

A Single-Channel Solution for Transmission Power Control in Wireless Ad Hoc Networks

Alaa Muqattash and Marwan Krunz

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

{alaa,krunz}@ece.arizona.edu

ABSTRACT

Transmission power control (TPC) has a great potential to increase the throughput of a mobile ad hoc network (MANET). Existing TPC schemes achieve this goal by using additional hardware (e.g., multiple transceivers), by compromising the collision avoidance property of the channel access scheme, or by imposing impractical requirements on the operation of the MAC protocol. In this paper, we present a novel power control MAC protocol, known as POWMAC, for MANETs that enjoys the same simple single-channel, single-transceiver design of the IEEE 802.11 ad hoc MAC protocol, but that achieves a significant throughput improvement over the 802.11 scheme. Collision avoidance is integrated into the design of POWMAC. Instead of alternating between the transmission of control (RTS/CTS) and data packets, as done in the 802.11 scheme, POWMAC uses an access window (AW) to allow for a series of RTS/CTS exchanges to take place before multiple, concurrent data packet transmissions can commence. The length of the AW is dynamically adjusted (based on local traffic load information) to allow for concurrent *interference-limited* transmissions to take place in the same vicinity of a receiving node. Collision avoidance information is inserted into the CTS packet and is used to *bound* the transmission powers of potential interferers, rather than to silence such nodes. Simulation results for “random-grid” and “clustered” topologies are used to demonstrate the significant throughput and energy gains that can be obtained under the POWMAC protocol.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless Communication; C.2.2 [Network Protocols]

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiHoc'04, May 24–26, 2004, Roppongi, Japan.

Copyright 2004 ACM 1-58113-849-0/04/0005 ...\$5.00.

General Terms

Algorithm, Design, Performance

Keywords

Power control, IEEE 802.11, ad hoc networks, throughput enhancement, load control, interference margin, multi-access interference.

1. INTRODUCTION

1.1 Motivation

Mobile ad hoc networks have recently received significant interest. This interest is attributed to the distributed, infrastructureless and robust nature of MANETs, which allows for instant deployment and rerouting of traffic around failed or congested nodes. Given the scarcity of the wireless spectrum and the need to share this spectrum efficiently among contending nodes, considerable research efforts are being dedicated to improving the throughput of a MANET.

So far, the ad hoc mode of the IEEE 802.11 standard [1] has been used as the *de facto* MAC protocol for MANETs. This protocol is based on CSMA/CA (carrier sense multiple access with collision avoidance), with extensions to allow for the exchange of RTS/CTS (request-to-send/clear-to-send) handshake packets between the transmitter and the receiver. The RTS/CTS exchange is used to reserve a *transmission floor* for the subsequent data packet transmission. Nodes transmit their control and data packets at a fixed (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 complete.

Despite its appealing simplicity, the 802.11 MAC approach can be overly conservative [20, 21], leading to an unnecessary reduction in network throughput. To illustrate, consider the situation in Figure 1, where node A uses its maximum transmission power (TP) to send packets to node B (we assume omnidirectional antennas, so a node's reserved floor is represented by a circle in the 2D space). Nodes D and E refrain from transmitting to avoid interfering with the ongoing transmission $A \rightarrow B$. However, it is easy to show that the three transmissions $A \rightarrow B$, $D \rightarrow C$, and $E \rightarrow F$ can overlap in time if nodes are able to select their transmission powers appropriately, consequently increasing network throughput




Figure 1: Inefficiency of the standard RTS/CTS approach. Nodes *A* and *B* are allowed to communicate, but nodes *D* and *E* are not allowed. Dashed circles indicate the maximum transmission ranges, while dotted ones indicate the minimum transmission ranges needed for coherent reception at the respective receivers.

and possibly reducing the overall energy consumption. The scheme proposed in this paper is intended to allow for such transmissions to take place.

The above example motivates the need for transmission power control (TPC) to improve network throughput by means of increasing the channel spatial reuse. Theoretical studies [13] and simulation results [20, 21] have demonstrated that TPC can provide significant gains in capacity and energy consumption, not to mention its benefits in providing admission control and in quality of service (QoS) provisioning [7]. Many TPC schemes for MANETs have been proposed in the literature. However, as explained in Section 2, these schemes suffer from one or more of the following deficiencies: (1) the TPC approach may yield energy reduction but not throughput gain, (2) lack of collision avoidance in the MAC design (resulting in the well-known hidden terminal problem), (3) extra hardware (e.g., multiple transceivers) is needed, and (4) the assumptions made in the MAC design are unrealistic. Accordingly, we introduce a new TPC scheme for MANETs that does not suffer from these deficiencies. In this scheme, a single-channel, single-transceiver power-controlled MAC protocol is used to improve throughput of the IEEE 802.11 scheme while preserving its collision avoidance properties.

The rest of the paper is organized as follows. In Section 2, we present and analyze related work in the area of TPC for MANETs, and show their limitations. The proposed protocol is presented in Section 3, followed by simulation results and discussion in Section 4. Finally, our main conclusions are drawn in Section 5.

2. RELATED WORK

TPC schemes for MANETs can be generally classified into two classes. In the first class (e.g., [27, 29, 32, 11]), TPC is used to control the network topology, indirectly impacting the set of next-hop neighbors of a node and the subsequent routing decisions taken by that node. The *same* TP is used by a node to transmit its packets to any of its neighbors (the TP can vary from one transmitting node to another). The TP is updated following a mobility-related topological change, which often occurs at the time scale of hundreds

of milliseconds to seconds (in contrast, packet transmission times occur at a time scale of few milliseconds). The main design issue here is how to determine the minimum amount of TP for a given node such that some topological properties (e.g., connectivity, bounded node degree, etc.) are guaranteed. One common limitation of this class of protocols is its sole reliance on CSMA for accessing/reserving the shared wireless channel. It has been shown [30] that using CSMA alone for accessing the channel can significantly degrade network performance (throughput, delay, and power consumption) because of the hidden node problem. Unfortunately, this problem cannot be overcome by simply incorporating a RTS/CTS-like channel reservation strategy in the MAC design (see [21] for details).

In the second class of TPC schemes, TPC is applied on a per-packet basis, with the TP being dependent on *both* the transmitting and receiving nodes. The TP in this case is not directly tied to the routing layer or the topological properties of the network (although some schemes in this class indirectly influence the decisions taken by the routing layer). For a given next hop that is provided by the routing layer, the main question here is what TP to be used for sending the current data packet to that next hop. This class of TPC schemes can be further divided into two subclasses: energy- and throughput-oriented schemes. The former subclass (e.g., [12, 15, 16, 25]) aims primarily at reducing energy consumption, with network throughput being a secondary factor. Nodes exchange their RTS and CTS packets at a maximum power (P_{\max}), but send their data and ACK packets at the minimum power needed for reliable communication (P_{\min}). The value of P_{\min} is determined based on the required QoS (i.e., signal-to-interference-and-noise ratio), the interference level at the receiver, and the channel gain between the transmitter and the receiver. In [15] the authors enhanced the performance of this approach by periodically increasing the TP of the data packet to P_{\max} for enough time to protect the reception of the ACK at the source node. While this class of TPC protocols achieves good reduction in energy consumption (relative to the 802.11 MAC protocol), at best it gives comparable throughput to that of the 802.11 scheme. The main reason is that, as in the 802.11 approach, RTS and CTS messages are used to *silence* neighboring nodes, preventing concurrent transmissions from taking place over the reserved floor¹.

Throughput-oriented TPC schemes (e.g., [20, 33, 21]) use per-packet TPC to increase the spatial channel reuse. These schemes allow for concurrent transmissions in the same vicinity of a receiver by locally broadcasting *collision avoidance information* (CAI) over a separate control channel. In [20] the authors proposed the PCMA protocol, in which the receiver advertises its *interference margin* by sending busytone pulses over a separate control channel. The use of a separate control channel in conjunction with a busytone scheme was proposed in [33], where the sender transmits the data packets and the busytones at reduced power, while the receiver transmits its busytones at maximum power. A node estimates the channel gain from the busytones and is al-

¹The *reserved floor* or the *maximum transmission range* is the largest area around the transmitter over which a packet can be successfully received in the absence of interference from other nodes.

lowed to transmit if its transmission is not expected to add more than a fixed interference to the ongoing receptions. The authors, however, make strong assumptions on the interference power. Specifically, they assume that the antenna is able to reject any interfering power that is less than the power of the “desired” signal (i.e., they assume perfect capture). Furthermore, the schemes in [20, 33] do not address the issue of contention among busytone or their energy overhead. The PCDC protocol [21] uses two frequency-separated channels for data and control. RTS and CTS packets are transmitted over the control channel, providing CAI that facilitates interference-limited concurrent transmissions in the same vicinity.

Although the simulations of the TPC schemes in [20, 21, 33] indicate impressive throughput performance, we see four major design issues with these schemes that make their practicality questionable:

- In [20, 21, 33], the channel gain is assumed to be the same for both the control (or busytone) and data channels, and that nodes are able to transmit on one channel and, simultaneously, receive on the other. It is very difficult to achieve these two assumptions simultaneously. For the first assumption to hold, the control channel must be within the coherence bandwidth of the data channel, which places an upper bound on the allowable frequency separation between the two channels. However, for the second assumption to hold, there must be some minimal frequency spacing between the two channels. Typically, a 5% guard-band of the nominal RF frequency is needed to keep the price and complexity of the transceiver at a reasonable level [28]. However, spacing the control (or busytone) and data channels by that much can make the first assumption invalid! Another issue is that the spectrums allocated to the control and data channels should not be equal, since otherwise a lot of bandwidth would be wasted on signaling and control. However, fading is spectrum dependent, so the two channels would experience different amounts of fading.
- To be able to receive/transmit and simultaneously receive/transmit over two channels, the mobile node must be equipped with two transceivers. The complexity and cost of the additional hardware may not justify the increase in throughput. Furthermore, it is unfair to compare the performance of these protocols to the single-channel, single-transceiver IEEE 802.11 scheme.
- Interoperability with existing standards and hardware is difficult if not impossible. Currently, most wireless devices implement the IEEE 802.11b standard. The class of two-channel protocols is not backward-compatible with the IEEE 802.11 standard, which makes it difficult to deploy such schemes in real networks.
- Finally, the optimal allocation of the total spectrum between the data and control channels is load dependent. So for the allocation to be optimal under various traffic loads, it has to be adjusted adaptively, which is not feasible in practice.

The protocol proposed in this paper addresses all the above issues and provides a comprehensive, throughput-oriented

MAC solution for MANETs using a single-transceiver, single-channel approach.

Before closing this section, we mention few other schemes in the literature that tackle the problem of power control from a completely different perspective. The COMPOW protocol [23] relies completely on routing-layer agents to converge to a *common* power level for all network nodes. However, for constantly moving nodes, the scheme (like any other routing-protocol-based scheme) incurs significant overhead, and convergence to a common power level may not be possible. Moreover, in situations where network density varies widely (i.e., nodes are clustered), restricting all nodes to converge to a common power level is a conservative approach. A clustering approach was proposed in [18], which simplifies the forwarding function for most nodes but at the expense of reducing network utilization (since all communications have to go through an elected node). This can also lead to the creation of bottlenecks. A joint clustering/TPC protocol was proposed in [17], 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. The routing overhead in this protocol grows with the number of routing agents, and can be significant even for simple mobility patterns (note that for the DSR routing protocol, for example, routing packets account for approximately 38% of the total received bytes [14]). The protocol in [5] is energy-oriented and is basically a mechanism to learn the minimum TP level required for a node to successfully transmit to a neighboring node. This approach, however, suffers from the hidden node problem (see [21] for more details). Another novel approach for TPC is based on joint scheduling and power control [10]. This approach requires a central controller to execute the scheduling algorithm, i.e., it is not a truly distributed solution. Furthermore, it makes a number of strong assumptions, including synchronization, that each node knows the geographical location of all other nodes, and that the SINR measurement at each receiver is known to *all* transmitters.

3. THE PROPOSED POWMAC PROTOCOL

3.1 Assumptions

In designing POWMAC, we assume that the channel gain is stationary for the duration of few control and one data packet transmission periods. As discussed in detail in Section 3.6, this assumption holds for typical mobility patterns and transmission rates. We also assume that the gain between two nodes is the same in both directions. This is the underlying assumption in any RTS/CTS-based protocol, including the IEEE 802.11 scheme. Finally, 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. Off-the-self wireless cards (e.g., [3]) readily provides such measured values. In POWMAC, each node is equipped with one transceiver and a carrier-sense hardware (i.e., a basic IEEE 802.11 compliant transceiver).

3.2 Overview of POWMAC

POWMAC is distributed, asynchronous, and adaptive to channel changes. It requires a very simple circuitry, and

is compatible with the IEEE 802.11 single-channel, single-transceiver architecture. It only requires software changes to the IEEE 802.11 Ad Hoc MAC protocol.

The key features of POWMAC are as follows. First, unlike the IEEE 802.11 approach (and the schemes in [5, 12, 15, 16, 25]), POWMAC does not use the control packets (i.e., RTS/CTS) to silence neighboring nodes. Instead, CAI is inserted in the control packets and is used in conjunction with the received signal strength of these packets to dynamically *bound* the TP of potentially interfering nodes in the vicinity of a receiving node. The exact details of this mechanism is presented in Section 3.4. The second main feature of POWMAC is that the required TP of a data packet is computed at its intended receiver, say node i , according to a predetermined network *load factor*. The rationale behind this approach is to allow for some interference tolerance at receiver i , so that multiple interference-limited transmissions can simultaneously take place in the neighborhood of i . The tradeoffs in determining this load factor are discussed in Section 3.3.

The third feature of POWMAC is that control packets are transmitted at an adjustable power level to reach all and only potentially interfering nodes. This improves the spatial reuse for the control packets themselves and reduces their collisions. Section 3.5 presents the details of this approach. Finally, in POWMAC, after nodes exchange their control packets, they refrain from transmitting their data packets for a certain duration, which we call the *access window* (AW). The AW allows other neighboring nodes to exchange their control packets such that (interfering) data transmissions can proceed simultaneously (as long as collisions are prevented). The AW consists of an adjustable number of fixed-duration time slots. As explained later, each node *adaptively* varies the number of access slots in its AW, depending on network load. The AW is needed for two reasons. First, it reduces the likelihood of collisions between control and data packets. Once the reception of a data packet at a node, say i , has started, all nodes in i 's vicinity should not transmit any *control* packets. Recall that control packets are often transmitted at a high power to reach all potential interferers. So, it is very likely that these packets will collide with the ongoing data reception at node i . On the other hand, by using the AW, node i can allow its neighbors to exchange their RTS and CTS packets before i 's data reception starts, and when possible, to have their own data packets proceed simultaneously with i 's reception. Note that data packets are transmitted at a reduced power level, and so multiple data packets can be transmitted concurrently and still be received correctly.

Another reason why the AW is needed is to be able to inform nodes that are currently transmitting or receiving of the ensuing data transmission. Because POWMAC uses a single-channel, single-transceiver architecture, nodes can either transmit or receive at a given time, but not both. As a result, a terminal, say i , is basically “deaf” while transmitting, so it does not hear any transmitted control packets in its vicinity. Consequently, when i becomes idle, its information about ongoing receptions in its vicinity may be outdated, leading to collisions (i.e., if i decides to transmit again). The protocols in [20, 21, 33] alleviate the above problem by using a two-channel, two-transceiver architec-




Figure 2: Basic operation of POWMAC.

ture; nodes are able to transmit/receive their data packets and still hear the control signals. In contrast, POWMAC uses an AW to allow other neighboring nodes to exchange their control packets so that the ensuing data transmissions can proceed simultaneously.

We note here that the use of a time window between the reservation and the data packet channels has also been used in the Medium Access via Collision Avoidance with Enhanced Parallelism (MACA-P) protocol [4]. However, in that work the objective was not to address TPC, but rather to prevent collisions between control and data packets.

We conclude this section with an example that illustrates the basic operation of POWMAC (see Figure 2). The network topology is the one shown in Figure 1. Node A transmits an RTS to node B at a maximum (known) power (P_{\max}). Node B replies back with a CTS packet that is sent at an adjustable power level to reach all and only potentially interfering nodes. The RTS/CTS exchange allows nodes A and B to decide on the TP of the data packet. It also provides a way to inform potentially interfering nodes (i.e., node E) of the power that they can use without disturbing the communication $A \rightarrow B$. Node A confirms that the transmission $A \rightarrow B$ can proceed using a third packet, which we call Decide-To-Send (DTS), that is sent at power P_{\max} . Besides other reasons mentioned in Section 3.4, the DTS packet is needed to inform the neighbors of A about the power level that A will use for its data transmission. As explained later, this information is needed so that A 's neighbors (i.e., node D) can determine whether or not they can receive a data packet from some other node (e.g., C) simultaneously while A is transmitting to B . After the RTS/CTS/DTS exchange, node A refrains from sending its data packet for the AW duration. This allows E and F to exchange control packets and, hence, start the transmission $E \rightarrow F$ if it does not disturb $A \rightarrow B$.

3.3 Load Control

Load control is a concept that allows receivers to determine the required TP of their intended data receptions. Note that in MANETs, power is allocated only once on a per-packet basis, i.e., the whole data packet is transmitted at

one power level. This power, therefore, must be selected carefully. If the power used to transmit a data packet to a node, say i , is only enough to overcome the current interference at i , then none of i 's neighbors should be allowed to transmit during i 's reception, which negatively impacts the aggregate throughput. On the other hand, if the data transmission power is too high, it induces high interference on other nodes in the vicinity of the transmitter, thus, preventing them from receiving.

The load factor at node i , denoted by $\xi^{(i)}$, is a measure of the activity in the neighborhood of i . Formally, it is defined as²:

$$\xi^{(i)} \stackrel{\text{def}}{=} \frac{P_{\text{thermal}} + P_{\text{MAI}}^{(i)}}{P_{\text{thermal}}} \quad (1)$$

where P_{thermal} is the thermal noise and $P_{\text{MAI}}^{(i)}$ is the multi-access interference (MAI) from interfering transmissions at receiver i . Now, consider the transmission of a packet from node j to node i . Let μ^* be the SINR required to achieve the target bit error rate (BER) at receiver i . The minimum TP that is needed to achieve this BER is $\mu^*(P_{\text{thermal}} + P_{\text{MAI}}^{(i)})/G_{ji} = \mu^*\xi^{(i)}P_{\text{thermal}}/G_{ji}$, where G_{ji} is the channel gain from node j to node i . While more capacity can be achieved by increasing the load factor (i.e., allowing larger $P_{\text{MAI}}^{(i)}$), this also increases the power needed to transmit the packet, which in turn increases energy consumption. Energy is a scarce resource in MANETs, so it is undesirable to trade it off for throughput. Moreover, the Federal Communications Commission (FCC) regulations put a limit on the maximum power that can be used by nodes in the 2.4 GHz spectrum (e.g., 1 Watt for 802.11 devices). Given this limit, as the load is increased, the channel gain must be increased (with μ^* and P_{thermal} being constants), and so the maximum range (or coverage) for reliable communication will decrease.

The above factors collectively necessitate *load planning*, i.e., imposing a maximum load factor (MLF), denoted by ξ_{max} , that nodes are allowed to use. This ξ_{max} is set at the design phase to reflect several goals (e.g., throughput, network lifetime, etc.). One possible design choice that takes the above factors into account is as follows. First, we require that the maximum range, say d_{max} , of POWMAC to be the same as that of the 802.11 scheme. At this maximum range, we let the TP in POWMAC to be equal to ξ_{max} times the power used in the 802.11 standard (assuming no MAI). Thus, ξ_{max} must be chosen such that it is not greater than the ratio of the power limit set by the FCC and the power used in the 802.11 scheme.

For the average energy consumption to be the same, we choose ξ_{max} in a manner that maintains the same *energy per bit* consumed in the 802.11 scheme. The value of ξ_{max} that achieves the above goals can be derived as follows. We assume that the TP attenuates with the distance d as k/d^n (k is a constant and $n \geq 2$ is the loss factor). The required TP in POWMAC is given by:

$$P_{\text{POWMAC}} = \frac{\mu^* \xi_{\text{max}} P_{\text{thermal}} d^n}{k} \quad (2)$$

One can assume that the distance d is uniformly distributed

²This definition is somewhat similar but not identical to the definition used in [24] for cellular systems.

from zero to d_{max} . Then,

$$E[P_{\text{POWMAC}}] = \frac{\mu^* \xi_{\text{max}} P_{\text{thermal}} d_{\text{max}}^n}{k(n+1)}. \quad (3)$$

As for the 802.11 protocol, its corresponding TP is:

$$P_{802.11} = \frac{\mu^* P_{\text{thermal}} d_{\text{max}}^n}{k}. \quad (4)$$

Note that $P_{802.11}$ does not depend on d since the 802.11 scheme uses a fixed TP. Now, to achieve the same average energy per bit, we equate (3) and (4). Accordingly, $\xi_{\text{max}} = n + 1$. As an example, consider the two-ray propagation model with $n = 4$. Then $\xi_{\text{max}} = 7$ dB, which lies within the range of values used in already deployed cellular systems [24].

3.4 Channel Access Mechanism

Given a predetermined MLF, the purpose of the channel access mechanism is to allow the transmitter and the receiver to agree on the required TP such that the MLF is not exceeded and such that the ensuing data transmission does not disturb any of the scheduled data receptions. We now describe the details of the distributed channel access mechanism used in POWMAC.

Each node i maintains a Power Constraint List (PCL(i)). This list is an extension of the Network Allocation Vector (NAV) used in the IEEE 802.11 scheme. Basically, PCL(i) encodes i 's knowledge about other *active* nodes (i.e., nodes that are either receiving, transmitting, or scheduled to do so) in i 's vicinity. For every active node u in i 's vicinity, PCL(i) contains the following entries (how node i gets this information will become clear shortly):

- The address of node u .
- The estimated channel gain G_{iu} between nodes i and u computed using the received signal strength of u 's control packet.
- The maximum tolerable interference (MTI) of node u , denoted by $P_{\text{MTI}}^{(u)}$. This is the maximum additional interference that node u can tolerate from node i such that u 's reception is not disturbed. As will be explained shortly, this information is advertised by node u .
- The start time and duration of u 's activity (reception or transmission), also advertised by node u in its control packet.

Let $\pi_i(u)$ be the maximum TP that node i can use without disturbing u 's reception. Using G_{iu} and $P_{\text{MTI}}^{(u)}$, node i computes $\pi_i(u)$ as:

$$\pi_i(u) = \min \left\{ \frac{P_{\text{MTI}}^{(u)}}{G_{iu}}, P_{\text{max}} \right\}. \quad (5)$$

We define $P_{\text{MAP}}(i)$ as the maximum allowable TP that node i can use without disturbing any of its neighbors:

$$P_{\text{MAP}}(i) = \min_{u \in \text{PCL}(i)} \{ \pi_i(u) \}. \quad (6)$$

When a node, say j , has a packet to transmit to node i and does not sense a carrier, it sends an RTS packet at P_{max} and

includes in this packet the value of $P_{\text{MAP}}^{(j)}$. The format of j 's RTS packet is:

$$\text{RTS}(j \rightarrow i) = \left\{ j, i, P_{\text{MAP}}^{(j)}, N_{\text{AW}}^{(j)}, T_{\text{data}}^{(ji)} \right\} \quad (7)$$

where $T_{\text{data}}^{(ji)}$ is the transmission duration of j 's data packet and $N_{\text{AW}}^{(j)}$ is the remaining number of the access slots in the current AW. $N_{\text{AW}}^{(j)}$ allows the neighbors of node j to determine when j will start its data transmission. The duration of each access slot is fixed and common to all nodes. Specifically, an AW slot consists of the amount of time it takes to send an RTS, CTS, and DTS packets, plus the maximum duration of the backoff interval (the exact contention resolution mechanism will be explained later in this section).

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

$$P_{\text{min}}^{(ji)} = \frac{\mu^* (P_{\text{thermal}} + P_{\text{MAI-current}}^{(i)})}{G_{ji}} \quad (8)$$

where $P_{\text{MAI-current}}^{(i)}$ is the *current* MAI from all already on-going (interfering) transmissions. This $P_{\text{min}}^{(ji)}$, however, does not allow for any interference tolerance at node i ³. Now, according to the load planning calculations in Section 3.3, the power that node j is allowed to use to send to i is given by:

$$P_{\text{loading}}^{(ji)} = \frac{\mu^* \xi_{\text{max}} P_{\text{thermal}}}{G_{ji}}. \quad (9)$$

If $P_{\text{loading}}^{(ji)} < P_{\text{min}}^{(ji)}$, then the MAI in the vicinity of node i is greater than the one allowed by the planned loading. In this case, i responds with a negative CTS, informing j that it cannot proceed with its transmission (the negative CTS is used to prevent multiple RTS retransmissions from j). The philosophy behind this design is to prevent transmissions from taking place over links that perceive high MAI. This consequently increases the number of active links in the network (subject to the available power constraints).

On the other hand, if $P_{\text{loading}}^{(ji)} > P_{\text{min}}^{(ji)}$, then it is possible for i to receive j 's signal but only if $P_{\text{loading}}^{(ji)}$ is less than $P_{\text{MAP}}^{(j)}$ (which is included in the RTS). This last condition is necessary so that transmitter j does not disturb any of the ongoing receptions in its own vicinity. If i finds that it is possible for it to receive j 's signal, it calculates the *maximum additional interference power* ($P_{\text{MAI-add}}^{(i)}$) that it can endure from future unintended transmitters so that the SINR at i does not drop below μ^* . It can be shown that $P_{\text{MAI-add}}^{(i)}$ is given by:

$$P_{\text{MAI-add}}^{(i)} = \frac{G_{ji}}{\mu^*} (P_{\text{loading}}^{(ji)} - P_{\text{min}}^{(ji)}). \quad (10)$$

³In [26] the authors derived a finite value for the interference range in the case of minimum TP. However, the thermal noise power was not taken into account in that derivation.

The next step is to equitably distribute this power tolerance $P_{\text{MAI-add}}^{(i)}$ among future potential interferers in the vicinity of i . The rationale behind this distribution is to prevent one neighbor from consuming the entire $P_{\text{MAI-add}}^{(i)}$. In other words, we think of $P_{\text{MAI-add}}^{(i)}$ as a network resource that should be shared among various nodes. Recall that j 's RTS contains $N_{\text{AW}}^{(j)}$. Obviously, the number of concurrent transmissions should not exceed $N_{\text{AW}}^{(j)}$. Thus, the number of future potential interferers is set to $N_{\text{AW}}^{(j)}$.

Now, the MAI at node i consists of MAI due to interferers that are within the maximum range of i (denoted by $P_{\text{MAI-within}}^{(i)}$) and MAI due to interferers outside that range (denoted by $P_{\text{MAI-other}}^{(i)}$). While node i can have control over $P_{\text{MAI-within}}^{(i)}$, it cannot influence $P_{\text{MAI-other}}^{(i)}$. We take this into account as follows. In line with cellular systems, we assume that $P_{\text{MAI-other}}^{(i)} = \alpha P_{\text{MAI-within}}^{(i)}$, where $\alpha < 1$ and depends mainly on the propagation path loss factor (in practice, $\alpha \approx 0.5$ for the two-ray model [24]). Accordingly, the maximum tolerable interference $P_{\text{MTI}}^{(i)}$ that *each* future neighbor can add to node i is given by

$$P_{\text{MTI}}^{(i)} = \frac{P_{\text{MAI-add}}^{(i)}}{(1 + \alpha) N_{\text{AW}}^{(j)}}. \quad (11)$$

When responding to j 's RTS, node i indicates in its CTS the power level $P_{\text{loading}}^{(ji)}$ that j must use for the data transmission. In addition, node i inserts $P_{\text{MTI}}^{(i)}$ in the CTS packet to inform its neighbors of the maximum power they can use such that i 's reception is not disturbed. The CTS is sent at an adjustable power ($P_{\text{CTS}}^{(i)}$), as explained in the next section. The format of the CTS packet is:

$$\text{CTS}(i \rightarrow j) = \left\{ i, j, P_{\text{loading}}^{(ji)}, P_{\text{MTI}}^{(i)}, P_{\text{CTS}}^{(i)}, N_{\text{AW}}^{(j)}, T_{\text{data}}^{(ji)} \right\}. \quad (12)$$

Upon receiving i 's CTS, node j replies back with a DTS packet that includes the value of $P_{\text{loading}}^{(ji)}$. The DTS is needed to inform j 's neighbors that may have not heard i 's CTS about $P_{\text{loading}}^{(ji)}$. Such neighbors can determine if they can receive data packets following the current AW depending on how high the interference is due to the scheduled transmission $j \rightarrow i$. In addition, as in [8], the DTS packet in POWMAC announces the success of the RTS/CTS exchange between j and i to those neighbors of j who have not heard i 's CTS. The IEEE 802.11 uses carrier sensing to achieve this goal; if the neighbors of j do not sense a carrier after hearing the RTS for some time, they assume that the RTS/CTS exchange was not successful. This same mechanism, however, cannot be used in POWMAC since the data packet is transmitted at a power less than the RTS power, and thus the carrier sense range of the data packet is much smaller than the RTS (or CTS) one. The DTS is sent at P_{max} with the following format:

$$\text{DTS}(j \rightarrow i) = \left\{ j, i, P_{\text{loading}}^{(ji)} \right\}. \quad (13)$$

A potentially interfering node, say r , that hears i 's CTS uses the signal strength of the received message and the value of the CTS transmission power (included in the CTS) to estimate the channel gain G_{ri} between itself and node i . The

channel gain and the broadcasted $P_{\text{MTI}}^{(i)}$ value are used to update the maximum power $P_{\text{MAP}}^{(r)}$ that r can use in its future transmissions, according to (6). If r wants to transmit a data packet, then the interference caused by r 's data transmission to i 's data reception must not exceed $P_{\text{MTI}}^{(i)}$. The CTS does not place any restriction on node r in case r wants to receive (after r has received an RTS).

A potentially interfering node, say u , that hears j 's DTS packet (u could be r) uses the predetermined P_{max} value and the power of the received DTS to estimate the channel gain G_{uj} between itself and node j . The channel gain and the broadcasted $P_{\text{loading}}^{(ji)}$ value are used to compute the interference that node u will perceive due to j 's data transmission. Using this information, node u determines whether or not it can receive any data packet following the current AW. Specifically, the value of $P_{\text{MAI-current}}^{(u)}$ is increased by the amount of interference due to j 's scheduled transmission. Now, if u wishes to receive a data packet (after receiving an RTS), then the interference caused by j 's data transmission to u 's data reception must not exceed the planned loading (i.e., not more than $\xi_{\text{max}} P_{\text{thermal}}$). The DTS does not place any restriction on node u in case u wants to transmit.

For contention resolution, we follow the work in [22]. If r (or u) wants to transmit a data packet, then r contends for the channel in the next slot of j 's AW with probability $p(r)$, which is the probability of network loses as perceived by node r ⁴. If successful, node r chooses a wait time B_r that is uniformly distributed in the interval $[0, B]$; B is a system-wide backoff counter. After this waiting time, node r senses the channel. If the channel is free, node r transmits its RTS. Note that when using this mechanism for POWMAC, we do *not* require any synchronization. Basically, once node r receives j 's RTS, it divides its time access into $N_{\text{AW}}^{(j)}$ slots of predetermined length, regardless of the absolute time at nodes r and j .

The AW size at a node, say j , is computed adaptively as a function of the load in the vicinity of j . At the end of the data reception at node j , if the actual interference perceived by node j is higher than a certain percentage (e.g., 75%) of the planned interference ($\xi_{\text{max}} P_{\text{thermal}}$), then the AW size need not be changed, since the allocated additional power to combat MAI was efficiently used. On the other hand, if less than that threshold was used, then node j checks the number of concurrent transmissions that actually took place in that AW (based on numbers of CTS and DTS packets). If this number is less than, say $\eta\%$, of the AW size, then either the load is low or the value of the AW size is too big to the extent that $P_{\text{MTI}}^{(j)}$ is too small (see (11)), i.e., $P_{\text{MTI}}^{(j)}$ is not large enough to allow for other nearby nodes to transmit. In both cases, it makes sense that node j decreases its AW size. To prevent large fluctuations in the AW values, the AW size is incremented or decremented in steps of 1.

Following a successful reception of a data packet, receiver i responds with an ACK packet, which is transmitted using the same power level that would have been used if i were to send a data packet to j . We assume that enough FEC code is used to protect ACK packets from most types of collisions (given the small size of the ACK packet, the FEC overhead

⁴The exact details of how this probability is calculated is beyond the scope of this paper. See [22] for details.

is not significant). A similar argument has been used in previously proposed protocols (e.g., [20, 21]).

The approach we discussed in this section provides a distributed mechanism for admission control. In contrast to cellular systems where the base station makes the admission decision, in here each node, and depending on previously heard RTS, CTS, and DTS packets, decides whether its transmission can proceed or not.

3.5 Adaptive Reservation Mechanism

According to the IEEE 802.11 scheme, nodes transmit the RTS and CTS packets using a fixed power P_{max} . As discussed in Section 1.1, this approach can be overly conservative. Recall that in POWMAC, a receiver, say i , sends a CTS packet that contains CAI, namely $P_{\text{MTI}}^{(i)}$, to bound the TP of potentially interfering neighbors. A node, say v , that hears this packet sets its $P_{\text{MAP}}^{(v)}$ according to (6). If P_{max} is less than $P_{\text{MTI}}^{(i)}/G_{vi}$, the CAI is actually irrelevant to node v , and the CTS packet has reached farther than necessary. In POWMAC, this issue is not harmful as in the IEEE 802.11 scheme, simply because control packets in POWMAC do not prevent neighbors from transmitting. Nonetheless, one way to further enhance the operation of POWMAC is to transmit control packets only to those nodes who can actually make use of the CAI. The advantage of doing so is that transmitting control packets at a reduced power reduces contention among the control packets themselves and increases the spatial reuse. Therefore, POWMAC uses the following adaptive TP approach for the control packets.

The farthest neighbor from node i , say v , that can actually make use of the CAI contained in i 's CTS is the one with channel gain $G_{iv} = P_{\text{MTI}}^{(i)}/P_{\text{max}}$. For any other node, say k , that is more than G_{iv} away from i , P_{max} is less than $P_{\text{MTI}}^{(i)}/G_{ik}$, and thus the CAI that is contained in i 's CTS is irrelevant to node k . Therefore, we set the range of the CTS of node i to $P_{\text{MTI}}^{(i)}/P_{\text{max}}$. Accordingly, the TP for the CTS packet of node i is:

$$P_{\text{CTS}}^{(i)} = \min \left\{ \mu^* P_{\text{thermal}} \frac{P_{\text{max}}}{P_{\text{MTI}}^{(i)}}, P_{\text{max}} \right\} \quad (14)$$

where the minimum is taken because of the hardware constraints of the wireless interface. This could mean that the control packets are not reaching all potential interferers. However, the IEEE 802.11 scheme itself does not solve this problem; it does not prevent nodes in the interference range from causing collisions with the data packet (see [15] for details). Thus, this problem is not introduced by the proposed protocol. Note that we take into account future MAI due to nodes that do not hear the control packets by using $\alpha = 0.5$ in (11).

3.6 Mobility and POWMAC

To determine the TP for data packets, POWMAC relies on the assumption that the channel gain determined at the time of the RTS/CTS/DTS exchange is stationary for the duration of the current AW and the ensuing data packets. Since the channel gain can change as a result of mobility, we now comment on the role of mobility in the design of POWMAC.

For *large-scale* channel variations (e.g., mean channel gain),

mobility has negligible impact on POWMAC since packet transmission times are on the order of few milliseconds while nodes mobility is often on the order few meters per second. So the time between a control packet and an ensuing data packet is small enough to make the estimation sufficiently accurate. As for *small-scale* channel variations, although their impact can be mitigated through diversity techniques at the physical layer (e.g., RAKE receivers [28]), we now argue that even if such techniques are not available, the “stationarity” of the channel gain in POWMAC is still valid. Consider a multipath environment, where multiple versions of the transmitted signal arrive at the receiver at slightly different times and combine to give a resultant signal that can vary widely in amplitude and phase. The spectral broadening caused by this variation is measured by the Doppler spread, which is a function of the relative velocity (v) of the mobile and the angle between the direction of motion and the directions of arrival of the multipath waves [28]. This variation can be equivalently measured in the time domain using the *coherence time* (T_c), which is basically a statistical measure of the time duration over which the channel can be assumed time invariant. As a rule of thumb in modern communication system, $T_c \approx 0.423/f_m$, where $f_m = v/\lambda$ is the maximum Doppler shift and λ is the wavelength of the carrier signal.

Now, at a mobile speed of $v = 1$ meter/sec and 2.4 GHz carrier frequency, $T_c \approx 52.89$ msec. This time reduces to 10.56 msec when $v = 5$ meters/sec. For the channel stationarity assumption to be valid, the access window and the data packet duration must not exceed T_c . At channel bandwidth of 2 Mbps, it takes 4 msec to transmit a 1000-byte packet. The propagation delay and the turnaround time (time it takes a node to switch from a receiving mode to a transmitting mode) are in the order of microseconds, and so they can be safely ignored. Thus, the assumption about channel stationarity is valid for moderate values of velocity (e.g., pedestrians). The IEEE 802.11 was designed for such velocity scenarios [6]. In cases when nodes move faster, the packet size can actually be shortened so that the stationarity assumption can still hold.

One final note about the practicality and delays of changing the power levels from the time the node sends a control packet until that node sends its data packet. 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 $2 \mu\text{sec}$,” and that “the transmit power-down ramp for 90% to 10% maximum power shall be no greater than $2 \mu\text{sec}$.” Given that the transmission duration of the RTS (or CTS) is in the order of tens of microseconds and the data transmission duration is in the order of milliseconds, the delay attributed to changing power levels ($\leq 2 \mu\text{sec}$) can be safely ignored. Note also that this delay is less than the turnaround period, which is approximately $5 \mu\text{sec}$. Hence, after a node, say 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.

4. PERFORMANCE EVALUATION

4.1 Simulation Setup

We now evaluate the performance of the POWMAC pro-




Figure 3: Toy topology where the two interfering transmissions $A \rightarrow B$ and $C \rightarrow D$ can proceed simultaneously if A 's and C 's transmission powers are appropriately chosen.

ocol and contrast it with the IEEE 802.11 scheme. Note that we do not compare POWMAC to energy-oriented protocols (e.g., [12, 15, 16, 25]), since at best these protocols give comparable throughput to that of the 802.11 scheme. Furthermore, since POWMAC uses a single-channel, single-transceiver design, it is unfair to compare it with two-channels, two-transceivers based protocols (e.g., [20, 21, 33]). 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. For simplicity, data packets are assumed to be of a fixed size.

The routing overhead is ignored since the goal here is to evaluate the performance improvements due to the MAC protocol. Furthermore, because the interference margin is chosen so that the maximum transmission range under the POWMAC and 802.11 protocols is the same, it is safe to assume that both protocols achieve the same forward progress per hop. Consequently, we can focus on the one hop throughput, i.e., the packet destination is restricted to one hop from the source. The two-ray propagation model is used, and the capture model is similar to the one in [31]. Other parameters used in the simulations are given in Table 1. These parameters correspond to realistic hardware settings [3].

Data packet size	2 KB
Data rate	1 Mbps
SINR threshold	6 dB
Transmission range	750 meter
Carrier-sense range	1500 meter
ξ_{max}	7 dB

Table 1: Parameters used in the simulations.

4.2 Macroscopic Results

We first simulate a set of simplified scenarios for the purpose of highlighting the advantages and operational details of POWMAC. Consider the line topology in Figure 3. The distances between the nodes are also shown in the figure. Node A is transmitting to node B , and node C is transmitting to node D . Persistent load is used in this experiment, i.e., nodes A and C always have packets to send. The transmissions from A and C interfere with data receptions at D and B respectively. However, the interference from A to D is much smaller than the one from C to B , and so we will only focus on the latter one.

In the first scenario, node B starts moving in the direction of node C at speed of 10 m/s. Figure 4(a) depicts the throughput of the network as a function of time. According to the 802.11 scheme, only one transmission can proceed at a time since all nodes are within the carrier-sense range of each other. However, according to POWMAC, for the first 16 seconds, the two transmissions $A \rightarrow B$ and $C \rightarrow D$ can proceed simultaneously, resulting in about 100% improvement in network throughput.

For the next ten seconds, as node B gets closer to node C , the channel gain G_{BC} increases and so $P_{\text{MAP}}^{(C)}$ decreases. Eventually, $P_{\text{MAP}}^{(C)}$ becomes less than the one required by node D to achieve its SINR threshold. Therefore, once node A exchanges RTS, CTS, and DTS packets with B , node C cannot transmit to D ⁵. On the other hand, if node C exchanges RTS/CTS/DTS packets with D before A does that with B , then node A actually increases its power to overcome the interference induced from C at node B . Hence, the two transmissions $A \rightarrow B$ and $C \rightarrow D$ can proceed simultaneously. Roughly, half of the time A starts before C and half of the time C starts before A . Therefore, the throughput enhancement is about 50% during the period between 16 and 26 seconds.

After 26 seconds, the interference at B due to C becomes larger than the one allowed by the planned loading and, thus, either $A \rightarrow B$ or $C \rightarrow D$ can proceed, but not both. The very small degradation in throughput after 26 seconds relative to the IEEE 802.11 scheme is due to the overhead of the AW when no simultaneous transmissions are taking place.

In the second scenario, node C moves in the direction of node B at a speed of 10 m/s, while all other nodes are stationary. Figure 4(b) shows the throughput of the network as a function of time. The difference between this scenario and the previous one is that this time, not only is $P_{\text{MAP}}^{(C)}$ decreasing (as a result of G_{BC} increasing), but $P_{\text{loading}}^{(CD)}$ is also increasing as a result of the decrease in G_{CD} . In the first 9 seconds, the two transmissions $A \rightarrow B$ and $C \rightarrow D$ can proceed simultaneously. Between 9 and 12 seconds, the throughput enhancement is 50% for the same reason given in the previous scenario. After that, only one transmission proceeds, and the throughput becomes comparable to that of the 802.11 scheme.

The last scenario we study is when D is moving away from C at a speed of 10 m/s, while the rest of the nodes are stationary. Figure 4(c) depicts the throughput of the network as a function of time. In here, as D moves away from C , G_{CD} decreases and so $P_{\text{loading}}^{(CD)}$ increases. For the first 15 seconds, the two transmissions $A \rightarrow B$ and $C \rightarrow D$ can proceed simultaneously. Between 15 and 20 seconds, the throughput enhancement is 50%. After that, $P_{\text{loading}}^{(CD)}$ becomes larger than $P_{\text{MAP}}^{(C)}$ in case A exchanges RTS/CTS/DTS with B before C does that with D , or the interference from C to B exceeds the planned loading at B , in the other case. Thus, only one transmission can proceed, and the throughput becomes comparable to that of the 802.11 scheme.

⁵When C sends an RTS to D , D replies with a negative CTS, since $P_{\text{MAP}}^{(C)}$ is less than $P_{\text{loading}}^{(CD)}$ as computed by node D .

4.3 Random Grid Topologies

We now study the performance under more generic network topologies. First, we consider a *random grid* topology, where 16 mobile nodes are placed within a square area of length 1500 meters. The square is split into 16 smaller squares, one for each node. The location of a mobile node within the small square is selected randomly. For each generated packet, the destination node is selected randomly from the one-hop neighbors. Each node generates packets according to a Poisson process with rate λ (same for all nodes). The Random Waypoint model [9] is used for mobility, with a node speed that is uniformly distributed between 0 and 2 meters/sec.

The performance is demonstrated in Figure 5. Part (a) of the figure depicts the throughput versus λ . It can be shown that at high loads, POWMAC achieves about 45% increase in throughput over the IEEE 802.11 scheme. This increase is attributed to the increase in the number of simultaneous transmissions.

Part (b) of Figure 5 depicts the energy consumption versus λ . This is the total energy used to *successfully* transmit a packet. It includes the energy used to transmit control packets and the lost energy in retransmitting data and control packets in case of collisions. For all cases, POWMAC requires roughly the same energy required by the 802.11 scheme. These results are in line with the analysis in Section 3.3, where the interference margin was chosen so that both protocols consume the same energy per bit.

4.4 Clustered Topologies

The authors in [19] argued that traffic locality is the key to determine the feasibility of large ad hoc networks. This motivates studying the performance of POWMAC under *clustered* topologies. In such topologies, a node 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 600×600 (in meters). Sixteen 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.

Part (a) of Figure 6 depicts the network throughput versus λ for $p = 0.25$. According to the 802.11 scheme, only one transmission can proceed at a time since all nodes are within the carrier-sense range of each other. However, according to POWMAC, two to three transmissions can proceed simultaneously, resulting in a significant improvement in network throughput. Part (b) of the figure shows that POWMAC saves a significant amount of energy relative to the 802.11 scheme. Since a node communicates mostly with nodes within its own cluster, the destination node is within 100 meters of the source node, thus requiring much less TP than P_{max} . This is the reason why the figure shows a huge advantage of POWMAC over the 802.11 scheme in terms of energy consumption. So, although POWMAC was designed with the goal of increasing throughput, significant energy may be saved as a consequence of reducing the TP.




Figure 4: Performance of the POWMAC and the 802.11 protocols (line topologies).

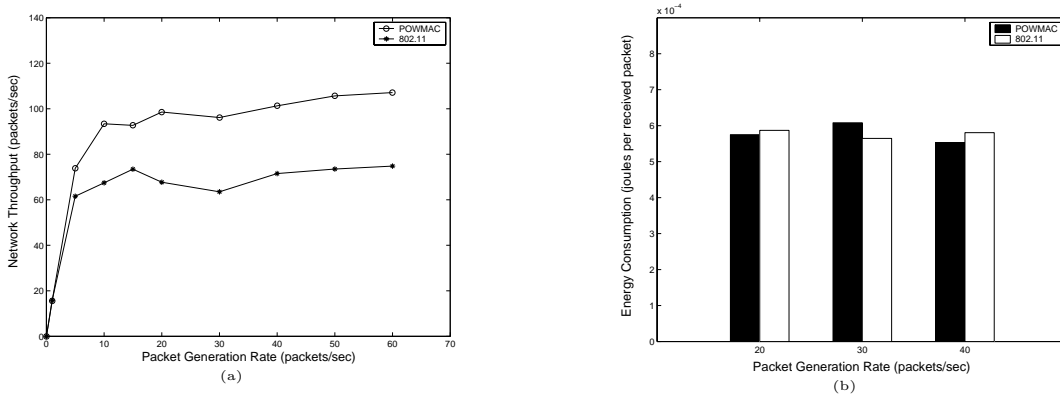


Figure 5: Performance of the POWMAC and 802.11 protocols as a function of λ (random grid topologies).

5. CONCLUSIONS AND FUTURE WORK

In this paper, we have proposed POWMAC, a power controlled MAC protocol for MANETs. Similar to the 802.11 scheme, POWMAC is based on a single-transceiver circuitry, and it operates over a single channel (for data and control packets). POWMAC adjusts the transmission powers of data packets to allow for some interference margin at the receiving nodes. Information about this interference margin is inserted into the CTS packet and is used to *bound* the transmission power of potentially interfering nodes in the vicinity of a receiver, rather than silencing such nodes. Multiple interference-limited transmissions in the vicinity of a receiver are allowed to overlap in time, provided that their MAI effects do not lead to collisions at that receiver.

We have compared the performance of POWMAC with that of the IEEE 802.11 scheme. Our simulation results showed that POWMAC can improve the network throughput by up to 45% in random grid topologies and much more than that in clustered topologies. Furthermore, POWMAC can achieve some reduction in the energy consumed to successfully deliver a packet from the source to the destination. To the best of our knowledge, POWMAC is the first single-channel protocol that utilizes TPC to increase network throughput while preserving the collision avoidance property of the 802.11 scheme.

Besides tuning the parameters of POWMAC and investi-

gating its performance under various scenarios and topologies, our future work will address other techniques for capacity improvement in POWMAC. Because of its demonstrated superior performance (compared to TDMA and FDMA), CDMA has been chosen as the access technology of choice in cellular systems including the recently adopted 3G systems. It is, therefore, natural to explore the potential of integrating CDMA into the design of the POWMAC protocol.

Acknowledgment

This work was supported in part by the National Science Foundation through grants ANI-0095626, ANI-0313234, and ANI-0325979; and in part by the Center for Low Power Electronics (CLPE) at the University of Arizona. CLPE is supported by NSF (grant # EEC-9523338), the State of Arizona, and a consortium of industrial partners.

6. 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>.




Figure 6: Performance of the POWMAC and 802.11 protocols as a function of λ (clustered topologies).

- [4] A. Acharya, A. Misra, and S. Bansal. MACA-P : a MAC for concurrent transmissions in multi-hop wireless networks. In *Proceedings of the First IEEE PerCom 2003 Conference*, pages 505–508, Mar. 2003.
- [5] 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.
- [6] K. Aretz, M. Haardt, W. Konhauser, and W. Mohr. The future of wireless communications beyond the third generation. *Computer Networks*, 37:83–92, 2001.
- [7] N. Bambos. Toward power-sensitive networks architecture in wireless communications: Concepts, issues, and design aspects. In *IEEE Personal Communications Magazine*, volume 5, pages 50–59, June 1998.
- [8] 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.
- [9] J. Broch, D. A. Maltz, D. B. Johnson, Y.-C. Hu, and J. Jetcheva. A performance comparison of multi-hop wireless ad hoc network routing protocols. In *Proceedings of the IEEE/ACM MobiCom Conference*, pages 85–97, 1998.
- [10] T. ElBatt and A. Ephremides. Joint scheduling and power control for wireless ad-hoc networks. In *Proceedings of the IEEE INFOCOM Conference*, pages 976–984, 2002.
- [11] 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*, pages 1506–1513, 2000.
- [12] J. Gomez, A. T. Campbell, M. Naghshineh, and C. Bisdikian. PARO: supporting dynamic power controlled routing in wireless ad hoc networks. *ACM/Kluwer Journal on Wireless Networks*, 9(5):443–460, 2003.
- [13] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2):388–404, Mar. 2000.
- [14] 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.
- [15] 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.
- [16] P. Karn. MACA - a new channel access method for packet radio. In *Proceedings of the 9th ARRL Computer Networking Conference*, pages 134–140, 1990.
- [17] V. Kawadia and P. R. Kumar. Power control and clustering in ad hoc networks. In *Proceedings of the IEEE INFOCOM Conference*, pages 459–469, 2003.
- [18] T. J. Kwon and M. Gerla. Clustering with power control. In *Proceedings of the IEEE MILCOM Conference*, pages 1424–1428, 1999.
- [19] 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.
- [20] 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*, pages 219–228, 2001.
- [21] A. Muqattash and M. Krunz. Power controlled dual channel (PCDC) medium access protocol for wireless ad hoc networks. In *Proceedings of the IEEE INFOCOM Conference*, pages 470–480, 2003.
- [22] T. Nandagopal, T.-E. Kim, X. Gao, and V. Bharghavan. Achieving mac layer fairness in wireless packet networks. In *Proceedings of the IEEE/ACM MobiCom Conference*, pages 87–98, 2000.
- [23] 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.
- [24] T. Ojanperä and R. Prasad. *Wideband CDMA For*

- Third Generation Mobile Communications*. Artech House, Incorporated, 1998.
- [25] M. B. Pursley, H. B. Russell, and J. S. Wycarski. 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.
- [26] D. Qiao, S. Choi, A. Jain, and K. G. Shin. Miser: an optimal low-energy transmission strategy for iee 802.11a/h. In *Proceedings of the 9th Annual International Conference on Mobile Computing and Networking*, pages 161–175, 2003.
- [27] R. Ramanathan and R. Rosales-Hain. Topology control of multihop wireless networks using transmit power adjustment. In *Proceedings of the IEEE INFOCOM Conference*, pages 404–413, 2000.
- [28] T. Rappaport. *Wireless Communications: Principles and Practice*. Prentice Hall, 2002.
- [29] V. Rodoplu and T. Meng. Minimum energy mobile wireless networks. *IEEE Journal on Selected Areas in Communications*, 17(8):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*, pages 1388–1397, 2001.
- [33] 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.