions were generated at each value of Δt . The default 36-byte packet length was used so packets were 17.9ms long, so Δt was changed by 0.5ms around 0ms and $\pm 18ms$ for higher resolution data. Time synchronization was only accurate to 1ms.

While the default 0dBm transmit power was used on both nodes, slightly moving one of the senders or adjusting the antenna orientation changed the power relationship between the two senders at the receiver. It was fairly difficult to find a "null" point at which neither transmitter was received due to the difference in received energy from both transmitters being below the SNR threshold of the receiver. In fact, we confirmed that the power difference required to cause one transmitter to be received over the other one was unobservable using the 10bit ADC to sample the RSSI pin on the radio. In our experiments, we deliberately move node B such that its signal was stronger than node A at the receiver.

Figure 5.1 summarizes the findings from this experiment. The Y axis is in percent (%) while the x axis is Δt . In the left half of the graphs B sent first while on the right half of the graphs A sent first. At the two edges of the graphs where $|\Delta t| > 17.9$ the messages do not overlap in time. At $x = 0$ they were both sent at the same time. Notice that many of the transitions in the graphs are not discrete because of the 1ms error in time synchronization.

Figure 4.a shows that B was almost always received. Figure 4.b shows that A was also almost always received, except when its message overlapped in time with B , establishing that B has both *interference* and *connectivity* with the receiver. However, A has connectivity but not interference, since it did not in-




Figure 4: **Three Node Experiment** This shows what the receiver can detect and recover during collisions. Node A and B are senders, and B overpowers A at the receiver.

interfere with B 's message. To verify that A had connectivity during the collisions, Figures 4.e and f show that, even though A 's message was clobbered, the destination and address fields could be reliably recovered when B 's message came sufficiently late, ie. when Δt was large enough. This experiment is therefore a simple and easy example to reproduce that breaks the connectivity-interference assumption; node A clearly had connectivity but it did not produce interference. Under the connectivity-interference assumption, node B 's messages would not have been received.

It is due to node B 's connectivity in the presence of the collision that makes collision detection possible, the results of which are shown in Figure 4.d. When B 's message arrives after the preamble and start bytes of A 's message, ie. when $\Delta t > 3ms$, a collision can be detected almost 100% of the time because two sets of preambles and start bytes are received. If the probability of Δt were uniformly distributed, collision detection would have a 42% success rate. This is essentially the theoretical limit that this technique can achieve, given the current preamble length. Note that this technique therefore has a corresponding 58% false negative rate; in 58% of all collisions it indicates that there is no collision

Recovery, on the other hand, fails when $1 < \Delta t < 3$, as shown in Figure 4.c. This is because, once the preamble of A is detected, the software waits up to five byte times for the arrival of the start bytes. However, they do not arrive because they are overpowered by B 's

preamble. By the time B 's start bytes arrive, the software has gone back to looking for preamble bytes. This software glitch will be fixed in the next version. Even with the glitch, however, collision recovery provides a 46% yield, meaning that almost half of the messages in a collision can be recovered. In this case, 25% are recovered by the standard network stack (the left half of the graphs) while 21% are recovered explicitly by the collision recovery technique (the right half of the graphs). This is near the theoretical limit of 50% yield for a two-packet collision.

5.2 Flooding-Induced Collisions

The second experiment was meant to explore collisions, shadowing, detection and recovery in a setting with many correlated transmissions but with a CSMA MAC protocol actively trying to prevent collisions. We chose to create a broadcast storm [11] scenario because of its simplicity and high likelihood of collisions. The results of this section, however, may apply to any traffic load with correlated transmission, such as event-detection, multi-hop routing, etc. that will also generate a large number of collisions.

Thirty-six nodes were deployed in a 9x4 grid with approximately two meter spacing. The nodes were on the floor in a large room with many obstacles such as tables, chairs and couches. One node in the corner emits a message with a sequence number to trigger the network flood. Every other node retransmits the message with

that sequence number once, with a random time up to 120ms after it first received that message. If the radio channel was busy, it took a random backoff of at most 12ms and retried, as implemented by the newest MAC protocol available with TinyOS [3]. Collisions should therefore only occur when the MAC layer fails or when there is a hidden terminal, ie. a single receiver can hear two transmitters that cannot hear each other. In order to explore the effects of power on shadowing and to change the number of hidden terminals in each flood, transmission power on all nodes was varied from -18dBm to 4.5dBm and 5 floods were done at each power setting. At -18dBm, the radio range was approximately two to three meters which at 4.5dBm the radio range covered the entire network.

With 36 nodes, each transmitting once per experiment, five experiments per power setting and five power settings, the experiments could have included up to 900 transmission if none of the floods died prematurely. In reality, we saw 857 transmission, 12687 receptions, 2036 collisions and 1142 instances of shadowing. In the experiment, each node recorded the timestamp of each message it sent and received so that the existence of collisions could be inferred after the experiment. Time synchronization was accurate to about 500 microseconds.

Figure 5 summarizes the findings from this experiment. The Y axes are in percent (%) while the x axes are power settings (dBm). These plots are box and whisker plots, which indicate a data set with four statistics. The central bar is the median while the box indicates the upper and lower quartiles. The whiskers extend to cover the rest of the data except for the outliers, which are indicated by + marks. Outliers are points that fall outside the box by more than 1.5 times the inner quartile range. Where there is no central bar, the median falls on the quartile mark.

In the analysis that follows, a transmitter-receiver pair is defined to be *connected* if at least 33% of the messages transmit were received at the receiver. A *collision* is said to take place at a receiver if two or more transmitters that are connected to the receiver transmit messages that overlap in time. Because collisions are defined at the receiver, there will be far more collisions per experiment than the number of packets that were transmit. Furthermore, note that this is an extremely conservative method of counting collisions, in that there will be more collisions counted than actually existed, due to normal packet loss. The statistics on how many collisions were detected or recovered will accordingly be lower than expected, but it is fundamentally impossible to know exactly how many collisions took place.

5.2.1 Collision Observations

The primary question pertaining to collisions, and one that helps motivate this paper is: how often do they happen? If collisions are not common, the shadowing phenomenon is not of great importance. Figure 5.a shows the percentage of all packets transmit that participated in a collision at at least one receiver. Each box plot represents the percentage values for each node during each round at a particular power setting, for a total of 180 data

points. The figure shows that up to 40-50% of all messages transmit in the flooding experiment were involved in at least one collision. This is a very high number caused partially by the fact that, in a flood, each message is destined to so many receivers. It does, however, represents an upper bound on the number of messages that can be lost due to collisions in bursty network traffic, in the case that all messages lost happened to be intended only for those nodes that lost them. Of course, the number of collisions is a function of the 120ms delay chosen, and experiments have shown that a delay of about 1000ms was enough to allow all messages to be sent without any collisions. In a real flooding scenario, however, it is unreasonable to assume one can know this delay value *a priori* since it depends on the node density, so flooding is usually likely to cause such collisions. Other traffic scenarios can even be more pathological; in flooding, hidden terminals arise only out of chance but in multihop routing, every step of the way involves avoiding a collision between one's parent and one's child.

5.2.2 Shadowing Observations

There are two looming questions about the shadowing phenomenon, the first pertaining to its prominence. With modulation schemes like ASK, it is assumed that shadowing almost never happens. Figure 5.c indicates the percentage of messages received in a collision (yield), all of which must have been due to shadowing. This number is fairly high compared to what would be expected with ASK: nominally 20% and up to 50% is a conservative estimate.

Another measure of the prominence of shadowing is the percent of one's neighbors for which the connectivity-interference assumption holds. One would normally expect this to be almost 100%, and this is what many networking protocols rely on. However, Figure 5.b indicates that the percentage of neighbors a node has for which it can count on the connectivity-interference assumption under shadowing is extremely low, between 0 and 10%. That is to say that, for almost all of node A's connected neighbors, there exists at least one other transmitter B which can overpower A at that receiver. This shows that the shadowing phenomenon is not pathological, but in fact a common case.

The second looming question about shadowing is how it is related to distance. Intuition indicates that closer nodes should overpower farther nodes most of the time. Figure 6 is a histogram of the ratios of distances between the two transmitters and the receiver in the 1142 instances of shadowing observed in this experiment. When the weaker node was the closer node, the ratio was added to the left side of the histogram. When the stronger node was the closer node, the ratio was added to the right side (the graph was plotted this way to avoid ratios with values less than one). One can see that it is almost equally likely for the closer node to be the stronger node as it is for it to be the weaker. The closer node is stronger only 57% of the time while it is weaker 43% of the time.

While the ratio of distances to the two transmitters is interesting from the receiver's point of view, a trans-




Figure 5: **Broadcast Storm.** Part a) on the upper left corner shows the extent of collision. Part b) shows the percent of neighbors that holds the connectivity-interference assumption. Part c) shows the yield, and part d) shows the rate of detecting collision. The boxes indicate the upper and lower quartiles, with the central dotted bar indicating the median. The whiskers show the extent of the outliers indicated as +.

mitter might be interested in knowing the probability of being overpowered at a node given its distance away, independent of which other nodes are transmitting. Intuition indicates that a node should be able to dominate its closer nodes most of the time, and farther nodes less often. However, Figure 7 indicates that the probability of a transmitter being overpowered given the receiver's distance is nearly uniform. These surprising results suggest that the shadowing effect is completely uncorrelated with distance.

5.2.3 Detection and Recovery Observations

Figure 5.c shows the percent yield of collision recovery to be nominally 20% and up to 50%, where yield is defined to be the ratio of the number of received packets in a collision, either due to stronger-first shadowing or stronger-second shadowing and collision recovery, to the number of contenders in a collision, i.e

$$\text{Yield} = \frac{\text{Recovered} + \text{Received Normally}}{\text{Collided}}$$

As mentioned earlier, the denominator here is being overestimated since many of the colliding messages would have been lost due to normal link failure. Even so, shadowing can yield up to 50% of the messages that would normally have been lost is a significant gain. The theoretical upper bound is not 50% since two or three messages can be recovered from a sequence of many colliding messages, which actually did occur.

Figure 5.d shows the detection rate of collision detection to be nominally 20% and up to 50%, where de-

tection is defined to be the ratio of detected or recovered messages to the number of contenders during a collision, i.e

$$\text{Detection} = \frac{\text{Recovered} + \text{Detected}}{\text{Collided}}$$

This value represents the number of collisions that are correctly identified by collision detection.

These last two statistics are much lower than the 46% and 42% we would expect from the data in Section 5.1. This is partially due to the fact that the probability of Δt is not uniformly distributed; many collisions actually occur within the 3ms window $1 < |\Delta t| < 3$ due to MAC backoff. One interesting observation pertaining to power variation is that the number of collisions, percent yield and detection rate all go up as the transmission power is set such that the radio range is greater than the network grid size but less than the entire grid. This is because, at high power settings there are no hidden terminals and the MAC layer can successfully avoid most collisions. At low power settings the radio range is so small that the number of nodes each transmitter must contend with is small.

6 Discussions

In this section, we focus on MAC protocols and suggest how they might be designed differently in light of the shadowing phenomenon. We then discuss some of the potential tradeoffs if shadowing were to be exploited. This discussion builds on the implications of shadowing on protocols as presented in Section 4.

6.1 MAC in Shadowing

For CSMA/Acks [7], shadowing not only changes the original semantic meaning of Acks for media access control but also leads to potential channel capturing. One approach to solve this problem is to have the receiver embed collision detection information in the Ack. If a collision is detected but no packet is received, a negative Ack (Nack) is sent such that contenders can back-off. In stronger-first shadowing, an Ack is sent normally since shadowing is undetected. Contenders can then infer stronger-first shadowing from the source and destination addresses of the Ack. In stronger-last shadowing, the receiver can detect the collision, and an Ack can be used to signal such an occurrence. It is unfortunate that upon reception of these Acks, contenders, in both cases of shadowing, can only backoff, which is ineffective.

Increasing their transmit power to avoid overpowering may also not work since it will create undesirable ripple effects. Instead, the receiver can actively mediate the traffic by feeding backoff information to the transmitters through the Acks. This requires the receiver to detect collision for both kinds of shadowing. Since stronger-last shadowing can be detected by default, we need a mechanism to infer stronger-first shadowing. Using termination symbols, as suggested in Section 3, receivers can also detect stronger-first shadowing. As a result, receivers can simply embed backoff information over Acks when shadowing is detected. Such techniques can potentially avoid the capturing problem even in the presence of shadowing. This protocol is advantageous because, aside from recovering data from collisions, channel sharing is controlled from a central point, allowing better optimization of channel usage than random backoffs.

RTS/CTS handshakes also encounters similar issues of channel capturing. When based on the connectivity-interference assumption, RTS contention is fair implicitly. However, with shadowing RTS/CTS can make no such guarantees since RTS packets may always be overpowered. Nonetheless, the techniques discussed for CSMA/Ack can also apply in this case. Alternatively, receivers can penalize more on the winning sender to avoid capturing by controlling each issue of CTS to maintain fairness. Another related issue is whether or not to exploit spatial reuse. With shadowing, it is possible to receive a RTS even though a node is exposed by its neighbor's communication. In general, the question is whether or not to grant a RTS if a node heard a CTS from its neighbor before. For the purpose of spatial reuse, it is desirable to grant such a RTS. However, the consequence is that it may clobber its neighbor's reception and result in low data delivery overall. While a node can ensure that no clobbering will occur, it may require maintaining much more neighbor information, such as fine-grain RSSI values and the corresponding signal-to-noise ratio between each pair of neighbors.

6.2 Costs and Benefits

The shadowing phenomenon introduces many new trade-offs in protocol design that were not possible to make otherwise. These tradeoffs can involve spatial-reuse, protocol overhead, and fairness. We observe clear ben-

efits of spatial reuse during flooding, without any extra protocol overhead. The spatial reuse benefits also apply to multihop routing, cluster-based topology, and base-station type infrastructure communication. Shadowing also enables information about collisions to be extracted at the receiver, which enhances the effectiveness of CSMA/Ack like protocols, and has a potential to decrease protocol overhead for solving the hidden terminal problem. For example, instead of using RTS/CTS, collision detection at the receiver makes phase-shifting type of media access control found in sensor networks attractive [16]. However, this technique assumes data rate is low and periodic. For normal Internet like traffic, more complex backoff mechanisms are required to achieve fairness under shadowing as discussed before. In general, to exploit spatial reuse and fairness under shadowing requires higher protocol overhead in protocol messages and the amount of state to maintain.

7 Related Work

Collision detection has been widely deployed in wired communication such as CSMA/CD in Ethernet (IEEE 802.3), where the sender monitor the difference between its transmit signal and the observed signal on the wire to infer collision. In wireless networks, this technique is difficult because wireless radio usually cannot receive while it is transmitting. Even if it is possible to detect collision at the sender, the ultimate goal is to detect collision at the receiver. In fact, researchers have used an out-of-band radio to provide collision feedback to the transmitter [6], and show significant improvement in simulations over the 802.11 MAC protocol. While this approach is useful, the cost of having two transceiver is high. An alternative is to create radios with full-duplex communication as discussed in [5], where the receiver can listen and provides feedback to the transmitter at the same time using the same channel. However, this work is simulation based and we cannot find such radios exist in the market. Up to our current knowledge, we are unable to identify other prior work that describes mechanism to either detect packet collision at the receiver using the same radio or recover packets during a collision. However, relevant research in information theory has shown the ability to recover information using intelligent coding techniques such as CDMA [15]. That is, the idea of using extra processing to extract or recover more information from the noise is not new. Nonetheless, many of these opportunities are either not exposed or implemented by the radio manufacturers. As a result, protocol designers can seldom exploit them.

8 Conclusion

Several surprising results arise in this paper. We empirically find that the common assumption of a node interfering with all its neighbors within its communication range can be false. In fact, a receiver can correctly receive the stronger message during a collision. Such shadowing phenomenon is not accounted for in many current networking protocols, even though stronger-first shadowing already exists over a class of significant radios.

We identify how these protocols may break given such connectivity-interference assumption is broken.

By exploiting the shadowing property and introducing additional processing at the link layer, we successfully enable a receiver to both detect collision and recover stronger-last packets in a collision. In the case of our FSK based radios, we have shown that shadowing is very common. Furthermore, experiments have shown how flooding is significantly benefited by spatial reuse enabled by shadowing, and it potentially can benefit other types of communications such as multihop routing.

The shadowing phenomenon imposes new complications over classical MAC protocols and techniques that rely on the connectivity-interference assumption. Potential solutions of how MAC protocol designs are different under shadowing is discussed. Spatial reuse benefited by shadowing comes at a cost of MAC fairness, and extra protocol overhead.

Many protocol issues remain to be solved in the future before shadowing can be fully exploited at the link and network layer. Although our study is based on one particular FSK radio, stronger-first shadowing also exists on other FSK and PSK radios Bluetooth and 802.11. We are interested in further pursuing the shadowing phenomenon on these radios in the future. To exploit the potential of stronger-last shadowing, cross-layer interactions down to the physical layer is required. We hope that radio manufacturers would expose the necessary hardware or firmware interfaces such that shadowing can be easily exploited by protocol designers.

References

- [1] Chipcon. http://www.chipcon.com/files/CC1000_Data_Sheet_2_1.pdf.
- [2] Chipcon. http://www.chipcon.com/files/CC2420_Data_Sheet_1_0.pdf.
- [3] Tinyos. <http://www.tinyos.net>.
- [4] Vaduvur Bharghavan, Alan Demers, Scott Shenker, and Lixia Zhang. Macaw: a media access protocol for wireless lan's. In *Proceedings of the conference on Communications architectures, protocols and applications*, pages 212–225. ACM Press, 1994.
- [5] Ajay Chandra, V. Gummalla, and John O. Limb. Receiver architecture for high-speed wireless lan protocols. In *Vehicular Technology Conference VTC 1999-Fall*, volume 2. IEEE, September 1999.
- [6] Ajay Chandra, V. Gummalla, and John O. Limb. Wireless collision detect (wcd): Multiple access with receiver initiated feedback and carrier detect signal. In *ICC*, pages 397–401. IEEE, 2000.
- [7] A. Colvin. CSMA with collision avoidance. pages 6(5):227–235, 1983.
- [8] Crossbow, Inc. www.xbow.com.
- [9] D. Johnson and D. Maltz. Dynamic source routing in ad hoc wireless networks. In *Mobile Computing*, pages 153–181. Kluwer Academic Publishers, 1996.
- [10] P. Karn. MACA - a new channel access method for packet radio. In *ARRL/CRRL Amateur Radio 9th Computer Networking Conference*, pages 134–140, 1990.
- [11] Sze-Yao Ni, Yu-Chee Tseng, Yuh-Shyan Chen, and Jang-Ping Sheu. The broadcast storm problem in a mobile ad hoc network. In *Proceedings of the 5th annual ACM/IEEE international conference on Mobile computing and networking*, pages 151–162. ACM Press, 1999.
- [12] Charles E. Perkins and Elizabeth M. Royer. Ad hoc on-demand distance-vector (aodv) routing. In *Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications*, 1999.
- [13] Nissanka B. Priyantha, Anit Chakraborty, and Hari Balakrishnan. The cricket location-support system. In *Proceedings of the 6th annual international conference on Mobile computing and networking*, pages 32–43. ACM Press, 2000.
- [14] Andreas Savvides, Chih-Chieh Han, and Mani B. Strivastava. Dynamic fine-grained localization in ad-hoc networks of sensors. In *Proceedings of the 7th annual international conference on Mobile computing and networking*, pages 166–179. ACM Press, 2001.
- [15] A. J. Viterbi. CDMA: Principles of spread spectrum communication. Addison Wesley, 1995.
- [16] Alec Woo and David E. Culler. A transmission control scheme for media access in sensor networks. In *Proceedings of the 7th annual international conference on Mobile computing and networking*, pages 221–235. ACM Press, 2001.




Figure 1: A typical hidden node scenario in which node A and node C cannot hear each other.




Figure 2: A scenario illustrating potential spatial reuse under shadowing.




Figure 3: A multihop traffic scenario.




Figure 6: **Distance Ratios of Transmitters Exhibiting Shadowing**




Figure 7: **Probability of Being Overpowered Given Distance**