Power Controlled Dual Channel (PCDC) Medium Access Protocol for Wireless 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. Directional and channel-gain information obtained mainly from overheard RTS and CTS packets is used to dynamically construct the network topology. 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, IEEE 802.11.**

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 [7]. 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 Control 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 other 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 terminal problem, it negatively impacts the channel utilization by not allowing concurrent transmissions 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 $A \to B$ and $C \to D$ 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.

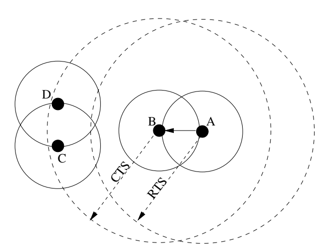

Fig. 1. Inefficiency of the standard RTS-CTS approach. Nodes A and B are allowed to communicate, but nodes C and D are not.

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.

[1] In addition to the DCF mode, the 802.11 standard also supports a Point Coordination Function (PCF) mode, which is essentially a polling scheme that is intended for delay-sensitive traffic.

While the idea of power control is simple, achieving it in a distributed manner is quite challenging. Consider, for example, the situation 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.

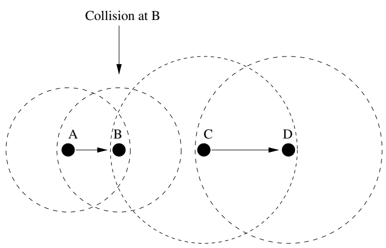

Fig. 2. Challenge of implementing power control in a distributed fashion. Node C is unaware of the communication $A \rightarrow B$, and hence it starts transmitting to node D at a power that destroys B' s reception.

From the above example, one can make the following observation: if nodes send their control (RTS-CTS) packets at a fixed maximum power level (P_{max}), but send their data packets at an adjustable (lower) power level, then the collision in the previous example could 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.

Power control for MANETs has been extensively studied (see Section III 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., [14], [24]), or by using it as a means of controlling the connectivity and topological properties of the network (e.g., [22], [16], [17], [19]). 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. In Section III we review related work in the area of power control for MANETs. The simulation results are presented and discussed in Section IV. Finally, our main conclusions are drawn in Section V.

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.

Radio channels are typically modeled using large- and small-scale propagation models [18]. Large-scale models are used to predict the mean signal strength for an arbitrary transmitter-receiver separation. Such models have 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 models characterize 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 models 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-random 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 a 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 hear the transmission, 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. Obviously, this 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 [8] 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 [7] 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. This allows for simultaneous transmissions in the neighborhood.

In addition to improving network throughput, reducing the transmission range plays a significant role in reducing the energy consumption. In [19] 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 [22]. Another advantage of power control that has not received much attention in the literature is related to reducing the power consumption at *irrelevant receivers* (those who are not addressed by the transmission). Significant power (more than 50% of the transmission power) is consumed in receiving a packet. Since reducing the transmission range results in a smaller number of nodes overhearing the transmission, less power will be consumed by those irrelevant 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 the 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, the second 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-D we develop a distributed strategy that dynamically adjusts the interference margin to maximize the number of simultaneous transmissions.

Finally, PCDC incorporates the above optimizations into a modified RTS-CTS reservation mechanism. According to this mechanism, a receiving node exploits knowledge of the power levels of the overheard RTS and CTS messages to determine the power that should be used by an intended transmitter. The details of this mechanism are explained in Section II-E.

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), the minimum set of nodes that guarantees connectivity of the node to the network. From this CS, the node can then decide on the set of next-hop nodes, as will be explained shortly. 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, assuring proper MAC functionality, 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 and angle of arrival (AOA) 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. In addition, techniques for AOA estimation without a positioning system (i.e., GPS) are available (see [11] for details). 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.

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 gain and AOA match the already stored ones, 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 or AOA *do not* match the already stored ones, then node i compares P_{ij} (the power required to communicate directly with node j) with $P_{iu} + P_{uj}$, where

$u \in \mathrm{CS}_i$. If $P_{ij} < P_{iu} + P_{uj}$ for every node $u \in \mathrm{CS}_i$, then node j is added to CS_i; otherwise, it is not. Let $P^{(i)}_{conn}$ denote the minimum power required for node i to reach the *farthest* node in CS_i. If node j is added to CS_i *and* $P_{ij} < P^{(i)}_{conn}$, then all other elements of CS_i must be re-examined. The reason is that a two-hop path between node i and a node $u \in \mathrm{CS}_i$ via 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 $P_{ij} \geq P^{(i)}_{conn}$, 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 $P^{(i)}_{conn}$ upon the receipt of an RTS/CTS packet from node j.

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

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

The computation of P_{ij} is easy since node i estimates the channel gain G_{ij} from j's signal. However, the computation of the power required for indirect communication is less obvious. To compute this power, we make use of the two-ray propagation model [18]. This model, which adequately characterizes the large-scale variations of the received signal power, implies a k/d^4 transmit power roll-off, where d is the transmitter-receiver separation and k is a constant that depends on the system parameters (e.g., antenna gain and height). Now, consider the situation in Figure 4. Suppose that $u \in \mathrm{CS}_i$, so node i already knows the channel gain G_{iu} and the direction of node u. Node i then uses the two-ray model to estimate the distances d_{iu} and d_{ij} using G_{iu} and G_{ij}, respectively. In addition, node i geometrically determines the distance d_{uj}:

$$d_{uj} = \sqrt{d^2_{iu} + d^2_{ij} - 2\, d_{iu}\, d_{ij} \cos w} \qquad (1)$$

Node i then uses d_{iu}, d_{uj}, and d_{ij} to find the power levels P_{ij} and $P_{iu} + P_{uj}$. Note that the value of the constant k is not needed for the power comparison, since k cancels from both sides of the inequality.

In deciding whether to add node j to CS_i or not, we only considered two-hop indirect communications. 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

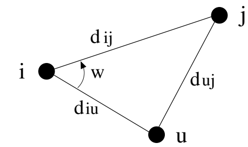

Fig. 4. Computation of the power required for indirect communication between nodes i and j via node u.

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 $x \in \mathrm{CS}_i$. 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 $i \to u \to v$ must require less power than $i \to v$, and hence, node v cannot be in CS_i. This contradicts the assumption that $v \in \mathrm{CS}_i$.

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 maximum power (P_{max}) to reach other nodes. Let $H = (V, E')$ be the undirected graph constructed based on our CS approach. If G is connected, then H is also connected.

Proof: See [15].

One nice feature of the algorithm is its symmetrical property: if $i \in \mathrm{CS}_j$ then $j \in \mathrm{CS}_i$, and vice versa (see [15] for details).

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 the control channel is mostly idle, and an auxiliary scheme is needed to ensure accurate computation of the connectivity set. In our protocol, if a node does not send any RTS or CTS packets for Δ seconds, this node shall broadcast a "hello" packet over the control channel at power P_{max}. The parameter Δ is a random variable that is uniformly distributed in the interval $[T/2, T]$, where 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. Randomization is needed to avoid collisions between synchronized "hello" transmissions. 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 used in the standard IEEE 802.11 CTS packet to indicate the receiver address is now used to indicate the transmitter address. Second, the duration field of the standard CTS packet is used here for a different purpose, which will be explained in Section II-D. Figure 5 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 ($\frac{14}{T}$

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

Fig. 5. Format of the “hello” packet in PCDC.

bytes per node per second). This is in contrast to the scheme in [22], where periodic or on-demand reconfiguration of the network topology is always needed if nodes are moving (the authors simulated only a static network). This affects network resources and increases packet delays, especially at peak load times.

Now that node i has computed the connectivity power $P^{(i)}_{conn}$, *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. Clearly, this reduces complexity and overhead. 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 $P^{(i)}_{conn}$. Take, for example, the network in Figure 6 (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 [9] 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,

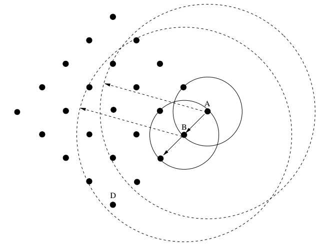

Fig. 6. Example showing the inefficiency of broadcasting RREQ packets at the maximum power.

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 is almost fixed, making it easier to design the contention window for RREQ packets. Therefore, PCDC results in lower overhead, less contention, and less consumed power in the process of finding the destination.

D. Interference Margin

In communicating with its connectivity set, a node has to know how much interference it can allow to account for *potential* transmissions in its neighborhood. Here, we develop a strategy that dynamically adjusts the interference margin to maximize network throughput. In [3] 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 received signal-to-interference ratio (SIR) at a receiving node i is given by:

$$\mathrm{SIR}^{(i)} = \frac{P^{(i)}_m}{\sum\limits_{n=1, n\neq m}^{N} P^{(i)}_n + \eta + \beta^{(i)}} \tag{2}$$

where $P^{(i)}_m$ is the “desired” power at the receiver i from the intended transmitter m, $P^{(i)}_n$ is the received power from an interfering (unintended) transmitter n, N is the number of active transmitters in the neighborhood of node i, η is the thermal noise, and $\beta^{(i)}$ is the interference margin of node i. It was proven in [3] that to increase channel capacity, $\beta^{(i)}$ must 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.

Applying the same algorithm in ad hoc networks is not so straightforward due to the absence of a centralized control.

Moreover, in an ad hoc network, 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 whom node i may interfere with.

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

$$\alpha^{(i)} = \frac{P_{max}}{P_{conn}^{(i)}} \tag{3}$$

This value represents the maximum scaling constant (and hence, the maximum interference margin) that node i can be asked to use without losing communication with its connectivity set. 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 energy. To account for these two conflicting goals, we use the ratio of the remaining energy ($E_{remain}^{(i)}$) to the full energy ($E_{full}^{(i)}$) of the battery to scale down the value of $\alpha^{(i)}$ as follows:

$$\alpha_{eff}^{(i)} = \max\left\{ 1,\ \alpha^{(i)} \frac{\left\lfloor \frac{4 \times E_{remain}^{(i)}}{E_{full}^{(i)}} \right\rfloor}{4} \right\} \tag{4}$$

The above relationship is chosen so that a node uses zero interference margin when its battery energy drops below 25% (note that $\alpha_{eff}^{(i)}$ must be greater than or equal to one, or otherwise coherent reception at node i is not possible). Node i then broadcasts the value of $\alpha_{eff}^{(i)}$ in the reserved field of the "hello" packets mentioned in Section II-C. It then chooses the minimum of all the $\alpha_{eff}^{(.)}$ values it has heard. Let $\alpha_{\min}^{(i)}$ be this minimum. The intuition is that if the scaling factor is made larger than $\alpha_{\min}^{(i)}$, 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.

E. Channel Access Mechanism

In our scheme, RTS and CTS packets are used to provide three functions. First, these packets allow nodes (transmitters and receivers) to determine the channel gain. Second, a receiver i uses the CTS packet to notify its neighbors of the additional noise power (denoted by $P_{noise}^{(i)}$) that each of the neighbors can add to node i without impacting its 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 a RTS packet over the control channel at P_{max} and includes in the RTS packet the maximum *allowable* power level ($MAP^{(j)}$) (among the M possible power levels provided by the hardware) that node j can 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 maximum allowable power level at the transmitter. Note that the additional field is enough to hold up to 256 levels.

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

$$P_{min}^{(ij)} = \frac{\text{SIR}_{TH} \times \delta^{(i)}}{G_{ij}} \tag{5}$$

where SIR_{TH} is the minimum SIR ratio that is needed for correct decoding (we assume SIR_{TH} is the same for all nodes, i.e., all nodes require the same QoS) and $\delta^{(i)}$ is the total undesired power, which includes the thermal noise plus the power received from all already ongoing transmissions that are not addressed to node i. This $P_{min}^{(ij)}$, however, does not allow for any interference margin at node i, so 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 neighborhood, node i requests that node j uses a transmission power that is larger than $P_{min}^{(ij)}$. Specifically, the power that must be used by node j is given by:

$$P_{requested}^{(ij)} = \frac{\text{SIR}_{TH}(\delta^{(i)} + K^{(i)} P_{noise}^{(i)})}{G_{ij}} \tag{6}$$

where $K^{(i)} P_{noise}^{(i)}$ is the total interference margin that node i can tolerate from unintended transmitters (the computation of $K^{(i)}$ will be discussed shortly). When responding to j's RTS, node i indicates in its CTS the power level (among the M possible levels) that is just above $P_{requested}^{(ij)}$. Node i then inserts a share of this value, namely $P_{noise}^{(i)}$, in the CTS packet and sends this packet back to sender j at P_{max} over the control channel. The rational behind inserting a share is to prevent one neighbor from consuming the entire interference margin. In other words, we think of the interference margin as a network resource that should be shared among various nodes. The format of the CTS packet is shown in Figure 7.

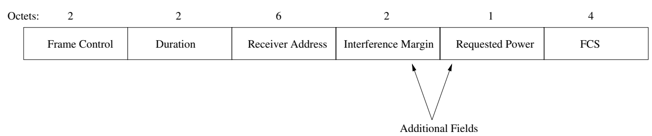

Fig. 7. Format of the CTS packet in the proposed protocol.

A potentially interfering node, say C, that hears the CTS message uses the signal strength of the received CTS to compute the channel gain between itself and node i. The channel gain along with the broadcasted interference margin $P_{noise}^{(i)}$ are used to compute C's maximum allowable power. This is the power that node C can use in its future transmissions that will not add more than $P_{noise}^{(i)}$ to the received noise at node i.

The parameter $K^{(i)}$ essentially represents the number of additional concurrent (unintended) transmissions in the neighborhood of a receiver. In here, we take $K^{(i)}$ as:

$$K^{(i)} = \frac{P_{max}}{P_{conn}^{(i)}} - 1 \quad (7)$$

The intuition behind this choice is as follows: suppose that the density of nodes in the region is low, and hence nodes can only communicate with each other using P_{max} with a zero interference margin. This means that $K^{(i)}$ has to be zero, which is the case in (7) since for a sparsely connected network $P_{conn}^{(i)} = P_{max}$.

We can express $P_{noise}^{(i)}$ in terms of $\alpha_{\text{min}}^{(i)}$, as follows:

$$\alpha_{\text{min}}^{(i)} \times \frac{SNR_{TH} \times \delta^{(i)}}{G_{ij}} = \frac{SNR_{TH}(\delta^{(i)} + K^{(i)} P_{noise}^{(i)})}{G_{ij}} \quad (8)$$

$$\Rightarrow P_{noise}^{(i)} = \delta^{(i)} \times \frac{\alpha_{\text{min}}^{(i)} - 1}{K^{(i)}} \quad (9)$$

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.

The protection of ACK packets was addressed in previous MAC protocols, 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 relatively small compared to the data-packet duration ($\approx 1\%$). 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 [4] can also be used.

G. Protocol Recovery

In [5] 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 a 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. In PCDC, we avoid the above scenario 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 $MAP^{(\cdot)}$ value that if used could cause an unacceptable 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

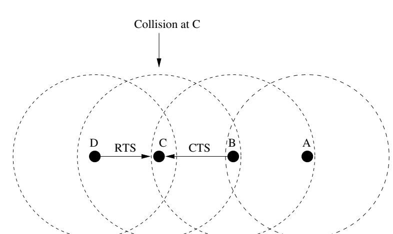

Fig. 8. Example of a collision between control packets.

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, ACK 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 [5]. 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. 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 [17] 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 [22] a cone-based solution that guarantees network connectivity was proposed. The authors in [6] 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 [19] 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 [20], [10] that using CSMA alone for accessing the wireless channel significantly degrades network performance.

The COMPOW protocol [16] 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 widely varies (i.e., nodes are clustered), restricting all nodes to converge to a common power level is very conservative, and achieves little gain, if any.

Clustering as proposed in [12] is another interesting approach for power control. It simplifies the forwarding function for most nodes, but at the expense reducing network utilization since all communications have to go through the cluster heads. This can also lead to the creation of bottlenecks. In [2], [23], 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.

Of the several schemes for power control, the ones in [14], [24] are the most relevant to our scheme. Our work is in line with [14] 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 [14] in the following ways. First, the protocol in [14] relies on the network layer to find a power efficient next hop. In dense networks, where power control is assumed 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 [14], 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 [14]. Finally, while PCDC dynamically adjusts the interference margin of the receiver, depending on the nodes density and battery energy left, in [14] the authors use a fixed interference margin value that is determined offline.

A busy-tone based power control protocol was proposed in [24], 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 [24].

IV. 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). 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 have 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 [21]. 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 $P_{conn}^{(\cdot)}$). The random waypoint model was used for mobility. Other parameters used in the simulations are given in Table IV-A.

Data packet size	2 KB
Data link speed	2 Mbps
Control link speed	200 Kbps
SIR threshold	6 dB
Reception threshold	-64 dBm
Carrier-sense threshold	-78 dBm
Maximum power	30 dBm
Number of power levels	12

TABLE I

PARAMETERS USED IN THE SIMULATIONS.

B. Simulation Results

We consider two types of topologies: *uniform* and *clustered*. In the uniform topology, 50 mobile hosts are placed across a square area of length 1000 meters. The square is split into 50 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 to be any node in the network.

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 the maximum power (P_{max}) is used to determine node connectivity. On average, we found that the node degree is 11.63 under the 802.11 scheme, compared to 4.12 under PCDC, which is a reduction of about 65%.

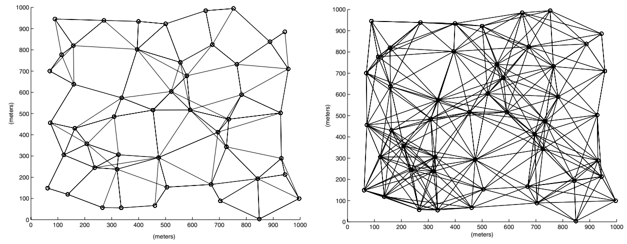

(a) PCDC. (b) 802.11.

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

The performance for uniform topologies is demonstrated in Figure 10. In this figure, we vary the packet generation rate (λ) of the Poisson process. It can be observed that under PCDC, U is about 2.4 times that of the 802.11 standard (on average). This significant increase in the utilization is due, in part, to the longer paths taken by packets under PCDC, and also to the increase in the number of simultaneous transmissions. 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 62% increase in the end-to-end throughput. Furthermore, PCDC saturates at about twice the load at which the 802.11 saturates.

Part (c) of Figure 10 depicts the energy consumption versus λ. Energy consumption is the total energy used 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 50% of the energy required under the 802.11 scheme. Note that in both protocols, the required energy increases as the load increases, but for different reasons. For the 802.11 standard, 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.

The authors in [13] argued that traffic locality is the key factor in determining the feasibility of large ad hoc networks. To investigate the effect of this factor on PCDC, we vary the end-to-end sender-destination separation distance and measure the end-to-end throughput. Indeed, as Figure 11 shows, the locality of the traffic can highly impact the network throughput. In fact, PCDC benefits from traffic locality, simply because it favors short ranges. Moreover, as the sender-destination separation gets larger, the overhead of the control packets associated with each extra hop the packet has to travel becomes considerable. This affects PCDC more significantly since control packets are transmitted at a lower speed than data packets. Therefore, as the destination gets farther, the throughput enhancement of PCDC gets less.

The above result motivates studying the performance of PCDC under *clustered* topologies. In such topologies, a node communicates mostly with nodes within its own cluster, and

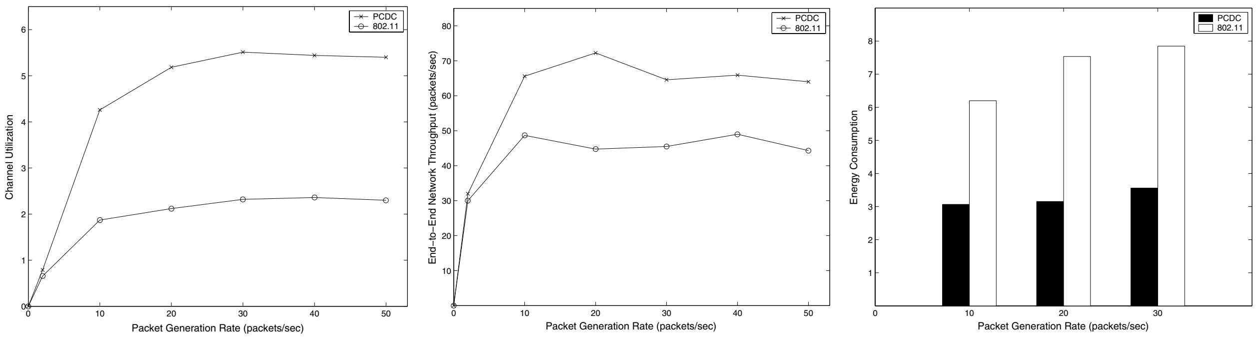

Fig. 10. Performance of PCDC and 802.11 as a function of λ (uniform topologies).

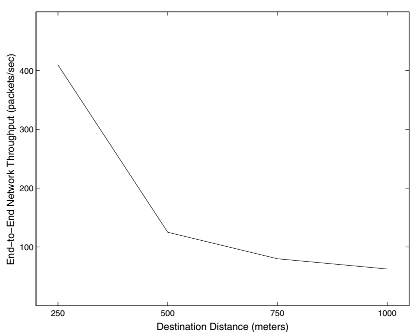

Fig. 11. End-to-end throughput versus the sender-destination separation distance for the PCDC protocol.

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 500×500 (in meters). Nodes are split into 4 groups, each occupying a 100×100 square in one of the corners of the complete area. As an extreme scenario, we let the source and destination nodes be randomly selected from the same group. Figure 12 depicts the performance versus λ. 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, four transmissions can proceed simultaneously, resulting in significant improvements in the channel utilization and the end-to-end throughput. Part (c) of the figure also shows that PCDC consumes much less energy to successfully deliver a data packet than the 802.11 standard. 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, the energy increases with the load, due to the increase in the interference from other concurrent transmissions.

V. 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 and the direction of arrival 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 240% and the end-to-end throughput by over 60%. At the same time, PCDC provides for more than 50% reduction in the energy consumed to successfully deliver a packet from the source to the destination. 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 the performance in more realistic topologies.

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] 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.

[3] 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.

[4] 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.

[5] 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.

[6] 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.

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

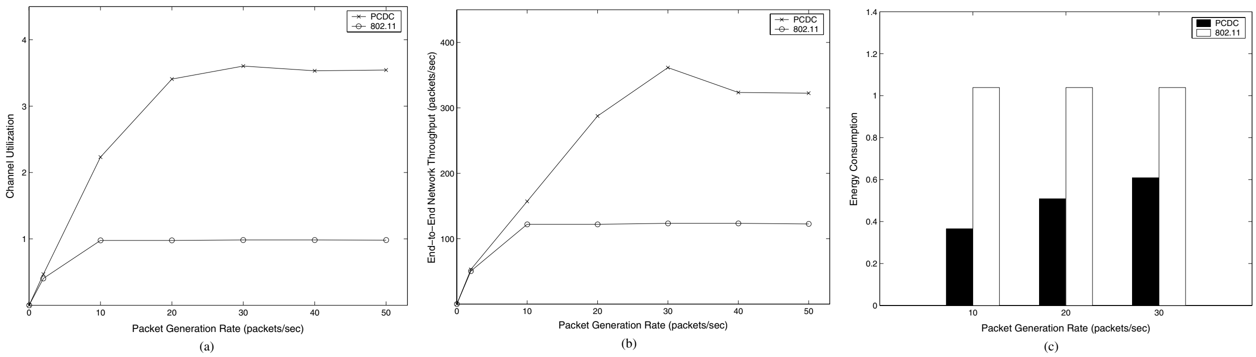

Fig. 12. Performance of PCDC and 802.11 as a function of λ (clustered topologies).

[8] 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.

[9] 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.

[10] 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.

[11] K. Krizman, T. E. Biedka, and T. S. Rappaport. Wireless position location: Fundamentals, implementation strategies, and source of error. In *Proceedings of the IEEE Vehicular Tech. Conference*, volume 2, pages 919–923, 1997.

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

[13] 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.

[14] 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.

[15] A. Muqattash and M. Krunz. Power controlled dual channel (PCDC) medium access protocol for wireless ad hoc networks (extended version). Technical report CENG-TR-245, University of Arizona, Dept. of Electrical & Computer Engineering, July 2002. http://www.ece.arizona.edu/∼krunz/tech_rep.html.

[16] 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.

[17] 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.

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

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

[20] 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.

[21] 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.

[22] 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.

[23] 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.

[24] 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.