A Distributed Transmission Power Control Protocol for Mobile Ad Hoc Networks

Alaa Muqattash and Marwan Krunz

Department of Electrical and Computer Engineering
The University of Arizona
Tucson, AZ 85721

{alaa,krunz}@ece.arizona.edu

***Abstract*— In this paper, we propose a comprehensive solution for power control in mobile ad hoc networks (MANETs). Our solution emphasizes the interplay between the MAC and network layers, whereby the MAC layer indirectly influences the selection of the next-hop by properly adjusting the power of *route request* packets. This is done while maintaining network connectivity. Channel-gain information obtained mainly from overheard RTS and CTS packets is used to dynamically construct the network topology. Unlike the IEEE 802.11 approach and previously proposed schemes, ours does not use the RTS/CTS packets to silence the neighboring nodes. Instead, collision avoidance information is inserted in the CTS packets and sent over an out-of-band control channel. This information is used to dynamically *bound* the transmission power of potentially interfering nodes in the vicinity of a receiver. By properly estimating the required transmission power for data packets, our protocol allows for interference-limited simultaneous transmissions to take place in the neighborhood of a receiving node. Simulation results indicate that compared to the IEEE 802.11 approach, the proposed protocol achieves a significant increase in the channel utilization and end-to-end network throughput, and a significant decrease in the total energy consumption.**

***Index Terms*—Power control, ad hoc networks, energy efficient routes, IEEE 802.11, interference margin.**

I. Introduction

Mobile ad hoc networks (MANETs) are multi-hop networks in which mobile nodes cooperate to maintain network connectivity and perform routing functions. These fast deployable, self-organizing networks are typically used in situations where network connectivity is temporarily needed or where it is infeasible (or expensive) to install a fixed infrastructure network. Power control in MANETs has recently received a lot of attention for two main reasons. First, power control has been shown to increase spatial channel reuse, hence increasing the overall (aggregate) channel utilization [15]. This issue is particularly critical given the ever-increasing demand for channel bandwidth in wireless environments. Second, power control improves the overall energy consumption in a MANET, consequently prolonging the lifetime of the network. Portable devices are often powered by batteries with limited weight and lifetime, and energy saving is a crucial factor that impacts the survivability of such devices.

The Distributed Coordination Function (DCF) of the IEEE 802.11 [1] standard is, by far, the most dominant MAC protocol for ad hoc networks[1]. This protocol generally follows the CSMA/CA paradigm, with extensions to allow for the exchange of RTS/CTS (request-to-send/clear-to-send) handshake packets between the transmitter and the receiver. These control packets are needed to reserve a *transmission floor* for the subsequent data packets. Nodes transmit their control and data packets at a common maximum power level, preventing all potentially interfering nodes from starting their own transmissions. Any node that hears the RTS or the CTS message defers its transmission until the ongoing transmission is over. While such an approach is fundamentally needed to avoid the hidden node problem, it negatively impacts the channel utilization by not allowing concurrent transmis-

This work was supported in part by the National Science Foundation under grants ANI 9733143, CCR 9979310, and ANI 0095626; and by the Center for Low Power Electronics (CLPE) at the University of Arizona. An abridged version of this paper was presented at the *IEEE Infocom 2003 Conference*, San Francisco, April 1–3,2003.

[1] In addition to the DCF, the 802.11 standard also supports a Point Coordination Function (PCF), which is essentially a polling scheme that is intended for delay-sensitive traffic. The DCF can be deployed in both the *Ad Hoc* and the *Infrastructure* modes. The former is assumed in this paper.

sions to take place over the reserved floor. This situation is exemplified in Figure 1, where node A uses its maximum transmission power to send its packets to node B (for simplicity, we assume omnidirectional antennas, so a node's reserved floor is represented by a circle in the 2D space). Nodes C and D hear B's CTS message and, therefore, refrain from transmitting. It is easy to see that both transmissions $A \rightarrow B$ and $C \rightarrow D$ can, in principle, take place at the same time if nodes are able to select their transmission powers in an appropriate manner. In Figure 1, the reserved floors based on the standard (fixed, maximum power) approach are indicated by dashed circles, while the ones that are based on the minimum *required* power for coherent reception are indicated by solid circles.

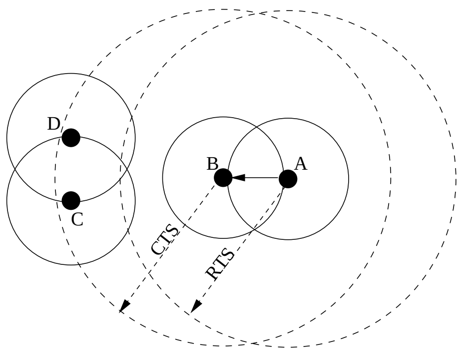

Fig. 1. Inefficiency of the 802.11 approach. Nodes A and B are allowed to communicate, but nodes C and D are not.

While the idea of power control is simple, achieving it in a distributed manner is challenging. More specifically, it is not enough to only adjust the transmission power of each transmitter according to the minimal power needed for coherent reception. This situation is shown in Figure 2, where node A has just started a transmission to node B at a power level that is just enough to ensure correct decoding at B. Suppose that node B uses the same power level to communicate with A. Nodes C and D are outside the floors of A and B, so they do not hear the RTS-CTS exchange between A and B. For nodes C and D to communicate, they have to use a power level that is reflected by the transmission floors in Figure 2 (the two circles centered at C and D). However, the transmission $C \rightarrow D$ will interfere with $A \rightarrow B$ transmission, causing a collision at B. In essence, the problem is caused by the asymmetry in the transmission floors.

From the above example, one can make the following observation: if nodes send their control (RTS-CTS) packets at a fixed power level ($P_{\max}$), but send their data packets at an adjustable (lower) power level, then the collision in the previous example can be avoided. This observation is the key to our proposal. However, to enable dynamic adjustment of the (data packet) transmission power, separate channels are needed for data and control packets. Control packets are transmitted at power level $P_{\max}$, and are received by all potentially interfering nodes, as in the IEEE 802.11 standard. However, in contrast to the IEEE 802.11, interfering nodes *may* be allowed to transmit concurrently, depending on some criteria that will be discussed later.

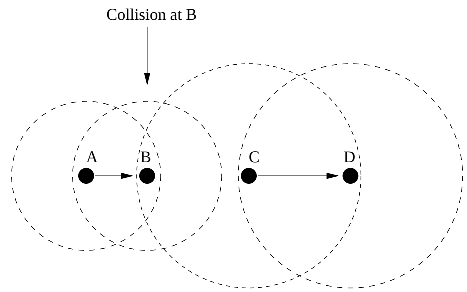

Fig. 2. Challenge of implementing power control in a distributed fashion. Node C is unaware of the communication $A \rightarrow B$, and hence it starts transmitting to node D at a power that destroys B' s reception.

Power control for MANETs has been extensively studied (see Section IV for related work). However, previously proposed protocols address the issue from a single-layer perspective, by either implementing power control with proper MAC functionality in mind (e.g., [24], [34]), or by using it as a means of controlling the connectivity and topological properties of the network (e.g., [32], [25], [27], [29]). While the two approaches may at first seem orthogonal, integrating them in one framework is, at best, highly inefficient. Consequently, none of these approaches offers a comprehensive solution to the problem. Our view is that *inter-layer dependence plays a critical role in providing an efficient and comprehensive solution to the power control problem*, and this view is a key design principle in our proposed protocol.

The rest of the paper is organized as follows. In Section II we present the proposed protocol, emphasizing the design considerations that were taken into account. The operational details are discussed in Section III. In Section IV we review related work in the area of power control for MANETs. The simulation results are presented and discussed in Section V. Finally, our main conclusions are drawn in Section VI.

II. Power Controlled Dual Channel (PCDC) Protocol

A. Channel Model and Protocol Assumptions

In designing our protocol, we assume that: (1) the channel gain is stationary for the duration of the control and the ensuing data packet transmission periods; (2) the gain between two nodes is the same in both directions; and (3) data and control packets between a pair of nodes observe similar channel gains. The justification of these assumptions follows next.

Radio channels typically exhibit large- and small-scale propagation behaviors [28]. Large-scale propagation characterizes the mean signal strength for an arbitrary transmitter-receiver separation. Such propagation behavior has no impact on the validity of our channel assumptions, since the distance and the level of clutter are the same in both directions and for both data and control channels; hence, the mean signal strength will also be the same. Also, the time needed for the RTS/CTS exchange followed by a data-packet transmission is typically in the order of tens of milliseconds. Within this time interval, very little change occurs in the locations of the mobile nodes, and consequently in the average signal strength.

Small-scale propagation characterizes the fluctuations of the received signal strength over very short time durations. These fluctuations result from multiple versions of the signal (i.e., multipath waves) arriving at the receiver at slightly different times and combining to give a resultant signal that can vary widely in amplitude and phase. Small-scale propagation may affect our protocol assumptions since signals may combine differently in both directions and for both channels. However, in a spread spectrum environment where the system spreads the signal into a relatively wide bandwidth using a pseudo-noise (PN) sequence, the receiver can exploit the multipath components to improve the performance of the system. This is accomplished by using several diversity techniques (such as RAKE receivers) that take advantage of the random nature of the signal by finding uncorrelated signal paths. Therefore, our proposed protocol relies on physical-layer techniques to mitigate the multipath effect, and in modest fading channels the assumptions will hold.

In addition to the above assumptions, we assume that the radio interface can provide the MAC layer with the average power of a received control signal as well as the average interference power. The radio interface is equipped with carrier-sense hardware that senses the control channel for any carrier signal. No carrier-sense is needed for the data channel. The control channel is further divided into two sub-channels: a RTS-CTS channel and an acknowledgement (ACK) channel. The carrier frequency spacing between the channels is enough to ensure that the outgoing signal on one channel does not interfere with the incoming signal on the other channel.

B. Protocol Overview and Design Considerations

The interaction between the network and MAC layers is fundamental for power control in MANETs. On the one hand, the power level determines who can receive the route request packets, and hence, it directly impacts the selection of the next hop. Obviously, this is a network-layer issue. On the other hand, the power level also determines the floor reserved for the node's transmission, which is a MAC-layer issue. Hence, we have to somehow introduce power control from the perspectives of both layers .

A power controlled MAC protocol reserves different floors for different uses of the channel, depending on the node's transmission power. The selection of the "best" transmission range has been investigated in the literature, but *not* in the context of collision-avoidance MAC protocols. In [16] the authors have shown that a higher network throughput can be achieved by transmitting packets to the nearest neighbor in the forward progress direction. In [15] the authors have proved that using a smaller transmission range increases network throughput. The intuition behind these results is that halving the transmission range increases the number of hops by two, but decreases the area of the reserved floor to one forth of its original value, allowing for more transmissions in the neighborhood.

In addition to improving network throughput, reducing the transmission range plays a significant role in reducing the energy consumption [12]. The power consumed by the radio frequency (RF) power amplifier of the network interface card (NIC) is directly proportional to the power of the transmitted signal, and thus, it is of great interest to control the signal transmission power to increase the operational lifetime of mobile nodes. Presently, the RF power amplifier consumes almost half (or more in the case of small computing devices such as sensors) of the total energy consumed by the NIC. This ratio is expected to increase in future NICs, as the processing components become more power efficient. Therefore, there is potential for a significant energy saving by reducing the signal transmission power. In [29] the authors have showed that power-efficient routes can be found by considering only the nodes in the "enclosure region" as potential next hops. Similar results have been provided in [32]. Another advantage of power control that has not received much attention in the literature is related to reducing the power consumption at *unintended receivers* (those who are not addressed by the transmission). Significant power is consumed in receiving a packet. For example, in the 2003 model of the Cisco Aironet 350 Series Client Adapter card operating at 5 volts, the reception power amounts to about 67% of the transmission power. Since reducing the transmission range results in a smaller number of nodes overhearing the transmission, less power will be consumed by those unintended receivers.

The above discussion provides sufficient motivation to dynamically adjust the transmission range for data packets. The question is how can a node select the lowest possible power that ensures network connectivity while simultaneously guaranteeing proper MAC functionality and introducing little overhead? Section II-C answers this question and explains how next-hop selection can be restricted by MAC-layer considerations.

Having decided on varying the transmission range, an access mechanism is required to avoid the type of collisions in Figure 2. For that, PCDC uses a modified RTS-CTS reservation mechanism. Unlike the 802.11 approach (and others, e.g., [4], [14], [18]), the RTS-CTS control signals in our scheme are not used to silence the neighbors of a receiving node. Instead, *the control signals in the proposed scheme are used to dynamically bound the transmission power of interfering nodes in the vicinity of a receiver.* The details of this mechanism are explained in Section II-D.

The third key consideration in PCDC is to provide cooperation among neighboring nodes at the MAC layer. A node that intends to transmit has to account for *potential* future transmissions in its neighborhood. This is achieved by having an *interference margin* that allows nodes at some interfering distance to start new transmissions. Nodes that are in the neighborhood may commence their transmissions if such transmissions will not disturb the ongoing ones. In Section II-E we develop a distributed strategy that dynamically adjusts the interference margin to maximize the number of simultaneous transmissions.

C. Connectivity Set

In PCDC, the MAC layer affects the performance of the network layer by controlling the power used to transmit the *route request* (RREQ) packets. These packets are broadcasted by a node to inquire about the next hop to a given destination. By controlling the transmission power of a RREQ packet, the MAC layer effectively controls the set of *candidate* next-hop nodes. From a power consumption standpoint, a smaller transmission power is preferable, which also means a smaller set of next-hop nodes. But reducing the size of this set may result in losing network connectivity. Hence, the goal is to provide a distributed mechanism by which a node can dynamically compute its *connectivity set* (CS) (defined below). From this CS, the node can then decide on the set of next-hop nodes. We now describe a localized algorithm for constructing the CS of an arbitrary node i (CS_i). This algorithm aims at producing power-efficient end-to-end routes while simultaneously maintaining network connectivity and introducing as little overhead as possible.

The intuition behind the algorithm is that CS_i must contain only the neighboring nodes with which direct communication requires less power than the indirect (two-hop) communication via any other node that is already in CS_i. To construct CS_i, node i continuously caches the estimated channel gain of every signal it receives over the control channel, *regardless of the intended destination of this signal*. Note that computing the gain is possible because control packets are transmitted at a fixed, known power, and hence, node i uses the reception power of the signal to determine the channel gain. Each node in CS_i is associated with a timer that expires T seconds from the time this node was added to CS_i. The value of T will be discussed later. If the timer expires, then the corresponding node is deleted from CS_i.

Let P_{uv} be the minimum power required to transmit a data packet from node u to node v at a given time instant. Upon receiving an RTS/CTS packet from another node, say j, node i does the following. If $j \in \mathrm{CS}_i$ *and* the newly computed channel gain matches the already stored one, then the timer associated with j's entry in CS_i is reset and no further action is taken. On the other hand, if $j \notin \mathrm{CS}_i$ or if $j \in \mathrm{CS}_i$ but the newly computed gain *does not* match the already stored one, then node i checks if $P_{ij} < P_{iu} + P_{uj}$ for every node $u \in \mathrm{CS}_i$, $u \neq j$. If so, then node j is added to CS_i; otherwise, it is not. Let $P_{\text{conn}}^{(i)} \stackrel{\text{def}}{=} \max_{j \in \mathrm{CS}_i} P_{ij}$. If node j is added to CS_i *and* $P_{ij} < P_{\text{conn}}^{(i)}$, then all other elements of CS_i must be re-examined. The reason is that a two-hop path between nodes i and u, $u \in \mathrm{CS}_i$, that goes through node j may now be more power efficient than the direct path between i and u. In this case, node u has to be deleted from CS_i. However, if $P_{ij} \geq P_{\text{conn}}^{(i)}$, then $P_{ij} + P_{ju} > P_{iu}$ for any $u \in \mathrm{CS}_i$ and hence, there is no need to re-examine CS_i. Figure 3 depicts the algorithm for updating CS_i and $P_{\text{conn}}^{(i)}$ upon the receipt of an RTS/CTS packet from node j.

The computation of P_{ij} is simple since node i estimates the channel gain G_{ij} from j's control signals (RTS/CTS/hello). However, computing the power required for indirect communication requires the exchange of additional information between one-hop neighbors, as we now explain. Each node i keeps a list N_i of all one-hop neighbors that are within the maximum transmission

UPDATE-CS(CS_i, j, $P_{\text{conn}}^{(i)}$)
for every node $u \in \text{CS}_i$ **do**
  if $P_{iu} + P_{uj} \leq P_{ij}$
    terminate UPDATE-CS
end-for
if $P_{ij} < P_{\text{conn}}^{(i)}$
  for every node $u \in \text{CS}_i$ **do**
    if $P_{ij} + P_{ju} \leq P_{iu}$
      $\text{CS}_i \leftarrow \text{CS}_i - \{u\}$
  end-for
$\text{CS}_i \leftarrow \text{CS}_i \cup \{j\}$
$P_{\text{conn}}^{(i)} \leftarrow \max\{P_{iu} : u \in \text{CS}_i\}$
end UPDATE-CS

Fig. 3. Algorithm for updating CS_i and $P_{\text{conn}}^{(i)}$ after receiving a control packet from node j.

range ($P_{\max}$) of node i. Note that $\text{CS}_i \subset N_i$. Node i also maintains the value of the minimum transmission power required to communicate *directly* with each node in N_i. The set N_i is updated dynamically whenever node i overhears a control packet (RTS, CTS, or hello) in its neighborhood. Initially, when node i comes up, it broadcasts its current N_i at the power $P_{\max}$. Subsequently, whenever a new node is added to N_i or whenever the minimum power required to communicate directly with an existing node in N_i changes by more than a threshold, say ϵ dB, then *only the updated information* is piggybacked on the next "hello" message that is broadcasted by node i at power $P_{\max}$. All nodes in N_i will receive this broadcast. Let $u \in N_i$. When u receives i's update, it uses that information along with N_u to update its connectivity set CS_u, as described in Figure 3.

An advantage of the above approach is that it accounts for the effect of shadowing (as part of the large-scale channel variations), independent of any specific propagation model. The communication overhead is relatively low since changes in channel gains due to shadowing occur on the time scale of seconds (channel gain is a characteristic of large-scale models). Furthermore, these channel gains are broadcasted locally, and are *not* flooded beyond the maximum-power neighborhood of a node.

In deciding whether to add node j to CS_i or not, we only considered the two-hop indirect paths. The reason is that if the two-hop path is less power-efficient than the direct path, then so are the L-hop paths, $L \geq 2$. We now prove this claim for the case $L = 3$, and the general case follows by induction. Suppose that node i has just heard a control signal from node j and that $P_{ix} + P_{xj} > P_{ij}$ for all $x \in \text{CS}_i$. We now show that $P_{iu} + P_{uv} + P_{vj} > P_{ij}$ for any nodes u and v in CS_i. The proof is by contradiction, i.e., suppose that $P_{iu} + P_{uv} + P_{vj} \leq P_{ij}$ for some nodes u and v in CS_i. Then the communication $i \rightarrow u \rightarrow v$ must require less power than $i \rightarrow v$, and hence, node v cannot be in CS_i. This contradicts the assumption that $v \in \text{CS}_i$.

As mentioned in Section II-B, maintaining network connectivity is crucial. The following theorem shows that if the network is connected under the standard maximum-power approach, then it must also be connected when each node communicates only with nodes in its connectivity set.

Theorem 1: Let $G = (V, E)$ be the undirected graph that results from using the power $P_{\max}$ to reach other nodes. Let $H = (V, E')$ be the undirected graph constructed based on the CS approach. If G is connected, then H is also connected.

Proof: See Appendix.

One nice feature of the algorithm is its symmetrical property: if $i \in \text{CS}_j$ then $j \in \text{CS}_i$, and vice versa. The reason for this property is that if the direct path from node i to node j is more power efficient than any other path, then so is the direct path from j to i.

At high loads, there is enough RTS-CTS activity to allow for the computation of the connectivity set at no extra bandwidth overhead. However, at light loads, channels are often idle, and an auxiliary scheme is needed to ensure accurate computation of the connectivity set. In our protocol, we let each node broadcast a "hello" packet over the data channel at power $P_{\max}$ every Δ seconds, where Δ is a random variable that is uniformly distributed in the interval $[T/2, T]$. Randomization is needed to avoid collisions between synchronized "hello" transmissions. The value of T is determined according to the overall mobility pattern in the network. For example, for conference room scenarios, the network topology hardly changes within a 3-second interval, so T can be set to, say, 4 seconds. The format of the "hello" packet is similar to that of the IEEE 802.11 CTS packet, except for two changes. First, the address field in the IEEE 802.11 CTS packet that indicates the receiver address is now used to indicate the transmitter address. Second, the duration field in the standard CTS packet is now used for a different purpose, which will be explained in Section II-E. Figure 4 shows the format of the "hello" packet. Note that initially the CS of a node is empty. However, it takes only T seconds in the worst case for the node to discover its neighborhood and start using a reduced power. The above "hello" approach incurs little overhead. This is in contrast to the scheme in [32], where periodic or on-demand reconfiguration of the network topology is always needed if nodes are moving (the

Octets:	2	6	2	4
	Frame Control	Transmitter Address	Reserved Field	FCS (32–bit CRC)

Fig. 4. Format of the "hello" packet in PCDC.

authors simulated only a static network). This affects network resources and increases packet delays, especially at peak load times. Our scheme, on the other hand, exploits information freely available through the control channel at those peak times.

Now that node i has computed the connectivity power $P_{\text{conn}}^{(i)}$, *it uses this power level to broadcast its RREQ packets*. This results in two significant improvements. First, any simple min-hop routing protocol, such as AODV or DSR, can now be used to produce routes that are very power efficient and that increase network throughput (i.e., reduce the total reserved floor). Hence, no intelligence is needed at the network layer and no link information (e.g., power) has to be exchanged or included in the RREQ packets in order to find power-efficient routes. Second, considering how RREQ packets are flooded throughout the network, significant improvements in throughput and power consumption can be achieved by limiting the broadcasting of these packets to nodes that are within the connectivity range $P_{\text{conn}}^{(i)}$. Take, for example, the network in Figure 5 (this topology approximates a classroom environment). Suppose that DSR is used for route discovery. Consider first the "standard" approach, whereby RREQ packets are transmitted at power P_{max}. If node A intends to send a packet to node D, it broadcasts a RREQ packet at P_{max}. Upon receiving A's RREQ packet, node B searches its route cache for the next hop to the destination node D. If no route is found, node B forwards the RREQ packet to its neighbors after adding its own address. Note that all nodes in A's maximum transmission range will perform the same procedure. The RREQ packet propagates through the network until it reaches the destination or a node with a route to the destination. Simulation results in [17] show that for DSR, the overhead of RREQ packets in *bytes* is approximately 38% of the total received bytes. Obviously, this overhead and the corresponding power consumption are significant. A close look at DSR reveals that these drawbacks become more significant as the range over which the RREQ packets are broadcasted is made larger. First, as this range increases, the number of receivers that receive multiple inquiries for the same destination also increases. As we pointed out earlier, a significant amount of energy is consumed in just receiving a transmission. More significantly, following the transmission of a RREQ packet, there will be a high contention period over the channel between nodes that intend to propagate the RREQ. This results in many collisions between RREQ packets (the transmissions of which are typically unacknowledged), which delays the process of finding the destination and requires retransmitting these packets. In contrast, in PCDC the RREQ packet is broadcasted to the connectivity set only (and not to the maximum reachable set) and hence, the number of contenders following a RREQ does not vary significantly, making it possible to design an efficient contention window for RREQ packets. Therefore, in the process of finding the destination, PCDC results in lower overhead, less contention, and less consumed power.

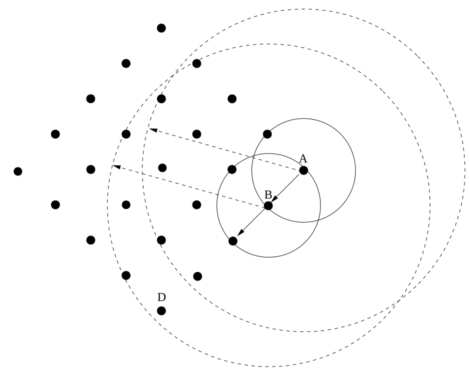

Fig. 5. Example showing the inefficiency of broadcasting RREQ packets at the maximum power.

D. Channel Access Mechanism

We now describe the admission control and channel access strategy in the PCDC protocol. RTS and CTS packets are used to provide three additional functions. First, they allow nodes to estimate the channel gains between transmitter-receiver pairs. Second, a receiver i uses the CTS packet to notify its neighbors of the additional interference power (denoted by $P_{\text{noise}}^{(i)}$) that each of the neighbors can add to node i without impacting i's current reception (hence, allowing for interference-limited concurrent transmissions). These neighbors constitute the set of *potentially interfering* nodes. Finally, each node keeps listening to the control channel regardless of the signal destination in order to keep track of its connectivity set, as explained in Section II-C. These functions are now explained in detail.

If node j has a packet to transmit, it sends an RTS packet over the control channel at power P_{max} and includes in this packet the maximum *allowable* power level ($P_{\text{map}}^{(j)}$) that node j is allowed to use without disturbing any ongoing reception in its neighborhood. The exact computation of this power will be discussed shortly. The format of the RTS packet is similar to that of the IEEE 802.11 except for an additional one-byte field that indicates the value of $P_{\text{map}}^{(j)}$.

Upon receiving the RTS packet, the intended receiver, say node i, uses the known P_{max} value and the power of the received signal $P_{\text{received}}^{(ji)}$ to estimate the channel gain $G_{ji} = P_{\text{received}}^{(ji)}/P_{\text{max}}$ between nodes i and j at that time (note that we assume channel reciprocity, and so $G_{ij} = G_{ji}$). Accordingly, node i will be able to correctly decode the data packet if this packet was transmitted at power $P_{\text{min}}^{(ji)}$ given by:

$$P_{\text{min}}^{(ji)} = \frac{\text{SNR}_{\text{th}}(P_{\text{thermal}} + P_{\text{MAI-current}}^{(i)})}{G_{ji}} = \frac{\text{SNR}_{\text{th}}\eta^{(i)}}{G_{ji}} \tag{1}$$

where SNR_{th} is the minimum SNR ratio that is needed to achieve the target bit error rate at that receiver (we assume SNR_{th} is the same for all nodes, i.e., all nodes require the same QoS), P_{thermal} is the thermal noise power, $P_{\text{MAI-current}}^{(i)}$ is the *current* multiple access interference (MAI) from all already ongoing (interfering) transmissions, and $\eta^{(i)} \stackrel{\text{def}}{=} P_{\text{thermal}} + P_{\text{MAI-current}}^{(i)}$. Note that because of the assumed stationarity in the channel gain over small time intervals, G_{ji} is approximately constant throughout the transmissions of the control packet and the ensuing data packet.

The value of $P_{\text{min}}^{(ji)}$ in (1) is the minimum power that node j must use for data transmission in order for node i to correctly decode the data packet *at the current level of interference*. This $P_{\text{min}}^{(ji)}$, however, does not allow for any interference tolerance at node i, thus all neighbors of node i will have to defer their transmissions during node i's ongoing reception (i.e., no simultaneous transmissions can take place in the neighborhood of i). To allow for a number of *future* interfering transmissions to take place in its vicinity, receiver i requests that node j scales up the transmission power $P_{\text{min}}^{(ji)}$ by the factor $\alpha_{\text{min}}^{(i)}$, where $\alpha_{\text{min}}^{(i)} \geq 1$. Therefore, node j must use a transmission power given by:

$$P_{\text{requested}}^{(ji)} = \alpha_{\text{min}}^{(i)} P_{\text{min}}^{(ji)}. \tag{2}$$

The computation of $\alpha_{\text{min}}^{(i)}$ will be explained in the next section. If $P_{\text{map}}^{(j)} < P_{\text{requested}}^{(ji)}$, then node i does not send a CTS, since the requested transmission power may cause a collision in the neighborhood of j. On the other hand, if $P_{\text{map}}^{(j)} > P_{\text{requested}}^{(ji)}$, then node j can transmit to node i using $P_{\text{requested}}^{(ji)}$ without disturbing any of the ongoing receptions in j's vicinity. Scaling up the transmission power by $\alpha_{\text{min}}^{(i)}$ amounts to "inflating" the total interference $\eta^{(i)}$ in (1) by the same factor. So the additional interference that node i can tolerate from *future* unintended transmissions is given by:

$$P_{\text{MAI-future}}^{(i)} = \eta^{(i)}(\alpha_{\text{min}}^{(i)} - 1). \tag{3}$$

We refer to $P_{\text{MAI-future}}^{(i)}$ as the interference margin at node i.

The next step is to equitably distribute this power tolerance among future potentially interfering users in the vicinity of i. The rational behind this distribution is to prevent one neighbor from consuming the entire $P_{\text{MAI-future}}^{(i)}$. In other words, we think of $P_{\text{MAI-future}}^{(i)}$ as a network resource that should be shared among the neighbors of i. Let $K^{(i)}$ be the number of nodes in the vicinity of i that are to share $P_{\text{MAI-future}}^{(i)}$. This number is determined as follows. Node i keeps track of the instantaneous number of simultaneously active transmissions (i.e., load) in its neighborhood, which we donate by $K_{\text{inst}}^{(i)}$. This can be easily achieved by monitoring the RTS/CTS exchanges over the control channel. In addition, i keeps track of a moving average of $K_{\text{inst}}^{(i)}$, denoted by $K_{\text{avg}}^{(i)}$. Then, $K^{(i)}$ is calculated as follows:

$$K^{(i)} = \begin{cases} \beta(K_{\text{avg}}^{(i)} - K_{\text{inst}}^{(i)}), & \text{if } K_{\text{avg}}^{(i)} > K_{\text{inst}}^{(i)} \\ \beta, & \text{otherwise} \end{cases} \tag{4}$$

where $\beta > 1$ is a safety margin. The rationale behind (4) is as follows. At a given time instant, there are $K_{\text{inst}}^{(i)}$ active transmissions in the neighborhood of i, the interference of which must have been accounted for in the current value of $\eta^{(i)}$. Note that before the start of their transmissions, these $K_{\text{inst}}^{(i)}$ interferers were accounted for in $P_{\text{MAI-future}}^{(i)}$, but once they have started their transmissions, their interference is now part of $\eta^{(i)}$. This leaves $K_{\text{avg}}^{(i)} - K_{\text{inst}}^{(i)}$ potential future interferers to share the current value of the interference margin. As $K_{\text{inst}}^{(i)}$ increases beyond $K_{\text{avg}}^{(i)}$, there are fewer inactive neighboring nodes that could become potential future interferers. We limit the number of such interferers to β.

Accordingly, the interference tolerance $P_{\text{noise}}^{(i)}$ that *each* future neighbor can add to node i is given by:

$$P_{\text{noise}}^{(i)} = \frac{P_{\text{MAI-future}}^{(i)}}{K^{(i)}}. \tag{5}$$

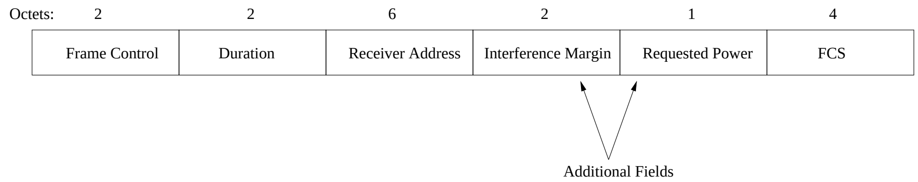

Fig. 6. Format of the CTS packet in the proposed protocol.

When responding to j's RTS, node i indicates in its CTS the power level $P^{(ji)}_{\text{requested}}$ that j must use. In addition, node i inserts $P^{(i)}_{\text{noise}}$ in the CTS packet and sends this packet back to node j at $P_{\max}$ over the control channel. The format of the CTS packet is shown in Figure 6.

A potentially interfering node, say s, that hears the CTS message uses the signal strength of the received CTS to compute the channel gain G_{si} between itself and node i. The channel gain along with the broadcasted $P^{(i)}_{\text{noise}}$ values are used to compute the maximum power $P^{(s)}_{\text{map}}$ that s can use in its future transmissions. This is the power that node s can use in its future transmissions that will not add more than $P^{(i)}_{\text{noise}}$ to the received noise at node i. Furthermore, $P^{(s)}_{\text{map}}$ is updated *dynamically* whenever s overhears a new CTS, and is taken as the minimum of the $P^{(k)}_{\text{noise}}/G_{sk}$ values, for all neighbors k of s. Note that it is possible for more than $K^{(i)}$ nodes to start transmitting during i's reception and this may result in MAI at i that is greater than $P^{(i)}_{\text{MAI-future}}$. We address this issue in Section II-G.

Obviously, the interference range is larger than the reception range set by $P_{\max}$ (in theory, any unintended transmission causes some interference), so collisions may still occur because of interferers outside the reception range. This problem is present in the IEEE 802.11 scheme as well. But compared to the 802.11 scheme, PCDC significantly *reduces* the severity of this problem, for the following reasons:

(i) According to the PCDC protocol, a node, say j, can communicate directly only with neighbors who are in j's connectivity set (CS$_j$), which is associated with the transmission power $P^{(j)}_{\text{conn}}$. This power is typically much smaller than $P_{\max}$. So when node j communicates with another node i (which must be in CS$_j$), it uses transmission power P_{ji} that is typically much less than $P_{\max}$. Since P$_{ji} < P_{\max}$, the interference that node j causes to nodes that are outside j's $P_{\max}$ range can be considered small, and vice versa. This is in contrast to the 802.11 approach, where nodes can communicate with any other node within the $P_{\max}$ range.
We use the example in Figure 7 to illustrate the above argument. Part (a) of the figure shows node A communicating with node B, and node C (which is just outside the *reception range* of node B) communicating with node D, all according to the IEEE 802.11 approach. The distance between B and the intended transmitter A is comparable to the distance between B and the interfering transmitter C. Since C is transmitting at power $P_{\max}$, the chances that this situation will cause a collision at node B are high.
Now, consider part (b) of Figure 7, where again A communicates with B and C with D, this time according to the proposed PCDC protocol. In this case, since A can only communicate with nodes within its connectivity set, the distance between B and the intended transmitter A is much larger than the distance between B and the interfering transmitter C. Furthermore, the interfering transmitter C is transmitting at a power that is much less than $P_{\max}$. So the interference that C causes at B is much less than in the case of the IEEE 802.11 scheme.

(ii) Let $P^{(i)}_{\text{MAI-other}}$ denote the part of the multiple access interference (MAI) at receiver i that is attributed to nodes outside i's $P_{\max}$ range. $P^{(i)}_{\text{MAI-other}}$ fluctuates as neighboring nodes finish their transmissions or start new ones. However, since $P^{(i)}_{\text{MAI-other}}$ is caused by a *large* number of interfering nodes (all nodes between $P_{\max}$ and ∞), we can assume, using the law of large numbers, that the mean value of $P^{(i)}_{\text{MAI-other}}$ is almost fixed. Now, this mean has already been accounted for in Equation (2) in the paper (in $P^{(i)}_{\text{MAI-current}}$; the current noise plus interference at node i). In other words, the receiver will request that the transmitter uses a power to combat $P^{(i)}_{\text{MAI-current}}$, which already includes the mean of $P^{(i)}_{\text{MAI-other}}$.

(iii) Node i receives the signal at a scaled up power level, allowing for some interference margin. If the interference power goes above the margin, then node i can respond with a special CTS packet over the control channel, preventing the RTS sender from commencing its transmission.

The approach we discussed in this section provides a distributed mechanism for admission control. In contrast

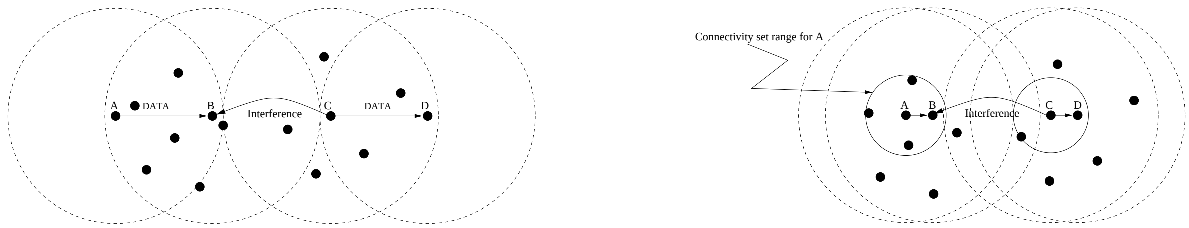

(a) IEEE 802.11 scheme.

(b) Proposed PCDC protocol.

Fig. 7. Impact of interference (the maximum power of each node is indicated by a dashed circle centered at that node, while in part (b), the solid circles indicate the connectivity set ranges of transmitters A and C).

to cellular systems where the base station makes the admission decision, in here each node, and depending on previously heard RTS and CTS packets, decides whether its transmission can proceed or not.

E. Interference Margin

In this section, we show how the scaling factor $\alpha_{\text{min}}^{(.)}$ is computed dynamically. In [5] a power-control algorithm was proposed for the uplink channel of a DS-CDMA cellular system. The purpose of that algorithm is to maintain the QoS of ongoing users while simultaneously maximizing the free capacity for new users. We propose a distributed algorithm to implement the idea in MANETs. First, note that the SNR at a receiving node i is given by:

$$\text{SNR}^{(i)} = \frac{P_j^{(i)}}{P_{\text{thermal}} + P_{\text{MAI-current}}^{(i)} + P_{\text{MAI-future}}^{(i)}} \tag{6}$$

where $P_j^{(i)}$ is the "desired" power at the receiver i from the intended transmitter j (see the previous section for the definition of the other variables). It was proven in [5] that to increase channel capacity, $P_{\text{MAI-future}}^{(i)}$ must be increased (and so, from (3), $\alpha_{\text{min}}^{(i)}$ must also be increased). The authors proposed an algorithm that scales up the power of active links (transmissions in progress) by the largest possible constant α. This constant α is calculated to accommodate the user with the maximum ratio of the currently used power over the peak power imposed by the hardware. If α is made larger than that, then at least one of the users will be peak-power limited (i.e., reaches its maximum power) and will be unable to attain its QoS.

The authors in [5] presented a centralized algorithm that implements the aforementioned power-scaling idea at the base station. Applying the same algorithm in MANETs is not so straightforward due to the absence of a centralized control. Moreover, in a MANET, the channel consists of overlapping regions where nodes do not hear all transmitted signals. This means that the power received at two different nodes consists of the power signals received from two different sets of transmitters. To account for these differences we treat the problem in a slightly different manner. First, while in the cellular scenario the base station applies the algorithm only to active users, in our case the notion of "users" is different, as it refers to the *expected number of future users*. Second, in our case, each node i tries to accommodate nodes that are within its own maximum transmission range, since those are the nodes which node i may interfere with.

To implement power scaling in a distributed manner, node i uses its dynamically computed connectivity power $P_{\text{conn}}^{(i)}$ to compute the maximum scaling constant $\alpha^{(i)}$ that node i can accommodate:

$$\alpha^{(i)} = \frac{P_{\text{max}}}{P_{\text{conn}}^{(i)}} \tag{7}$$

This value represents the maximum scaling constant that node i can be asked to use. A larger value implies that one node will have to transmit at power greater than P_{max}, which is not possible. Note here that a more clustered topology would result in a larger interference margin, and hence, more simultaneous transmissions.

While the maximum available capacity for prospective transmitters can be achieved by maximizing $\alpha^{(i)}$, this has a negative effect on the node's battery life. The reasons is because the transmitter scales up the transmission power by $\alpha^{(i)}$, thus, as $\alpha^{(i)}$ increases, more energy is consumed to deliver the packet. To account for these two conflicting goals, we use the ratio of the remaining energy $(E_{\text{remain}}^{(i)})$ to the full energy $(E_{\text{full}}^{(i)})$ of the battery to scale down the value of $\alpha^{(i)}$ as follows:

$$\alpha_{\text{eff}}^{(i)} = \max\left\{ 1,\ \alpha^{(i)\frac{\left\lfloor \frac{\xi \times E_{\text{remain}}^{(i)}}{E_{\text{full}}^{(i)}} \right\rfloor}{\xi}} \right\} \tag{8}$$

where ξ is a pre-specified positive integer, say 4. Note that

$$\alpha^{(i)} \frac{\left\lfloor \frac{\xi \times E_{\text{remain}}^{(i)}}{E_{\text{full}}^{(i)}} \right\rfloor}{\xi} = \begin{cases} 0 & \text{if } 0 \leq E_{\text{remain}}^{(i)}/E_{\text{full}}^{(i)} < 1/\xi \\ \alpha^{(i)}/\xi & \text{if } 1/\xi \leq E_{\text{remain}}^{(i)}/E_{\text{full}}^{(i)} < 2/\xi \\ 2\alpha^{(i)}/\xi & \text{if } 2/\xi \leq E_{\text{remain}}^{(i)}/E_{\text{full}}^{(i)} < 3/\xi \\ \vdots & \\ (\xi-1)\alpha^{(i)}/\xi & \text{if } (\xi-1)/\xi \leq E_{\text{remain}}^{(i)}/E_{\text{full}}^{(i)} < 1 \\ \alpha^{(i)} & \text{if } E_{\text{remain}}^{(i)}/E_{\text{full}}^{(i)} = 1 \end{cases} \quad (9)$$

The value of ξ is battery dependent, and can be selected by the system designer to reflect any given battery model (different values of ξ can be used by different nodes in the same network). Other forms of utility functions can be used to control the throughput/battery life tradeoff (e.g., exponentially decreasing the value of $\alpha_{\text{min}}^{(i)}$ to one as $E_{\text{remain}}^{(i)}/E_{\text{full}}^{(i)}$ approaches zero, etc.). Note also that $\alpha_{\text{eff}}^{(i)}$ must be greater than or equal to one, or otherwise coherent reception at node i is not possible. Node i broadcasts the value of $\alpha_{\text{eff}}^{(i)}$ in the reserved field of the "hello" packets mentioned in Section II-C. The value of $\alpha_{\text{min}}^{(i)}$ is set to the minimum of the $\alpha_{\text{eff}}^{(\cdot)}$ values that node i receives from its neighbors, i.e., $\alpha_{\text{min}}^{(i)} = \min_{j \in N_i} \alpha_{\text{eff}}^{(j)}$. The intuition is that if the scaling factor is made larger than $\alpha_{\text{min}}^{(i)}$, then at least one of the nodes that is within the maximum range of node i will be peak-power limited (or battery limited) and will be unable to attain its QoS while conserving its battery energy *if* it needs to start a communication with one of its connectivity set neighbors.

F. Link Layer Reliability

Providing link-layer error control is important not only because it provides faster recovery than transport-layer error control, but also because the performance of traditional transport layer protocols (such as TCP) degrades significantly over wireless links, resulting in a large number of unnecessary retransmissions [8]. This can reduce throughput, incur unacceptable delays, and consume battery energy.

The protection of ACK packets was addressed in previous MAC protocols (e.g., [1], [6]), but in the absence of power control. For example, in the IEEE 802.11 standard, a node that hears an RTS packet must defer its transmission, since it may destroy the reception of the ACK at the sender. While such an approach is fundamentally needed to protect the ACK, it reserves the floor around the transmitter for the whole duration of the data and ACK transmissions, when, in fact, the floor needs to be reserved for the duration of the ACK packet only. In practice, the ACK transmission period is small compared to the data-packet duration (≈ 1%). Hence, we propose the use of a second control channel for sending ACK messages.

In our scheme, if a node, say i, hears an RTS that is intended for some other node, then node i defers from transmitting *over the ACK control channel* for the duration of an ACK packet. This deference duration starts right after the end of transmission of the data packet (computed from the information in the RTS). In case of two neighboring nodes that start their data receptions at different times but complete them at the same time, the one with the later start-of-reception must wait for the duration of an ACK packet before acknowledging the receipt of the data packet.

Although PCDC uses a collision avoidance backoff algorithm similar to the IEEE 802.11b standard, more sophisticated backoff algorithms such as the one in [7] can also be used.

G. Protocol Recovery

In [11] the authors observed that when the transmission and propagation times of control packets are long, the likelihood of a collision between a CTS packet and an RTS packet of another contending node increases dramatically; the vulnerable period being twice the transmission duration of a control packet. At high loads, such a collision can lead to collisions with data packets, as illustrated in Figure 8. Suppose that node D starts sending a RTS to node C while C is receiving B's CTS that is intended to A. A collision happens at C, and hence, C is unaware of B's subsequent data reception. Afterwards, if C decides to transmit a CTS to D, it will destroy B's reception. Another problem that was mentioned earlier is if the interference power goes above $P_{\text{MAI-future}}^{(\cdot)}$. In PCDC, we solve the above two problems as follows. If while receiving a data packet, node i hears over the control channel a RTS message (destined to any node) that contains an allowable power $P_{\text{map}}^{(\cdot)}$ value that if used could cause an unacceptable

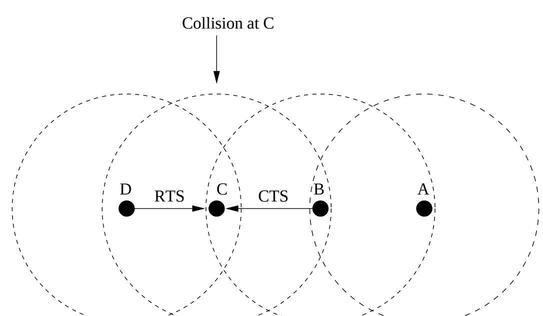

Fig. 8. Example of a collision between control packets that eventually leads to a collision with a data packet.

interference with node i's ongoing reception, then node i shall respond *immediately* with a special CTS packet over the control channel, preventing the RTS sender from commencing its transmission. The duration field of the CTS packet contains the time left for node i to finish its ongoing reception. To see how this solution helps in reducing the likelihood of collisions with data packets, consider the situation in Figure 8. Suppose that node A sends a RTS to node B, and B responds back with a CTS that collides at C with a RTS from node D. Now, C does not know about B's ongoing reception. Two scenarios can happen. In the first, node C may later wish to send a packet to, say, node D. It sends a RTS, which will be heard by node B. Node B responds back to node C with a special CTS. Note that there is a good chance that B's special CTS will collide with the CTS reply from D; however, this is desirable since C will fail to recover D's CTS packet, and will therefore defer its transmission and invoke its backoff procedure. In essence, B's special CTS acts as a jamming signal to prevent C from proceedings with its transmission. The second possible scenario is that D (or any other node that is out of the maximum range of node B) may send a new RTS to node C. Node C will respond to node D with a CTS, and D will start sending data to node C. Simultaneously, node A may be sending to B, without any collision. This is possible because in PCDC, DATA and RTS/CTS packets are sent on separate channels.

Note that in PCDC we try to avoid highly probable collision scenarios like the one mentioned in [11]. However, there will still be few complicated (and definitely much less probable) scenarios where data packets may collide; recovery from such collisions is left to the upper layers.

III. PCDC at Work

In this section, we provide the operational details of PCDC. As is the case in practice, we assume that each node i has M transmission power levels $P^{(i)}[1], \ldots, P^{(i)}[M]$ (for example, CISCO 350 series Aironet has six levels [3]). Each node i maintains the following variables:

1) $\text{NAV}^{(i)}[x]$, $x = 0, 1, \ldots, M$ (data channel virtual-sense mechanisms for node i). This is the amount of time during which node i is not allowed to use the power level $P^{(i)}[x]$. It is implemented as a counter that counts down to zero at a uniform rate. The $\text{NAV}^{(i)}[x]$ values constitute the network allocation vector NAV for node i, which is similar to the one used in the IEEE 802.11 standard [1].
2) $\text{NAV}^{(i)}_{\text{cont}}$ (control channel virtual-sense mechanism for node i). This is the time period during which the control channel will be busy (another counter).
3) $K^{(i)}$: The expected number of simultaneous transmissions around node i, as computed in (4).
4) CS_i: A table that contains the CS of node i. For each node $v \in \text{CS}_i$, the table contains v's address, the channel gain G_{vi}, and the associated time stamp.
5) N_i: A table that contains the neighbors of node i. For each node $v \in N_i$, the table contains v's address, the channel gain G_{vi}, and the associated time stamp.

Table I lists the parameters used in the rest of this section. The actions taken by a node at various stages of its operation are as follows:

$P^{(i)}_{\text{recv\_cont}}$	Received power of a control packet at node i
T_{data}	Duration of a data packet
T_{CTS}	Duration of a CTS packet
T_{ACK}	Duration of an acknowledgement packet
T_{PROP}	Maximum propagation delay
T_{SIFS}	Short interframe space[1]
T_{DIFS}	Distributed interframe space[1]

TABLE I
Parameters used in implementing the PCDC protocol.

Step 1: When node j intends to send a packet to node i:

- Node j must wait for both the physical carrier-sense mechanism and the $\text{NAV}^{(j)}_{\text{cont}}$ to indicate that the control channel has been idle for a duration of T_{DIFS} . After that, node j generates a random backoff period for an additional deferral time before transmitting (unless the backoff timer already contains a nonzero value). If the medium is determined to be busy at any time during a backoff slot, then the backoff procedure is frozen and only allowed to resume after the medium is determined to be idle for a T_{DIFS} interval. At this point, node j sends an RTS packet over the control channel at $P_{\max}$. In that packet, j includes its $P^{(j)}_{\text{map}}$ and the time duration T_{data} of the yet to be transmitted data packet.
- After transmitting the RTS packet, node j sets a timer to a timeout value of $2T_{\text{PROP}} + T_{\text{SIFS}} + T_{\text{CTS}}$ seconds. This value is the sum of the time for the RTS packet to reach the destination (T_{PROP}), the time the receiver must wait before sending back the CTS (T_{SIFS}), the time it takes the CTS to reach the sender (T_{PROP}), and the CTS transmission duration (T_{CTS}). If after this period node j has not received a correct CTS packet, it concludes that the transmission of the RTS

[1] As defined in the IEEE 802.11 standard [1].

has failed, and hence, it invokes its backoff procedure.

Step 2: When a node n receives j's RTS that is intended for node i:

- It updates its $\text{NAV}_{\text{cont}}^{(n)}$ as follows: $\text{NAV}_{\text{cont}}^{(n)} = 2T_{\text{PROP}} + T_{\text{SIFS}} + T_{\text{CTS}}$, which is the maximum duration it takes receiver i to respond with a CTS.
- It computes the channel gain $G_{nj} = \frac{P_{\text{recv\_cont}}^{(n)}}{P_{\text{max}}}$, and uses this value to update both N_n and CS_n, as explained in Section II-C.

Step 3: When node i receives an intended RTS packet from node j:

- It computes the channel gain $G_{ij} = \frac{P_{\text{recv\_cont}}^{(i)}}{P_{\text{max}}}$, and uses this value to compute the power $P_{\text{requested}}^{(ji)}$ that the data packet should be sent at using (2).
- If $P_{\text{requested}}^{(ji)} < P_{\text{map}}^{(j)}$ (which is provided in the RTS packet), then node i sends a CTS packet over the control channel at P_{max}. The CTS is transmitted after a T_{SIFS} period if both the physical carrier-sense mechanism and $\text{NAV}_{\text{cont}}^{(i)}$ indicate that the control channel is idle. The CTS packet includes the interference margin $P_{\text{noise}}^{(i)}$, the requested power $P_{\text{requested}}^{(ji)}$, and the length of the yet to be received data packet T_{data} (copied from the RTS packet).
- After transmitting the CTS packet, node i sets a timer to $2T_{\text{PROP}} + T_{\text{SIFS}}$ seconds. If after this time node i has not started receiving a data packet (recognized from its header), it concludes that the transmission of the CTS has failed. At this point, node i sends a control packet announcing the release of the channel. This packet is just another CTS packet with zero in the duration field. Note that although the transmitter sends the data immediately after receiving the CTS packet, a T_{SIFS} period is needed for the receiver i to process the packet header.

Step 4: When an irrelevant node m receives i's CTS:

- It updates its $\text{NAV}_{\text{cont}}^{(m)}$ as follows: $\text{NAV}_{\text{cont}}^{(m)} = T_{\text{PROP}}$ which allows the CTS to reach back to sender j.
- It computes the channel gain $G_{mi} = \frac{P_{\text{recv\_cont}}^{(m)}}{P_{\text{max}}}$ and uses this value to update both N_m and CS_m, as described before.
- It finds the maximum power level that it can use without adding more than $P_{\text{noise}}^{(i)}$ to the ongoing reception at node i. This is calculated as follows:

$$P_{\text{map}}^{(m)} = \frac{P_{\text{noise}}^{(i)}}{G_{mi}}$$

Accordingly, node m updates its $\text{NAV}^{(m)}[x]$ vector as follows:

$$\text{NAV}^{(m)}[x] = 2T_{\text{PROP}} + T_{\text{data}} \quad \forall x : P^{(m)}[x] > P_{\text{map}}^{(m)}$$

but only if $\text{NAV}^{(m)}[x]$ already contains a smaller value.

Step 5: When node j receives i's CTS:

- It transmits the data packet at power level $P_{\text{requested}}^{(ji)}$ (included in the CTS packet).
- It computes the channel gain $G_{ij} = \frac{P_{\text{recv\_cont}}^{(j)}}{P_{\text{max}}}$ and uses this value to update both N_j and CS_j, as described before.
- After transmitting the data packet, node j sets a timer to $2T_{\text{PROP}} + T_{ACK}$ seconds. If after this period node j has not received a correct ACK packet, it concludes that the transmission of the data packet has failed, and hence, it invokes its backoff procedure.

Step 6: Once the destination node i has successfully received the data packet, it immediately transmits the ACK packet over the ACK control channel.

In parallel with the above actions, each node, say i, must do the following:

- Decide on the value of $\alpha_{\text{eff}}^{(i)}$ using (7) and (8).
- Broadcast the computed $\alpha_{\text{eff}}^{(i)}$ as part of the "hello" packet, as explained in Section II-E.
- Cache the minimum of all the $\alpha_{\text{eff}}^{(.)}$ values it has heard.
- Compute $K^{(i)}$ using (4), and $P_{\text{noise}}^{(i)}$ using (5)

The IEEE 802.11 specifications [1] state that "the transmit power-on ramp for 10% to 90% of maximum power shall be no greater than 2 μsec," and that "the transmit power-down ramp for 90% to 10% maximum power shall be no greater than 2 μsec." Given that the RTS (or CTS) transmission duration is in the order of 100s of μsec and the data transmission duration is in the order of 1000s of μsec, the delay attributed to changing power levels (≤ 2 μsec) can be safely ignored. Note also that this delay is less than the turn-around period (period it takes a node to switch from a receiving mode to a transmitting mode); the later is approximately 5 μsec. Hence, after node j transmits an RTS packet at power P_{max}, the data packet can go out on a different power level with a very small delay that has negligible effect on the system efficiency.

IV. Related Work

Previous schemes for power control in MANETs have focused on either throughput enhancement or energy consumption. None of these schemes provide a comprehensive solution that enables a node to communicate via energy efficient links using different transmission ranges while still maintaining exclusive use of the channel (i.e., proper MAC functionality). In [27] the authors suggested a protocol that exploits global topological information provided by the routing protocol to reduce the nodes transmission powers such that the degree of each node is upper- and lower-bounded. In [32] a cone-based solution that guarantees network connectivity was proposed. The authors in [13] proposed the use of a synchronized global signaling channel to build a global network topology information where each node communicates only with its nearest N neighbors (N is a design parameter). In [29] the authors proposed a position-based distributed algorithm aided by a GPS system to allow each node to communicate only with its enclosure region. One common deficiency in the above protocols is that they rely solely on CSMA for accessing the wireless channel. It has been shown in [30], [21] that using CSMA alone for accessing the wireless channel significantly degrades network performance.

The COMPOW protocol [25] relies completely on routing-layer agents to converge to a *common* lowest power level for all network nodes. However, for constantly moving nodes, the scheme (like any routing-protocol-based scheme) incurs significant overhead, and convergence to a common power level may not be possible, leading to a situation like the one described in Figure 2. Moreover, in situations where network density varies widely (i.e., nodes are clustered), restricting all nodes to converge to a common power level can be conservative, and may achieve little gain.

Clustering as proposed in [22] is another interesting approach for power control. An elected cluster head (CH) performs the function of a base station in a cellular system. It uses closed-loop power control to adjust the transmission powers of nodes in the cluster. Communications between different clusters occur via GateWays, which are nodes that belong to more than one cluster. This approach simplifies the forwarding function for most nodes, but at the expense of reducing network utilization since all communications have to go through the CHs. This can also lead to the creation of bottlenecks. A joint clustering and power control protocol was proposed in [20], where each node runs several routing-layer agents that correspond to different power levels. These agents build their own routing tables by communicating with their peer routing agents at other nodes (i.e., the protocol is distributed with no CHs). Each node along the packet route determines the lowest-power routing table in which the destination is reachable. The routing overhead in this protocol grows in proportion to the number of routing agents, and can be significant even for simple mobility patterns (recall that for DSR, RREQ packets account for approximately 38% of the total received bytes [17]).

The problem of adjusting the transmission power of *broadcast* messages was addressed in [9]. The proposed approach relies on using the distance information between nodes to construct a restricted neighborhood graph (RNG). A node adjusts its transmission power to reach only those nodes that are in its RNG. This approach was improved in [10] by taking into account physical shadowing, realistic MAC protocols and the effect of collisions.

In [4], [33], a *single* channel was used to send the RTS-CTS control packets but at different power levels. This again results in the situation in Figure 2. In [19], [14], [26] the authors proposed that communicating nodes exchange their RTS and CTS packets at power P_{max}, but send their DATA/ACK packets at the minimum power P_{min} needed for reliable communication. The energy consumption in this approach is expectedly less. However, similar to the 802.11 scheme, control signals are used to silence neighboring nodes, preventing concurrent transmissions in the neighborhood of a receiver. In fact, as pointed out in [18], such schemes achieve less throughput than the IEEE 802.11 since they introduce a new problem; interference with the reception of the ACK message at the source node. The authors in [18] proposed a solution to that problem. The proposed protocol in that paper was shown to preserve energy without decreasing the network throughout below that of the IEEE 802.11. However, the protocol in [18] does not allow for any concurrent transmissions to take place over the reserved floor, where the reserved floor is the maximum transmission range (i.e., the control packets transmission range).

Of the several schemes for power control, the ones in [24], [34] are the most relevant to our scheme. Our work is in line with [24] in the sense that we use the signal strength of a received control message to bound the transmission power of neighboring nodes. However, our scheme differs from [24] in the following ways. First, the protocol in [24] relies on the network layer to find a power efficient next hop. In dense networks, where power control is supposed to achieve a higher channel reuse factor, the next hop will be in the maximum range region, and hence, little gain (if any) will be achieved in using power control. Even if we assume that a more intelligent power-aware routing protocol runs on top of the scheme in [24], this incurs the overhead of exchanging link-power information. In addition, routing packets will still have to be broadcasted at maximum power; something we avoid in PCDC. It is worth mentioning that the connectivity set that each node builds in PCDC is a result of sending the control packets (RTS-CTS) over a separate control channel at fixed power. Hence, this set cannot be built with protocols like the one described in [24]. Finally, while PCDC dynamically adjusts the interference margin of the receiver, depending on the nodes density and battery energy left, in [24] the authors use a fixed interference margin value that is determined offline.

A busy-tone based power control protocol was proposed in [34], where the sender transmits the data and the busy tone at minimum power. The receiver transmits its busy tone at maximum power. A neighbor estimates the channel gain from the busy tone and is allowed to transmit if its transmission is not expected to add more than a fixed "noise" value to the ongoing reception. However, in the suggested protocol, the receiver does not take into account the additional noise that future transmitters add to the ongoing reception. Consequently, the criterion for correct reception will simply not be met as soon as neighbors start their transmissions. In addition, a similar argument to the one mentioned above concerning next-hop selection also applies to the protocol in [34].

V. Protocol Evaluation

A. Simulation Setup

We now evaluate the performance of the PCDC protocol and contrast it with the IEEE 802.11 scheme. Our results are based on simulation experiments conducted using CSIM programs (CSIM is a C-based process-oriented discrete-event simulation package) [2]. In our simulations, we investigate both the network throughput as well as the energy consumption. We use two measures for the throughput: *channel utilization* (U) and *end-to-end throughput*. Channel utilization refers to the average number of successfully received packets per packet transmission time. Essentially, it is a measure of the one-hop goodput. Note that according to this definition, U can be greater than 1, since multiple transmissions can occur simultaneously.

For simplicity, data packets are assumed to be of a fixed size. Each node generates data packets according to a Poisson process with rate λ (same for all nodes). The capture model is similar to the one in [31]. We use a min-hop routing policy, but we ignore the routing overhead. For the 802.11 scheme, the next-hop candidates are nodes that are within the maximum power range of the sender. For PCDC, these candidates are nodes that are within the connectivity power range (based on $P_{\text{conn}}^{(.)}$). The random waypoint model is used for mobility, with a host speed that is uniformly distributed between 0 and 2 meters/sec. Note, however, that mobility has a little effect on our protocol, since an RTS-CTS exchange precedes every packet transmission. The transmission periods for the RTS, CTS, data, and ACK packets are all in tens of milliseconds, so no significant changes in topology take place within these periods. Nodes are assumed to have full energy in their batteries during the simulation time and thus, (8) was not implemented for simplicity. Other parameters used in the simulations are given in Table II. These paraments correspond to realistic hardware settings [3].

Data packet size	2 KB
802.11 data rate	2 Mbps
PCDC data rate	1.6 Mbps
Control channel rate	400 Kbps
SNR threshold	6 dB
Reception threshold	-94 dBm
Carrier-sense threshold	-108 dBm
Thermal+receiver noise	-169 dBm/Hz
P_{max}	20 dBm

TABLE II
Parameters used in the simulations.

B. Simulation Results

We consider two types of topologies: *random grid* and *clustered*. In the random grid topology, 49 mobile hosts are placed across a square area of length 3000 meters. The square is split into 49 smaller squares. The location of a mobile user is selected randomly within each of these squares. For each generated packet, the destination node is randomly selected.

Figure 9 depicts two instances of the network topology as constructed under PCDC and 802.11, respectively. As expected, the topology is much denser in the case of 802.11 because P_{max} is used to determine node connectivity. On average, we found that the node degree is 12.74 under the 802.11 scheme, compared to 4.81 under PCDC, which is a reduction of about 62%.

The performance for random grid topologies is demonstrated in Figure 10 as a function of the packet generation rate. It can be observed that under PCDC, U is about 2.5 times that of the 802.11 standard (on average). This significant increase in the utilization is achieved mainly because in PCDC, communicating nodes reserve smaller floors to achieve successful communications. This allows for a tighter packing of source-destination pairs within a network environment, thereby improving channel spectral reuse.

Part (b) of the figure depicts the end-to-end throughput, which is a more significant measure of effectiveness than the utilization. It is shown that PCDC achieves up to 45% increase in the end-to-end throughput. Furthermore, PCDC saturates at about twice the load at which the 802.11 scheme saturates.

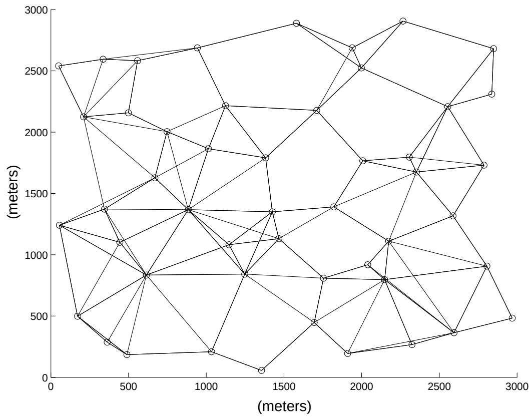

(a) PCDC.

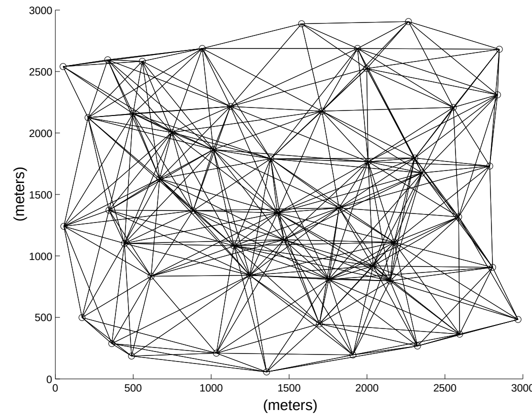

(b) 802.11.

Fig. 9. Instances of generated network topologies under PCDC and the IEEE 802.11 standard.

Note that PCDC has a more significant impact on utilization than on end-to-end throughput. The reason for this is that PCDC limits the set of possible next-hop nodes to ones that are within the $P_{\text{conn}}^{(.)}$ range, and thus, it forces shorter transmissions and longer routes. On the other hand, the 802.11 scheme allows for communication with any node that is within the maximum range, and hence, produces higher *progress* toward the destination per hop.

Part (a) of Figure 11 depicts the *transmission* energy consumption versus λ. This is the total energy used by the network to successfully *transmit* a packet end-to-end, normalized by the energy needed to send the data packet one hop at maximum power. It includes the energy lost in retransmitting data and control packets in case of collisions. For almost all cases, PCDC requires less than 23% of the energy required under the 802.11 scheme. The reason is that PCDC enforces shorter hops. In fact, our results indicate that, for the random grid topologies, the average number of hops generated by PCDC is about 1.76 that of the 802.11 scheme (using min-hop routing). Given that the signal power decreases as $1/d^4$, where d is the transmitter-receiver separation, PCDC should, on average, consume about $\nu = 21\%$ of the energy consumed by the 802.11 scheme. However, there are two other factors that contribute to ν. First, PCDC scales up the transmission power by $\alpha_{\text{min}}^{(i)}$ as explained in Section II-E. This increases ν. On the other hand, the node density is finite (49 nodes) in the studied topologies, and thus the transmitter-receiver separation distance is not continuous. Therefore, nodes using the 802.11 scheme do not achieve the maximum range by using P_{max} and energy is wasted. This decreases the value of ν.

Note that in both protocols, the required energy increases as the load increases, but for different reasons. For the 802.11 scheme, as λ increases the probability of collisions also increases, and hence more energy has to be spent on retransmissions. For PCDC, as λ increases the interference increases, and so, more power will be requested by receivers to achieve their SNR thresholds.

Part (b) of Figure 11 depicts the *reception* energy consumption versus λ. This is the average energy spent on receiving a data packet by both intended and unintended (or irrelevant, see Section II-B) receivers, normalized by the energy needed to receive a data packet by one node. As the transmission range of data packets is reduced, fewer irrelevant nodes receive this packet and thus, a significant amount of energy is preserved.

The packet delay performance of the PCDC and the 802.11 scheme is shown in Figure 12. For both protocols, the delay increases as the traffic rate is increased. Note, however, that as soon as λ exceeds a certain value, PCDC incurs less delay than the 802.11 scheme. At first this may seem counterintuitive. One would think that since packets travel longer routes under PCDC, it would take these packets longer times to reach their destinations than it would under the 802.11 scheme. However, a close look at the network reveals that as λ increases, the channel *contention period* also increases. In fact, as the traffic picks up, this period becomes much longer than the packet transmission period and dominates the end-to-end packet delay. The reason why PCDC incurs larger delays for very low load is that there is not enough packets to utilize the available space efficiently. If these very low traffic situations happen in a non-energy constrained networks, then it may be desirable to disable power control, possibly through the wireless network card interface. The disabling of power control in a node, say i, is equivalent to setting $P_{\text{conn}}^{(i)} = P_{\text{max}}$.

The authors in [23] argued that traffic locality is the key factor for determining the feasibility of large ad hoc networks. This motivates studying the performance of PCDC under *clustered* topologies. In such topologies, a node

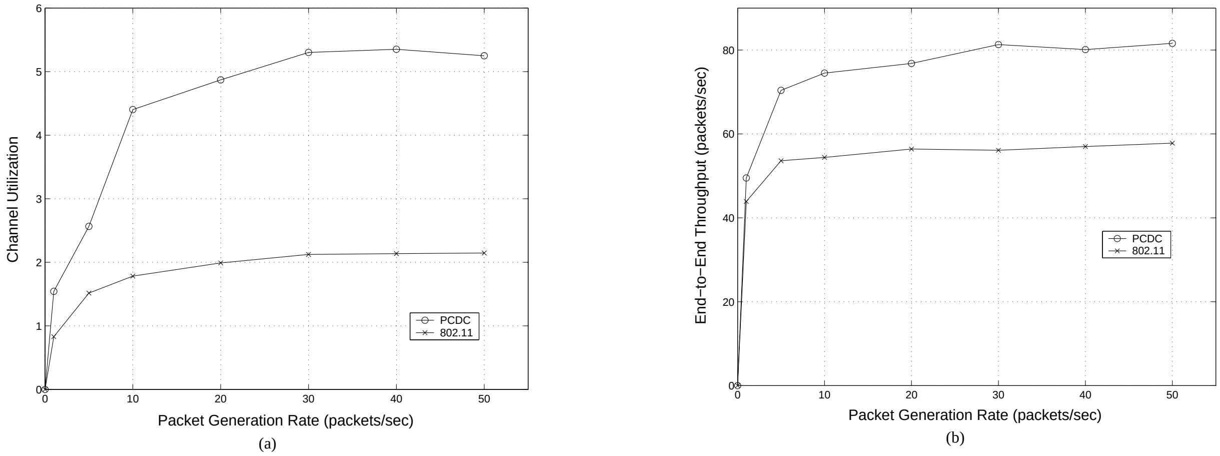

Fig. 10. Throughput performance of PCDC and 802.11 as a function of λ (random grid topologies).

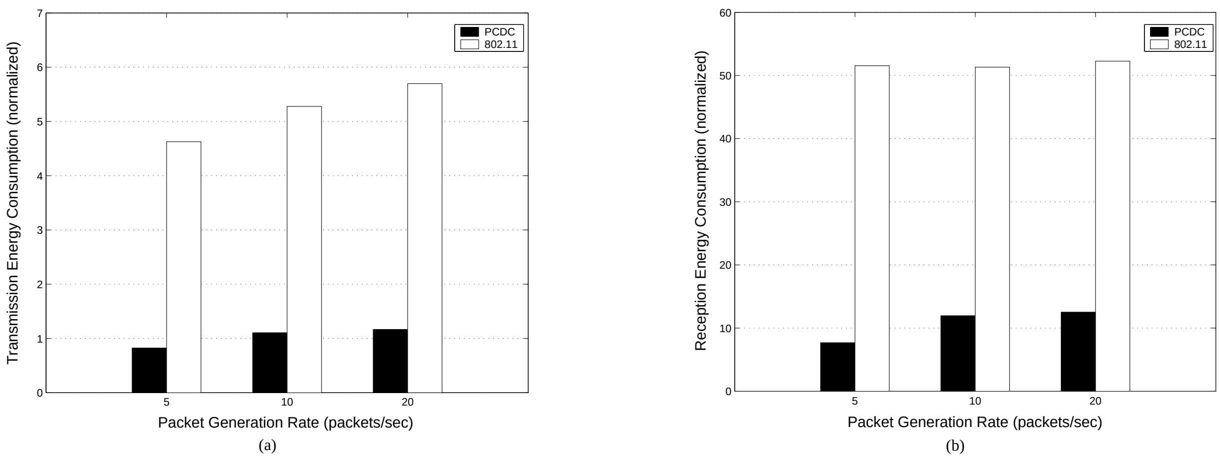

Fig. 11. Energy consumption in PCDC and 802.11 as a function of λ (random grid topologies).

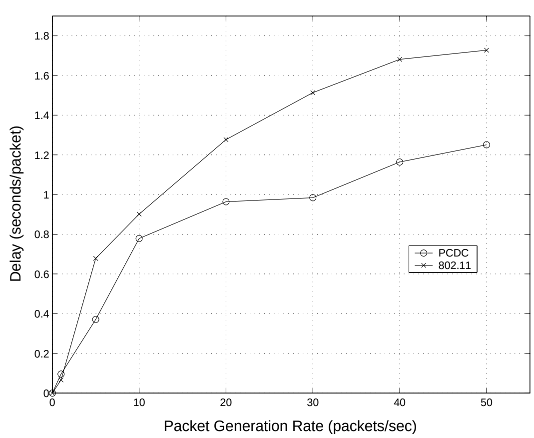

Fig. 12. Delay performance of PCDC and 802.11 as a function of λ (random grid topologies).

communicates mostly with nodes within its own cluster, and rarely with neighboring cluster nodes. These topologies are common in practice (e.g., a historical site where users of wireless devices move in groups). To generate a clustered topology, we consider an area of dimensions 1000×1000 (in meters). 24 nodes are split into 4 equal groups, each occupying a 100×100 square in one of the corners of the complete area. For a given source node, the destination is selected from the same cluster with probability $1 - p$ or from a different cluster with probability p. In each case, the selection from within the given cluster(s) is done randomly.

Figure 13 depicts the channel utilization and network throughput versus λ for $p = 0.25$. According to the 802.11 standard, only one transmission proceeds at a time since all nodes are within the carrier-sense range of each other. However, according to PCDC, up to four transmissions can proceed simultaneously, resulting in significant improvements in channel utilization and end-to-end throughput. Notice that for low traffic rates, and since the number of nodes is only 24 nodes, there is not enough packets to utilize the channel efficiently, which reduces the throughput of PCDC. This means that controlling the transmission range (and thus the next hop selection) can

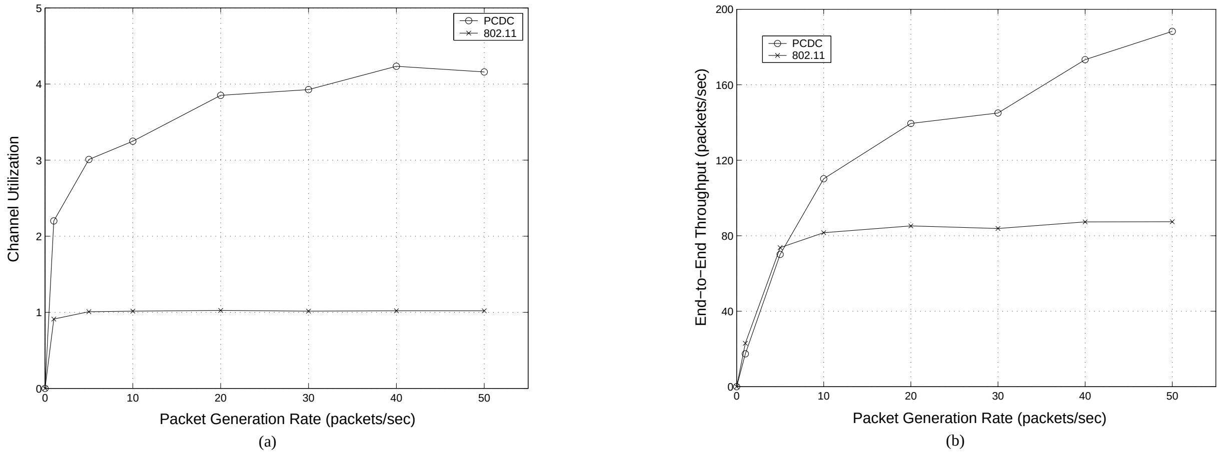

Fig. 13. Throughput performance of PCDC and 802.11 as a function of λ (clustered topologies with $p = 0.25$).

reduce the throughout but only at very low traffic rates.

Figure 14 also shows that PCDC consumes much less energy to successfully deliver a data packet than the 802.11 standard. The figures shows the total energy dissipated in transmissions and receptions, normalized by the energy needed to send the data packet one hop at maximum power. Note that in the case of clustered topologies, the energy consumption for the 802.11 does not vary with λ. The reason is that all the nodes are within each other's transmission range, which significantly reduces collisions. For PCDC, concurrent transmissions in different clusters add little interference to each other and so, the energy consumption does not vary significantly with λ.

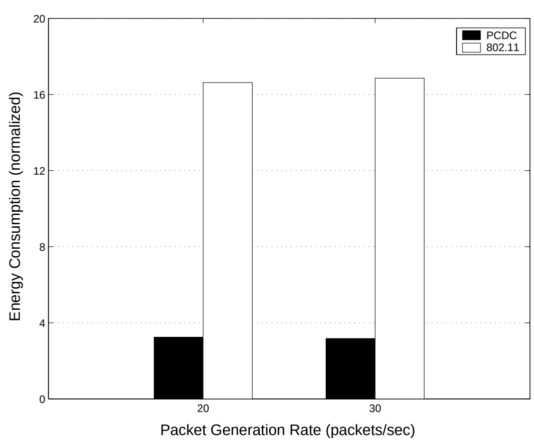

Fig. 14. Total energy consumption (transmission and reception) under PCDC and 802.11 as a function of λ (clustered topologies with $p = 0.25$).

Finally, the delay performance of the clustered topology is shown in Figure 15. Again, the figure shows that PCDC incurs much less delay for most values of λ, and performs worse only for very small values of λ.

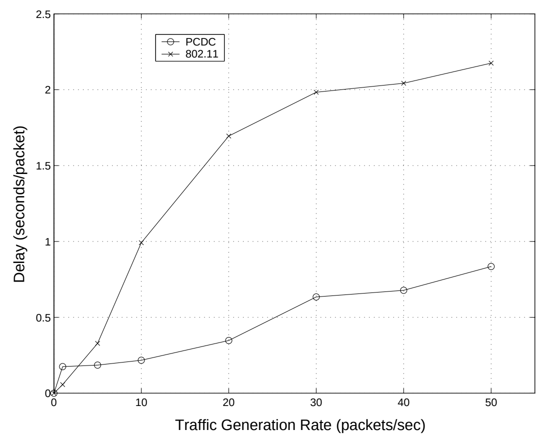

Fig. 15. Delay performance of PCDC and 802.11 as a function of λ (clustered topologies with $p = 0.25$).

VI. Conclusion

In this paper, we proposed a power controlled dual channel (PCDC) MAC protocol for wireless ad hoc networks. To produce power-efficient routes, PCDC allows the MAC layer to indirectly influence the routing decision at the network layer by controlling the power level of the broadcasted RREQ packets. PCDC uses the signal strength of the overheard control (RTS/CTS) signal to build a power-efficient network topology. By allowing for a receiver-specific, dynamically computed interference margin, PCDC enables simultaneous interference-limited transmissions to take place in the vicinity of a receiver.

We compared the performance of PCDC to that of the IEEE 802.11 standard. Our simulation results showed that PCDC can improve the channel utilization by up to 250% and the end-to-end throughput by over 45% (for random grid topologies). At the same time, PCDC provides for more than 76% reduction in the energy consumed to successfully deliver a packet from the source to the destination. It also reduces the end-to-end packet delay. To the best of our knowledge, PCDC is the first protocol to provide a comprehensive *and* efficient solution to the power control problem in MANETs. Our future work will focus on tuning the parameters of PCDC, studying a number of design issues, and investigating its performance under various mobility scenarios.

APPENDIX

PROOF OF THEOREM 1

Let x_1 and x_n be any pair of nodes in V. Since G is connected, then there exist a path in G from x_1 to x_n. Let this path be $\{x_1, x_2, \ldots, x_i, \ldots, x_{n-1}, x_n\}$. To prove that x_1 and x_n are connected in H, it is sufficient to prove that there exists a path in H between every pair of *successive* nodes, x_i and x_{i+1}, $i = 1, 2, \ldots, n-1$.

Consider any two successive nodes x_i and x_{i+1} on the G-path from x_1 to x_n. These two nodes are one-hop apart in G, that is, they are within each other's maximum transmission range. Therefore, according to our algorithm, one of the following two cases must hold:

- Case 1: $x_{i+1} \in CS_i$, and hence x_i and x_{i+1} are already connected in H.
- Case 2: $x_{i+1} \notin CS_i$. Then, there exists a node $u \in \mathrm{CS}_i$ for which the indirect communication $x_i \to u \to x_{i+1}$ requires less power than the direct one $x_i \to x_{i+1}$. Then, our problem now reduces to showing that there is a path in H between nodes u and x_{i+1}. Considering that the power required for transmission is proportional to d^4, then it must be the case that $d_{u,x_{i+1}} < d_{x_i,x_{i+1}}$, i.e., u and x_{i+1} are also within the transmission range of one another. By repeating the same process, one can prove that either $x_{i+1} \in \mathrm{CS}_u$ or that u can reach x_{i+1} through another (closer) node. This process continues until we reach a node that is in sufficiently close to x_{i+1} (i.e., in the CS of x_{i+1}), which completes the proof.

REFERENCES

[1] International Standard ISO/IEC 8802-11; ANSI/IEEE Std 802.11, 1999 Edn. Part 11: wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications.

[2] Mesquite Software Incorporation, http://www.mesquite.com.

[3] The Cisco Aironet 350 Series of wireless LAN, http://www.cisco.com/warp/public/cc/pd/witc/ao350ap/prodlit/a350c_ds.pdf.

[4] S. Agarwal, R. H. Katz, S. V. Krishnamurthy, and S. k. Dao. Distributed power control in ad-hoc wireless networks. In *IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, volume 2, pages 59–66, Oct. 2001.

[5] D. Ayyagari and A. Ephramides. Power control for link quality protection in cellular DS-CDMA networks with integrated (packet and circuit) services. In *Proceedings of the IEEE/ACM MOBICOM Conference*, pages 96–101, 1999.

[6] V. Bharghavan. Performance evaluation of algorithms for wireless medium access. In *Proceedings of the IEEE Performance and Dependability Symposium*, pages 86–95, Aug. 1998.

[7] V. Bharghavan, A. Demers, S. Shenker, and L. Zhang. MACAW: A media access protocol for wireless LANs. In *Proceedings of the ACM SIGCOMM Conference*, volume 24, pages 212–225, Oct. 1994.

[8] R. Caceres and L. Iftode. Improving the performance of reliable transport protocols in mobile computing environment. *IEEE Journal on Selected Areas in Communications*, 13:850–857, June 1995.

[9] J. Cartigny, D. Simplot, and I. Stojmenović. Localized minimum-energy broadcasting in ad-hoc networks. In *Proceedings of the IEEE INFOCOM Conference*, volume 3, pages 2210–2217, Apr. 2003.

[10] X. Chen, M. Faloutsos, and S. V. Krishnamurthy. Power adaptive broadcasting with local information in ad hoc networks. In *Proceedings of the IEEE ICNP Conference*, pages 168–178, Nov. 2003.

[11] J. Deng and Z. Haas. Dual busy tone multiple access (DBTMA): A new medium access control for packet radio networks. In *Proceedings of the IEEE ICUPC*, pages 973–977, Oct. 1998.

[12] S. Doshi, S. Bhandare, and T. X. Brown. An on-demand minimum energy routing protocol for a wireless ad hoc network. *ACM SIGMOBILE Mobile Computing and Communications Review*, 6(3):50–66, 2002.

[13] T. A. ElBatt, S. V. Krishnamurthy, D. Connors, and S. Dao. Power management for throughput enhancement in wireless ad-hoc networks. In *Proceedings of the IEEE ICC Conference*, volume 3, pages 1506–1513, 2000.

[14] J. Gomez, A. T. Campbell, M. Naghshineh, and C. Bisdikian. PARO: Supporting dynamic power controlled routing in wireless ad hoc networks. *ACM/Kluwer Wireless Networks*, 2003, to appear.

[15] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2):388–404, March 2000.

[16] T.-C. Hou and V. O. K. Li. Transmission range control in multiple packet radio networks. *IEEE Transactions on Communications*, 34(1):38–44, Jan. 1986.

[17] P. Johansson, T. Larsson, N. Hedman, B. Mielczarek, and M. Degermark. Scenario-based performance analysis of routing protocols for mobile ad-hoc networks. In *Proceedings of the IEEE/ACM MOBICOM Conference*, pages 195–206, 1999.

[18] E.-S. Jung and N. H. Vaidya. A power control mac protocol for ad hoc networks. In *Proceedings of the IEEE/ACM MOBICOM Conference*, pages 36–47, 2002.

[19] P. Karn. MACA - a new channel access method for packet radio. In *Proceedings of the 9th ARRL Computer Networking Conference*, pages 134–140, 1990.

[20] V. Kawadia and P. R. Kumar. Power control and clustering in ad hoc networks. In *Proceedings of the IEEE INFOCOM Conference*, 2003.

[21] S. Khurana, A. Kahol, and A. P. Jayasumana. Effect of hidden terminals on the performance of IEEE 802.11 MAC protocol. In *Proceedings of the IEEE LCN Conference*, pages 12–20, 1998.

[22] T. J. Kwon and M. Gerla. Clustering with power control. In *Proceedings of the IEEE MILCOM Conference*, volume 2, pages 1424–1428, 1999.

[23] J. Li, C. Blake, D. S. Couto, H. I. Lee, and R. Morris. Capacity of ad hoc wireless networks. In *Proceedings of the IEEE/ACM MOBICOM Conference*, pages 61–69, 2001.

[24] J. Monks, V. Bharghavan, and W.-M. Hwu. A power controlled multiple access protocol for wireless packet networks. In *Proceedings of the IEEE INFOCOM Conference*, volume 1, pages 219–228, 2001.

[25] S. Narayanaswamy, V. Kawadia, R. S. Sreenivas, and P. R. Kumar. Power control in ad-hoc networks: Theory, architecture, algorithm and implementation of the COMPOW protocol. In *Proceedings of the European Wireless Conference*, pages 156–162, Feb. 2002.

[26] M. B. Pursley, H. B. Russell, and J. S. Wysocarski. Energy-efficient transmission and routing protocols for wireless multiple-hop networks and spread spectrum radios. In *Proceedings of the EUROCOMM Conference*, pages 1–5, 2000.

[27] R. Ramanathan and R. Rosales-Hain. Topology control of multihop wireless networks using transmit power adjustment. In *Proceedings of the IEEE INFOCOM Conference*, volume 2, pages 404–413, 2000.

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

[29] V. Rodoplu and T. Meng. Minimum energy mobile wireless networks. *IEEE Journal on Selected Areas in Communications*, 17:1333–1344, Aug. 1999.

[30] F. A. Tobagi and L. Kleinrock. Packet switching in radio channels: Part ii–the hidden terminal problem in carrier sense multiple-access and the busy-tone solution. *IEEE Transactions on Communications*, 23(12):1417–1433, Dec 1975.

[31] C. Ware, J. Chicharo, and T. Wysocki. Simulation of capture behaviour in IEEE 802.11 radio modems. In *Proceedings of the IEEE Vehicular Tech. Conference*, volume 3, pages 1393–1397, Fall 2001.

[32] R. Wattenhofer, L. Li, P. Bahl, and Y.-M. Wang. Distributed topology control for power efficient operation in multihop wireless ad hoc networks. In *Proceedings of the IEEE INFOCOM Conference*, volume 3, pages 1388–1397, 2001.

[33] J. Whitehead. Distributed packet dynamic resourse allocation (DRA) for wireless networks. In *Proceedings of the IEEE Vehicular Tech. Conference*, volume 1, pages 111–115, 1996.

[34] S.-L. Wu, Y.-C. Tseng, and J.-P. Sheu. Intelligent medium access for mobile ad hoc networks with busy tones and power control. *IEEE Journal on Selected Areas in Communications*, 18(9):1647–1657, 2000.