

A Power Controlled Multiple Access Protocol for Wireless Packet Networks

Jeffrey P. Monks, Vaduvur Bharghavan, and Wen-mei W. Hwu

University of Illinois
Urbana-Champaign, IL 61801

Email: {jmonks, bharghav, hwu}@crhc.uiuc.edu

Abstract— *Multiple access-based collision avoidance MAC protocols have typically used fixed transmission power, and have not considered power control mechanisms based on the distance of the transmitter and receiver in order to improve spatial channel reuse.*

This work proposes PCMA, a power controlled multiple access wireless MAC protocol within the collision avoidance framework. PCMA generalizes the transmit-or-defer “on/off” collision avoidance model of current protocols to a more flexible “variable bounded power” collision suppression model. The algorithm is provisioned for ad hoc networks and does not require the presence of base stations to manage transmission power (i.e. it is decentralized). The advantage of implementing a power controlled protocol in an ad-hoc network is that source-destination pairs can be more tightly packed into the network allowing a greater number of simultaneous transmissions (spectral reuse).

Our initial simulation results show that the PCMA can improve the throughput performance of the non-power controlled IEEE 802.11 by a factor of 2 with potential for additional scalability as source destination pairs become more localized, thus providing a compelling reason for migrating to a new power controlled multiple access wireless MAC protocol standard.

Keywords— *Multiplex access, ad hoc wireless networks, power control*

I. INTRODUCTION

A major issue in wireless networks is developing efficient medium access protocols that optimize spectral reuse, and hence, maximize aggregate channel utilization. Recent theoretical studies have shown that ideal medium access protocols using optimal power control can improve aggregate channel utilization by up to a factor of $O(\sqrt{\rho})$, where ρ is the density of nodes in the region (using fluid model approximations) [1]. This motivates the study of power-controlled wireless medium access protocols.

Past work on power control has primarily dealt with cellular networks, where separate frequency bands are typically allocated for uplink and downlink channels and base stations provide centralized control [2], [3]. Distributed power control algorithms have also been presented [4], [5] in the sense that individual base stations control the power. However, these techniques still require the fundamental cellular configuration (mobile users communicate through base stations - centralized access). Other work has focused on MAC protocols that control the transmission power level to conserve power consumption [6], [7].

This paper differs from related work in two significant ways: (a) we focus on wireless multiple access networks, where all nodes share a single channel and there is no centralized control or access, and (b) we focus on power control as a mechanism for increasing channel efficiency rather than as a mechanism for increasing battery life (though that may be a desirable offshoot of our approach).

The dominant wireless multiple access protocol is currently the IEEE 802.11 standard, which follows the “carrier sense multiple access with collision avoidance (CSMA/CA)” paradigm. Our goal is to propose power-controlled multiple access protocols that follow the same collision avoidance principle. To the best of our knowledge, there exists no power-controlled MAC protocol that fits within the collision avoidance framework. We show in Section II that this is due to fundamental characteristics of the handshake and collision suppression mechanisms in the CSMA/CA class of protocols, which require that stations transmit all control packets at the same power level. The main contribution of this work is to achieve power controlled transmission while still preserving the collision avoidance property of multiple access protocols. Our proposed protocol, PCMA (power controlled multiple access) demonstrates improvements in aggregate channel utilization by more than a factor of 2 compared to the IEEE 802.11 protocol standard.

The rest of the paper is organized as follows. Section II discusses the challenges in achieving power controlled CSMA/CA protocols, and presents the overview of the PCMA approach. Section III discusses the channel characteristics and power constraints that must be considered when implementing a power controlled wireless MAC protocol. Section IV defines the PCMA protocol. In Section V, PCMA is compared to IEEE 802.11 and an ideal power controlled protocol using an implementation in the *ns2* wireless network simulator. Finally, Section VI summarizes key results and issues.

II. THE PROBLEM AND APPROACH TO THE SOLUTION

Multiple access-based collision avoidance MAC protocols have made the case that a sender-receiver pair should first “acquire the floor” before initiating a data packet transmission [8], [9], [10], [11]. Acquiring the floor allows the

sender-receiver pair to avoid collisions due to hidden and exposed stations in shared channel wireless networks (Figure 1 illustrates the scenario). The protocol mechanism that is used to achieve such collision avoidance typically involves preceding a data packet transmission with the exchange of a RTS/CTS (request-to-send/clear-to-send) control packet handshake between the sender and receiver. This handshake allows any station that either hears a control packet or senses a busy carrier to avoid a collision by *deferring* its own transmissions while the ongoing data transmission is in progress (as shown in Figure 1).

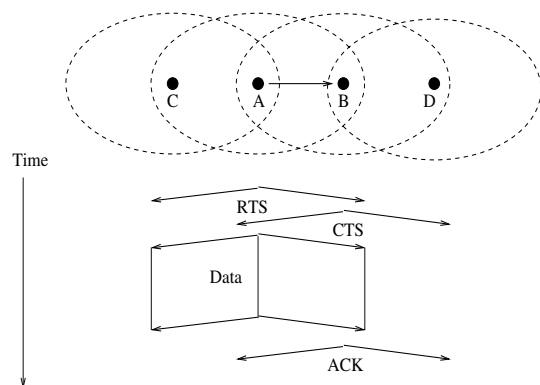


Fig. 1. General Protocol Operation for Multiple Access with Collision Avoidance. The top part of the figure shows four wireless nodes that have a transmission range shown by the dashed ellipses. A is the sender, B is the receiver, C is the exposed station (within range of sender, but not receiver) and D is the hidden station (within range of receiver, but not sender). Note that for a successful A-B transmission, D must not transmit. When A wants to send a data packet to B, it senses the channel to see if it is free. Then A sends an RTS to B. If C hears the RTS, it defers until A can hear B's CTS. If B is free to receive, it sends back a CTS to A. When D hears the CTS, it defers transmission until A finishes sending data to B. When C hears a busy carrier, it defers transmission. After B receives the data packet correctly, it sends back an ACK to A. This is the ideal operation of the protocol.

While acquiring the floor to enable collision avoidance from hidden and exposed stations is certainly a fundamental requirement for the efficient operation of wireless medium access, this method precludes multiple concurrent transmissions over the region of the acquired floor. To optimize spatial channel reuse in a shared wireless channel network, a pair of communicating nodes must only acquire the *minimum area* of the floor that is needed for it to successfully complete a data transmission (Figure 2 illustrates this scenario). Unfortunately, it turns out that for the collision avoidance mechanisms considered above (for 802.11) to work correctly, the control and data packets must be transmitted with a *fixed power* because of the following reason. When A is sending data to B, B's CTS must reach every hidden station whose transmission can cause a collision at B. Likewise, A's RTS must reach every exposed station with whom its data transmission can

collide. This means that an RTS-CTS exchange must acquire the channel over the maximum range over which any hidden or exposed station can cause collisions (a function of the maximum transmit power of an interfering station). Thus, even if A's data transmission is sent at a lower power (for purposes of power conservation), the A-B pair must acquire the channel assuming the worst case transmission power of all other (potentially interfering) stations in their region. From the perspective of channel reuse, this implies that adjusting transmission for data has no impact in terms of increasing channel reuse, and is equivalent to a "fixed power" MAC protocol. In summary, because control packets will need to be transmitted at the same fixed (maximum) power, current multiple access MAC protocols that follow the above framework cannot adaptively change the floor size acquired depending on how close the transmitter and receiver are to each other.




Fig. 2. Motivation for Power Control in Collision Avoidance-based Medium Access. For the given placement of nodes, in a traditional collision avoidance-based MAC protocol, if C is sending to D then A could not send to B since B would hear the RTS from C and sense the ongoing transmission. However, if C reduced its transmission power such that it would be just enough for D to capture its signal then other nodes in the region (e.g. A) could also proceed with their transmission. Such a protocol would allow for a tighter packing of source destination pairs within a network environment, thereby improving the spectral reuse.

Our goal is to change the "on/off fixed power" transmission model of the existing protocols to a more flexible "bounded and variable power controlled" transmission model, thereby changing the fixed floor acquisition model to an adaptive floor acquisition model for collision avoidance.

The fundamental change that we make in the existing approach is the following: unlike current protocols that use the reception of control packets as an on-off trigger for transmission/deferral by hidden and exposed stations, our approach is to use the *signal strength of a received control message to bound* the transmission power of these stations. This control message is a "generalized version of CTS", which we describe in Section IV as a signal pulse in a "busy tone" channel. However, its reception by a hidden station does not preclude transmission. In this case,

each hidden node bounds its transmission power by a function of the received signal strength of the generalized CTS at the node. Note that past research, most recently by Deng and Haas [12], also used busy tones, but as an added (“on-off”) collision avoidance mechanism. Given a mechanism that allows nodes to advertise their tolerance to interference by manipulating the transmission power of the generalized CTS, we are able to achieve power controlled multiple access by adhering to two key principles:

- (a) the *power conserving principle* dictates that each station must transmit at the minimum power level that is required to be successfully heard by its intended receiver under current network conditions (i.e. channel gain between source-destination pair and noise power observed at the destination), and
- (b) the *cooperation principle* dictates that no station that commences a new transmission must transmit loud enough to disrupt ongoing transmissions.

Enforcing these two principles, in concert with the mechanism for advertising interference tolerance through the generalization of CTS, achieves efficient power controlled multiple access within the framework of collision avoidance protocols.

III. THE NETWORK AND CHANNEL MODELS

As in IEEE 802.11 and other multiple access protocols [8], [9], [10], [11], [12], we assume a shared channel model in which simultaneous transmissions in the neighborhood of the receiver will result in a collision at the receiver. In a spread spectrum physical layer environment, this (shared channel access) model corresponds to a group of nodes accessing the medium with the same frequency hopping pattern in Frequency Hopping Spread Spectrum, or the same pseudo-random sequence number in Direct Sequence Spread Spectrum. At the MAC layer, we do not assume a cellular model, and we do not constrain designated “base stations” to be senders or receivers of data.

In the rest of this section, we first describe the channel propagation model, and then describe the transmission power constraints (network model) that must be satisfied by each node in the shared channel network such that the basic operating principles (the *power conserving principle* and the *cooperative principle*) are not violated.

A. Channel Propagation Models

We now describe the channel propagation model that we assume for a typical wireless channel. While our focus is on the MAC layer rather than the physical layer in this paper, we have taken into account the channel propagation characteristics at a sufficient level of detail that the power controlled MAC should work reasonably well in practice.

The amount of spatial reuse and transmission power required for a node to send a valid signal to its destination will depend on the *gain* between each source and destination, which models the *attenuation of the transmitter*

power over distance. We define two path loss field regions: the region where the gain drops proportional to the distance squared (inside the Fresnel zone) and refer to it simply as the $1/d^2$ field and the region (outside of the Fresnel zone or beyond the cross-over distance) where the gain is proportional to the distance to the fourth power and refer to it as the $1/d^4$ field. There are other channel effects that we outline below though a detailed discussion of these effects is outside of the scope of this paper and interested readers should refer to [13], [14].

In the protocol design, we measure the actual gain G_{ij} based on the sender power (advertised in the packet) and the receiver power. We then overcompensate to account for the distortions introduced from fading, by a factor of the fast fading amplitude (a function of the channel model).

Let us now investigate the channel assumptions of the protocol, and see how they hold for the channel propagation model above. In PCMA, we assume that:

1. The data and busy tone channels observe similar gains.
2. Channel reciprocity holds so that the gain between two nodes is approximately the same in both directions.
3. The channel gain is stationary for the duration of the control and data packet transmissions.

To ensure that the gain on the data and busy tone channels are similar (the first assumption), the busy tone frequency components must be within the coherence bandwidth of the data channel. The coherence bandwidth is inversely proportional to the multipath delay spread, which may vary greatly depending on the environment. In many outdoor environments the delay spread can be greater than 1 μ s resulting in a coherence bandwidth of less than 1 MHz [13], [14]. However, there are also minimal channel spacing constraints imposed as a result of the protocol requiring the busy tone be transmitted from receivers at the same time data is received. Therefore, to avoid the outgoing busy tone pulses from degrading the data it is necessary to place the busy tone channel outside of the coherence band. One method that would allow us to get around this is to allow the busy tone pulse to corrupt the data and assume the data has a sufficient amount of redundancy to correct the errors. Note that this is simplified by the fact that the receiver knows the bit positions that may be corrupted by the busy tone transmission. That is since the position of the errors are known they are actually erasures and only one bit is required to correct each erasure [15].

There are three basic channel effects [13]: path loss which is directly related to the separation between source and destination, shadowing which accounts for objects between the source and destination attenuating the signal, and multipath which accounts for the result of multiple paths (between sender and receiver) combining at the receiver. The distance is the same in both directions (i.e. source to destination and destination to source) and the

objects impeding the paths are the same in both directions. However, the way the paths refract off objects and combine at the source and destination receivers can differ depending on the extent of the multipath effects. Therefore, only multipath effects the validity of the second assumption. However, as long as the multipath effects are small the assumption holds. The third assumption guarantees that the channel gain measured from sending the initial request (control packet) is still valid for the duration of the data packet and ACK following. Path loss and shadowing will have little effect since the distance a node moves in the duration of a control and data transmission (on the order of a few milliseconds) is small. With multipath effects the gain will not be stationary for the duration of a packet. However, the (short term - on the order of a few bits) average gain measured for the RTS (or equivalent source request packet) is also valid in the data and ACK packets that follow since the short term average gain is primarily a factor of path loss and shadowing effects (slow fading). Even in cellular environments the power adjustments are not quick enough to keep up with the fast fading (multipath effects). Therefore, fast fading degradations must be overcome by physical layer techniques such as RAKE receivers, OFDMA [16], or channel coding, or be tolerated with additional overcompensation in transmission power.

In Section IV, we give a detailed description of the protocol under these assumptions, but later (in Section V) show that these assumptions can be violated to a certain degree with only modest degradations in performance¹. Further, we show that over compensation in transmission power can help to alleviate the problems to a considerable extent. In fact, adaptive over compensation techniques can potentially effectively address the fast fading problem, though a detailed investigation of these techniques is beyond the scope of this paper.

B. Power Constraints

Let Pt_Max and Pt_Min denote the maximum and minimum transmission powers for a transmitter on the data channel, respectively. Let RX_Thresh and CS_Thresh denote the minimum received signal power for receiving a valid packet and for sensing a carrier, respectively. Let SIR_Thresh denote the “capture threshold”, i.e. the minimum signal to interference ratio for which the receiver can successfully receive a packet.

Given the transmitter and receiver power parameters and the channel propagation characteristics, a transmitter i must transmit a packet to a receiver j at the minimum transmission power Pt_i that satisfies the following power constraints.

¹In fact, any protocol that makes the commutativity assumption (i.e. A can hear B implies B can hear A) has the same problem since fast fading can cause this assumption to be violated. We show through simulations in Section V that both PCMA and 802.11 are susceptible to this problem.

1. The transmission power of i must be within its parameter range, $Pt_Min \leq Pt_i \leq Pt_Max$.
2. The received power at j must at least be equal to the minimum received power threshold, $G_{ij}Pt_i \geq RX_Thresh$.
3. The observed signal to noise ratio for the transmission at j must at least be equal to the minimum SIR threshold, $SIR_j = \frac{G_{ij}Pt_i}{Pn_j} \geq SIR_Thresh$, where Pn_j is the total noise that node j observes on the data channel and is defined as $Pn_j = \sum_{l \neq i} G_{lj}Pt_l + N_j$. The term N_j is the power of the thermal noise (the power observed at a receiver when no nodes are transmitting) observed at node j .
4. Let E_k be the “noise tolerance” of any receiver k that is receiving an ongoing transmission in the neighborhood of i . E_k is thus the additional noise power that k (currently receiving data from some other node at power Pr_k) can tolerate before its SIR drops below its SIR_Thresh , and is defined as $E_k = \frac{Pr_k}{SIR_Thresh} - Pn_k$. Since the transmission power of i should not disrupt any ongoing transmission $Pt_i \leq \min_k \{ \frac{E_k}{G_{ik}} \} = Pt_bound_i$.

If the above four constraints can be met, then i can successfully transmit to j without disrupting any ongoing transmissions. The critical issues are therefore (a) handshaking between a transmitter-receiver pair to determine the minimum transmission power that satisfies constraints 2 and 3 (i.e. the power conserving principle), and (b) for every receiver to advertise its noise tolerance so that no potential transmitter will disrupt its ongoing reception applying constraint 4 (i.e. the cooperative principle). These problems are addressed in the PCMA protocol section.

IV. THE PCMA PROTOCOL

The goal of PCMA is to achieve power controlled multiple access within the framework of CSMA/CA based multiple access protocols. In these protocols, there are two main components: (a) *collision avoidance*, and (b) *collision resolution*. Collision avoidance takes place by means of a combination of carrier sensing by the transmitter and deferral of transmissions by hidden and exposed stations when they hear RTS/CTS packets. Collision resolution takes place by means of a backoff-based algorithm. In this section, an over view of the protocol is first given and then the protocol steps are described.

A. PCMA Protocol Overview

In PCMA, collision avoidance is generalized to power control. Conventional collision avoidance methods had an “on/off model”, wherein a node can either transmit (if it is not deferring and does not sense a busy carrier) or not. However, in Section III-B we determined that a node can transmit to its intended receiver so long as it satisfies four constraints. Thus, the on/off model is generalized to a “bounded-power model”. In order to achieve the bounded-

power model, the power control component in PCMA has two main mechanisms:

- A request-power-to-send(RPTS)/acceptable-power-to-send(APTS) handshake between the data sender and receiver, which is used to determine the minimum transmission power that will result in a successful packet reception at the receiver. The RPTS/APTS handshake occurs in the *data channel* and precedes the data transmission. After the successful reception of the data, the receiver sends back an ACK packet confirming its reception.
- The noise tolerance advertisement is used by each active receiver to advertise the maximum additional noise power it can tolerate, given its current received signal and noise power levels. The noise tolerance advertisement or busy tone is periodically *pulsed* by each receiver in the *busy tone channel*, where the signal strength of the pulse indicates the tolerance to additional noise. A potential transmitter first “senses the carrier” by listening to the busy tone for a minimum time period to detect the upper bound of its transmit power for all control (RPTS, APTS, ACK) and data packets.

The packet handshake sequence on the data channel is RPTS-APTS-DATA-ACK. Here we note that there is an issue in how to properly protect the ACK from collision since the noise power observed at the source cannot be updated during the data transmission. However, this is a fundamental problem associated with all power control methods since carrier sensing while transmitting is extremely expensive.

The last major component in PCMA is collision resolution, which is backoff-based. While a simple backoff algorithm similar to 802.11 was implemented to facilitate a one-to-one comparison with 802.11 and focus on power control, we can certainly use more sophisticated collision resolution algorithms as suggested in [9], [17].

To summarize, PCMA has one-to-one analogs of the key components of standard CSMA/CA protocols. At the sender, monitoring the busy tone is equivalent to sensing the carrier. At the receiver, periodically pulsing the busy tone is equivalent to sending a CTS for collision avoidance. The RPTS/APTS handshake that precedes the data transmission is similar to the RTS/CTS handshake, except that its purpose is not to force hidden senders to backoff. Thus PCMA can improve efficiency of channel access without changing the fundamental MAC paradigm.

B. PCMA Protocol Steps

Here we present the detailed PCMA protocol steps. The protocol steps correspond to some source node i sending to a destination node j and a potential interfering transmitter l , as show in Figure 3.

Step 1: A node i in its IDLE state monitors the busy tone channel to determine its power bound Pt_bound_i by

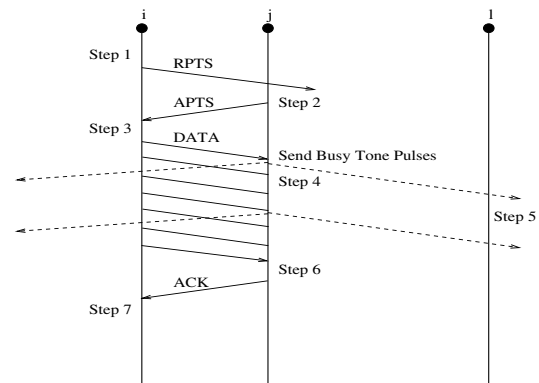


Fig. 3. PCMA protocol steps

measuring the maximum power received on the *busy tone* channel over a threshold time window. When i seeks to transmit a data packet, it waits until γPt_bound_i is greater than Pt_Min , and then backs off for a random interval bounded by its backoff counter to allow for contention resolution. The term γ is a constant (set to .9 for simulation results) that keeps the power level slightly below the threshold (Pt_bound). The node continues to sense the busy tone during its backoff. If at the end of the backoff the transmission power bound, Pt_bound , is still greater than the minimum transmit power, Pt_Min , by a factor of $1/\gamma$, then i sends a Request-Power-To-Send (RPTS) control message at the transmission power level $Pt = \gamma Pt_bound$ on the *DATA* channel. The RPTS packet contains the transmission power level, Pt , and source noise power, Pn_S_i (obtained from the air interface), placed in the packet.

Step 2: When the destination receives the RPTS, it measures the received power, Pr . The channel gain, G_{ij} , is computed to be the received signal power over the transmitted power (advertised in the RPTS packet). The receiver then requires the data be sent at

$$Pt_{i_des} = \max\left\{\frac{RX_Des}{G_{ij}}, \frac{SIR_Des \cdot Pn_D_j}{G_{ij}}\right\}, \quad (1)$$

in order to satisfy both its received power threshold and its SIR threshold. Here the constraints $RX_Des > RX_Thresh$ and $SIR_Des > SIR_Thresh$ ensure the constraints from Section III-B are enforced, and Pn_D_j is the noise power measured at the receiver. Pt_{i_des} is placed in an Acceptable-Power-To-Send (APTS) control packet so that the source can be notified of the power level to send its data packet. Assuming the same gain in both directions, the transmission power for the APTS packet is computed to be

$$Pt_j = \max\left\{\frac{RX_Des}{G_{ij}}, \frac{SIR_Des \cdot Pn_S_i}{G_{ij}}\right\}, \quad (2)$$

where the destination’s noise power is replaced with that of the source (extracted from the RPTS packet). If this power

is less than Pt_bound calculated at the receiver, then the APTS is sent at Pt_j on the *DATA* channel.

Step 3: When the source receives the APTS packet, it checks if the desired transmission power is below its current power bound, and transmits the *DATA* packet at Pt_{i_des} on the *DATA* channel if the bound is satisfied. If the source times out before receiving the APTS, it multiplicatively increases its backoff bound and starts over.

Step 4: The receiver starts sending busy tone pulses on the *busy tone* channel after starting to receive the data packet. The busy tone power, Pt_BT_j , sent from node j depends on the noise tolerance, E_j , (see definition in the previous section) is calculated

$$Pt_BT_j = \frac{C}{E_j}. \quad (3)$$

The value of $C = Pt_Max \cdot CS_Thresh$ is such that a node at a distance that would add exactly E_j additional noise when transmitting at max power Pt_Max would receive the busy tone at exactly the CS_Thresh . Note the busy tones only have to be received at the detection threshold since their power only has to be measured and no data bits need be received. Also since the busy tone's power can be no greater than Pt_BT_Max there is a minimum noise tolerance,

$$E_min = \frac{C}{Pt_BT_Max}. \quad (4)$$

This limits the ability of the busy tone to bound far away stations when the receiver is very sensitive to any small increase in noise. If there was not a minimum noise tolerance the busy tone power could potentially approach infinity and force nodes infinitely far away to not transmit at all. Note if E_min was plugged into Equation 3 the resulting busy tone power would then be $Pt_BT_j = Pt_BT_Max$, which conforms to the physical limitations. The resulting noise tolerance is then

$$E_j = \max\left\{\frac{Pr}{SIR_Thresh} - Pn_j, E_min\right\}. \quad (5)$$

Step 5: When a node l receives the busy tone at a power of $Pr_BT_l = \frac{C}{E_j} G_{jl}$ it calculates its transmission power bound imposed by node j as

$$Pr_bound_j = \frac{C}{\frac{C}{E_j} G_{jl}} = \frac{E_j}{G_{jl}}. \quad (6)$$

Then node j can receive at most $Pr_j = \frac{E_j}{G_{lj}} G_{jl}$, from node l since we assume that $G_{lj} \cong G_{jl}$, and $Pr_l = E_j$. Since there may be busy tones received from multiple receivers, the transmission power bound at a node is defined by the most sensitive receiver (receiver that can tolerate the least transmission power from this node)

$$Pt_bound = \min\left\{\min_j \left\{\frac{E_j}{G_{jl}}\right\}, Pt_max\right\}. \quad (7)$$

The receivers periodically send busy tone pulses (as oppose to a solid tones) in order to minimize the probability of destructive interference of busy tones (i.e. collisions). The width of the pulse is based on the signal capture interval of the receiver. Sending separate pulses also allows receivers to periodically update their noise tolerance advertisement to avoid collisions with new transmitters. The needed frequency of busy tone pulses is based on the rate of change of background noise (traffic load) and is evaluated in Figure 4, where sending a busy tone after every 128 bytes of data is found to be a sufficient update interval. There is another problem that may happen (particular at high traffic loads): multiple potential transmitters, upon hearing a receiver's busy tone, may locally decide that it is acceptable to transmit and commence transmission simultaneously (within one period of the busy tone advertisement), thereby cumulatively creating enough noise to disrupt an ongoing packet reception. This problem is similar to contention, except that failure of contention resolution disrupts ongoing transmissions rather than the contending packets themselves. A simple solution to reduce such collisions is for a receiver to immediately pulse a busy tone whenever it sees a change in its noise tolerance by a threshold level.

Step 6: When the destination receives the entire data packet without errors, it sends an ACK at the power level needed to get back to the source on the *DATA* channel.

Step 7: If the source receives a valid ACK it resets the max backoff and returns to the *IDLE* state otherwise, it increases the max backoff and starts over.

V. PERFORMANCE OF PCMA

In this section, we investigate the performance of PCMA under various network and channel conditions. The performance of PCMA is compared with the current standard MAC IEEE 802.11 for wireless networks in an ad hoc environment and an ideal power controlled protocol, IPC. The IPC protocol is provided with perfect (global) knowledge of the link gain between any two nodes, the noise at any potential destination, and the upper bound on a transmitter's signal power needed to protect other receivers (maximum transmission power that neighboring receivers can tolerate). The protocol, like IEEE 802.11, follows the RTS-CTS-DATA-ACK exchange. However, all messages are sent with only enough power needed to reach the destination, and like PCMA it backs off if the destination requires more power than a neighboring node can tolerate. It also like PCMA starts with in initial over compensated transmission power instead of making power adjustments through the transmission since in a multiple access environment the later is not practical due to contention delay. In this case, IPC demonstrates the upper bound on the performance of transmission power controlled protocols for the multiple access environment. To evaluate the performance of PCMA *ns2* (a commonly used network simulator)

was used and both PCMA and IPC were integrated into the CMU wireless extensions [18]. For these simulations the routing overhead was removed (since the goal of this paper is to evaluate the performance of MAC protocols and not routing protocols) and the destinations were restricted to within one hop of source nodes. Later work will evaluate the performance for multiphop wireless networks.

Parameter Type	Parameter Value
Packet Size	2 KB
Data rate	2 Mbps
Channel carrier frequency	916 MHz
RTS	20 Bytes
CTS, ACK	14 Bytes
RPTS	28 Bytes
APTS	18 Bytes
Max retransmissions	4
SIR_Thresh	6 dB
SIR_Des	10 dB
CS_Thresh	-78 dBm
RX_Thresh	-64 dBm
RX_Des	-60 dBm
Noise Floor	-104 dBm
Pt_min	-7.5 dBm
Pt_max	28.5 dBm
Pt	24.5 dBm

TABLE I
SIMULATION PARAMETER SETTINGS

The parameter values used in the simulation are shown in Table I. Here PCMA and IPC can send at a minimum power of -7.5 dBm and a maximum power of 28.5 dBm, and 802.11 sends at a fixed power of 24.5 dBm. The maximum power of PCMA and IPC are set to be 4 dB above the fixed power of 802.11 so that a destination at maximum transmission range for 802.11 will also be at maximum transmission range for PCMA and IPC allowing for a 4 dB compensation in transmission power. This allows the same scenario files (that determine the node connectivity) to be used for all three protocols. These parameters are reasonable and correspond to realistic settings in the hardware of a commercial wireless vendor. Our traffic model is simple: sources generate arrivals according to independent Poisson processes. The source node is picked randomly from the set of all nodes and the destination is picked randomly from the set of all nodes one hop away (in transmission range). Each data transmission between source and destination will be referred to as a flow, and each flow will have a specified rate that refers to the number of packets sent per second.

For the figures demonstrating the performance, the throughput is normalized by the carrier sense range (i.e.

550 meters) and the slot time such that the total number of arrivals and departures is divided by a scaling factor, sf , defined as follows

$$sf = \frac{\text{network area}}{\text{carrier range area}} \frac{1}{\text{data slot size}}. \quad (8)$$

This demonstrates the utilization with respect to the non-power-controlled MAC with optimal (best case) node placement. For a 1000 by 1000 meter network with the parameter settings in Table I, the resulting scaling factor is then $sf = \frac{1000^2}{550^2} \frac{1}{.008} = 413.22$.

We now present our results and compare the performance of PCMA to 802.11 and IPC. In Figure 4, the throughput is shown as the arrival rate is increased for a 1000 by 1000 meter network where nodes are uniformly distributed over the area. The performance of PCMA is demonstrated for differing number of busy tone pulses sent per data transmission period (1, 4, 16, 64). The performance increases as the number of busy tone pulses increases (approaching the performance of an ideal power controlled protocol (IPC)) since the feed back information (neighbor information) will be more up to date with more frequent busy tone pulses. However, the amount of improvement decreases and 16 busy tone pulses (i.e. sending one busy tone pulse for every 128 bytes of data since the data packets are 2048 bytes) is sufficient and the remaining PCMA results will be for one busy tone pulse sent every 128 bytes of data. This demonstrates that the bandwidth required for the busy tone channel is small with respect to the data channel. Figure 4 shows that the performance of PCMA is significantly better than 802.11 if a reasonable number of busy tone pulses are sent. Notice that with power control the utilization may go above 100% since the output is normalized with respect a non-power controlled protocol. PCMA restricts node transmissions according to variable interference region determined by the distance between source and destination (the *power conserving principle*) and busy tones (the *cooperative principle*), which can be significantly less than the fixed transmission range of 802.11.




Fig. 4. 100 nodes in a 1000x1000 meter network with 100 flows each sending 2 KB packets, and a connectivity range of 250 meters




Fig. 5. Throughput for a 100x100 meter network with 100 flows each sending 2 KB packets, and a connectivity range of 250 meters

If we now consider a network where 100 nodes are distributed over a 100 by 100 meter region the resulting throughput is then shown in Figure 5. Here we see that PCMA and 802.11 yields almost the same throughput performance, except at extremely high loads where PCMA does only slightly better. For this type of configuration the region is smaller than the transmission range and for 802.11 all nodes hear the RTS and CTS. As a result, there will be few collisions since there are no hidden stations. Although, in this region there is also significantly less spatial reuse for PCMA to take advantage of (since most nodes in the network are in the $1/d^2$ field instead of the $1/d^4$ field, see Section III-A). In addition, PCMA does not backoff based on the carrier sense so that it may take advantage of spectral reuse, causing it to more quickly reach its maximum number of retransmissions and give up early. This is a tradeoff made for not carrier sensing to improving spectral reuse.

However, we argue that uniform distributions do not well define the distributions of users in a typical environment. In most situations, we expect a more cluster grouping of nodes. Figure 6 demonstrates the throughput for a region containing both two and four clusters each 25 meters in diameter and located 50 meters apart. Here each node chooses a cluster at random and a random position within the cluster. A sender is chosen at random and the destination is chosen at random from the other nodes in the sender's cluster. In this configuration PCMA does significantly better than 802.11. The improvement is because PCMA can sent packets simultaneously in both clusters by reducing its transmission power, while in 802.11 each node in a cluster must always contend with the nodes in the other cluster (in addition to its own). Figure 6 shows that as the network becomes more clustered the throughput increases since a greater number of simultaneous transmissions are possible and less nodes compete within each cluster.

In the previously mentioned figures note that the performance of the power controlled protocols continues to increase even under very high loads due to long range transmissions being blocked by the transmission power bound




Fig. 6. Throughput for a 100x100 meter network with nodes separated into clustered regions




Fig. 7. Destination range distribution for PCMA with $P_{t,max} = P_t + 4dB$

allowing a greater number of short range transmissions. As the network load increases, the probability of a node requiring more power than the transmission power bound (set by the *cooperation principle*) also increases. The expected power for a source to reach its destination will increase as the network load increases due to an increase in background noise, and the expected transmission power bound decreases as the network load increase because there will be an increases in the number of exposed receivers in the network. Then sources requiring more transmission powers (i.e. greater transmission ranges for a simple path loss channel) will be more likely to backoff allowing a greater number of short range transmissions. Therefore, a power controlled MAC operating in a multiple access environment will result in unfair favoritism toward source-destination pairs sending over shorter distances. This phenomenon is particularly evident over the 250 meter connectivity range for PCMA as demonstrated in Figure 7, where the fraction of total packets received by destinations in five distance ranges (0-50, 50-100, 100-150, 150-200, and 200-250 meters) from their sources is shown for 100 flows sending 1, 4, 16, and 64 packets per second. A perfectly fair protocol would result in a linearly increasing number of packets sent to each range since the number of destinations within each range increases as $2\pi r$, where r is the distance from the source node. Notice for a very low transmission rate such as 1 packet per second, the number of packets sent to each range is linearly increasing. However,

as the network load increases the ratio of packets sent over a greater distance decreases, and for extreme loads we observe that the majority of connections are short range.




Fig. 8. Destination range distribution for 802.11




Fig. 9. Destination range distribution for PCMA with $P_{t,max} = P_t + 8dB$

The fraction of packets sent to each range for 802.11 is shown in Figure 8. The protocol also becomes less fair (sending fewer packets to greater ranges) as the load increase, but not to the extent that a power controlled protocol like PCMA does. The 802.11 protocol has an equal probability of sending packets to destinations at any distance since the transmission power is not taken into account while contending. However, because all transmissions are sent at a fixed power level there is less noise protection for destinations further from their sources resulting in a greater number of lost packets at greater network loads. PCMA on the other hand, has the same amount of protection (the compensation factor) for destinations at all ranges, however, the probability of sending a packet to further destinations decreases as the network load increases, as described above.

If we now increase the maximum transmission power to be 8 dB more than is required (P_t for 802.11) for a node at maximum transmission range instead of 4 dB, we see from Figure 9 that the fairness is improved over the flows shown in Figure 7. In particular, the 4 and 16 rate flows are now almost ideal (linearly increasing for increasing ranges) for all but the furthest range, and for the greatest flow rate the middle ranges are improved. The idea here is to increase

the range distributions while still limiting the transmission ranges to the same distance. This has the effect of stretching out the fairness plots and limiting transmissions in further ranges where a significant throughput disparity would be observed. This is a way to improve the fairness for power controlled multiple access protocols. However, it will also reduce the spectral reuse (throughput) and demand additional power from the transmitters reducing the node's lifetime (power reserves). Whether additional compensation is used or the scheduling algorithm is changed under heavy loads more work is needed to investigate techniques that overcome the fairness implications for power controlled protocols.




Fig. 10. Throughput for different amounts of gain distortion with varying compensations in a 1000x1000 meter network




Fig. 11. Throughput for different amounts of busy tone distortion with varying compensations in a 1000x1000 meter network

Thus far we demonstrated the performance of the PCMA protocol under ideal situation, where the channel gains are similar from one sample to the next and reciprocity holds. However, in many environments where we have multipath fading these assumptions may not hold so we demonstrate the robustness as these assumptions are relaxed. Implementing multipath fading into the simulation is very complex since this would require sample by sample deviations and is also difficult to generalize over all types of environments. Therefore, we alter the gain (degrading the channel) on a per packet basis by a factor χ (dB), where χ takes on the values $-\lambda$ with a probability .25, a value λ with probability .25, and a value of 0 (dB) with a probability of .5. The gain distortion factor λ will

be referred to as the distortion amplitude. While other methods for introducing distortion are more appropriate for specific environments this method shows the distortions that may be observed by the *MAC layer* and is sufficient to demonstrate the protocols performance as the assumptions are violated. Here we wish to evaluate the performance of PCMA as compared to 802.11 as the data channel is distorted to test its dependence on channel reciprocity and gain stability dependence, and as the busy tone channel is distorted to test the protocol under misinformation of neighboring node states. The throughput curves are shown in Figures 10 and 11 for both a distorted data and busy tone channels, respectively. For both cases the arrival rate is fixed to 32 packets per second and varying amounts of over compensation (0, 4, 8, 12, 16, and 20 dB) points are plotted with different distortion amplitudes (0, 4, 8, and 12 dB for the data channel distortion and 0, 4, 8, 12, and 16 dB for the busy tone channel distortion) for each plot. Figure 10 shows that both PCMA and 802.11 degrade with increasing distortions in the data channel. However, PCMA (with some over compensation) outperforms 802.11 up to 8 dB and after which it does slightly worse. Note that 802.11 does not have an over compensation factor so the plots shown for the protocol are straight lines. Both protocols have reciprocity and stability assumptions, but PCMA depends on it to make the correct power level settings. That is 802.11 assumes if I can hear you, you can hear me and PCMA assumes if I can hear you you can hear me at the same level (or actually within the over compensation factor) making it more sensitive to gain distortions. In Figure 11, PCMA outperforms 802.11 (with some over compensation) up to about 12 dB and does slightly worse with addition distortion to the busy tone channel. There is no busy tone channel in 802.11 so PCMA is compared to a single 802.11 plot. The distortion figures together show that PCMA is able to handle modest deviations from the assumptions stated in Section III-B, demonstrating that it can operate under various channel conditions.

The figures above demonstrated that for dense networks with a spatial reuse to be exploited PCMA performs significantly better than 802.11. When users generally communicate locally we observe that the protocol provides improvements in throughput and increases scalability. This demonstrates that there are compelling reasons for integrating power controlled MAC protocols into ad hoc networks.

VI. CONCLUSION

In this paper, we presented the PCMA power controlled medium access protocol within the collision-avoidance multiple access framework. We have demonstrated that PCMA allows for a greater number of simultaneous senders than 802.11 by adapting the transmission ranges to be the minimum value required to satisfy successful reception at the intended destination.

Of course, PCMA is still a protocol design in progress, we are working on a number of protocol design issues as well as engineering issues. Ongoing work includes rethinking of the fairness properties of PCMA, performance under mobility, and evaluating the protocol in a multihop wireless network.

Our preliminary performance results show that PCMA can achieve more than a 2 times improvement in aggregate bandwidth compared to 802.11 for highly dense networks. As the connectivity range is reduced, the aggregate throughput gain over 802.11 continues to increase. With over compensation in transmission power, PCMA can be designed to degrade approximately the same as 802.11 under channel distortion. These results lead us to believe that if engineered correctly, PCMA can achieve significant performance gains without significant compromises in robustness, and hence provides a powerful motivation to migrate towards power-controlled MAC protocol standards.

REFERENCES

- [1] P. Gupta and P. Kumar, "The capacity of wireless networks," *IEEE Transactions on Information Theory*, Nov. 1998.
- [2] F. Rashid-Farrokh and et. al., "Downlink power control and base station assignment," *IEEE Communications Letters*, vol. 1, pp. 102–104, July 1997.
- [3] J. Oh and K. Wasserman, "Optimality of greedy control in DS-CDMA mobile networks," *Mobicom '99*, 1999.
- [4] N. Bambos, "Toward power-sensitive network architectures in wireless communications: Concepts, issues, and design aspects," *IEEE Personal Communications*, vol. 5, pp. 50–59, June 1998.
- [5] J. Oh, T. Olsen, and K. Wasserman, "Distributed power control and spreading gain allocation in CDMA data networks," *INFOCOM '00*, pp. 379–85, 2000.
- [6] J. Chen and et. al., "A comparison of mac protocols for wireless local area networks based on battery consumption," *INFOCOM '98*, 1998.
- [7] R. Ramanathan and R. Hain, "Topology control of multihop radio networks using transmit power adjustment," *Proc. IEEE Infocom*, Mar. 2000.
- [8] P. Karn, "MACA - a new channel access method for packet radio," *ARRL/CRRL Amateur Radio 9th Computer Networking Conference*, pp. 134–40, Sept. 1990.
- [9] V. Bharghavan, A. Demers, S. Shenker, and L. Zhang, "MACAW: A media access protocol for wireless lan's," *Proceedings of ACM SIGCOMM '94*, vol. 24, pp. 212–25, Oct. 1994.
- [10] IEEE Std 802.11 - 1997, *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, 1997.
- [11] C. Fullmer and J. Garcia-Luna-Aceves, "Floor acquisition multiple access (FAMA) for packet-radio networks," *Proceedings of ACM SIGCOMM '95*, 1995.
- [12] J. Deng and Z. Haas, "Dual busy tone multiple access (DBTMA): A new medium access control for packet radio networks," *IEEE ICUPC*, pp. 973–7, Oct. 1998.
- [13] T. Rappaport, *Wireless Communications: Principles and Practice*, Prentice Hall, New Jersey, 1996.
- [14] J. Proakis, *Digital Communications*, McGraw Hill, New York, 1995.
- [15] R. Blahut, *Algebraic Codes For Data Transmission*, draft of book produced at The University of Illinois IEEE Student Branch, Urbana, IL, 1999.
- [16] K. Cheng and S. Wu, "A programmable architecture for OFDM-CDMA," *IEEE Communications Magazine*, vol. 37, Nov. 1999.
- [17] V. Bharghavan, "Performance evaluation of a medium access protocol for wireless packet networks," *IEEE Performance and Dependability Symposium*, pp. 86–95, Aug. 1998.
- [18] "CMU Monarch Project," <http://www.monarch.cs.cmu.edu/>, Dec. 1999.