GMAC: A Game-theoretic MAC Protocol for Mobile Ad Hoc Networks

Fan Wang, Ossama Younis, and Marwan Krunz
Department of Electrical & Computer Engineering
University of Arizona
Tucson, AZ 85721
E-mail:{wangfan,younis,krunz}@ece.arizona.edu

***Abstract*— The conservative nature of the 802.11 ad hoc scheme has instigated extensive research whose goal is to improve the spatial reuse and/or energy consumption of this scheme. Transmission power control (TPC) was shown to be effective in achieving this goal. Despite their demonstrated performance gains, previously proposed power-controlled channel access protocols often incur extra hardware cost (e.g., multiple transceivers). Furthermore, they do not fully exploit the potential of power control due to the heuristic nature of power allocation and "interference margin" computations. In this paper, we propose a distributed, single-channel MAC protocol (GMAC) that is inspired by game theory. In GMAC, multiple potential transmitters are first involved in an admission phase, which enables terminals to compute the transmission powers that achieve a Nash equilibrium (NE) for the given utility function. Subsequently, successful contenders can simultaneously proceed with their transmissions. Simulation results indicate that GMAC improves the network throughput over the 802.11 scheme by about 80%, and over another single-channel power-controlled MAC protocol (POWMAC) by about 40%. These gains are achieved at no extra energy cost.**

I. Introduction

Mobile ad hoc networks (MANETs) are infrastructureless networks that provide an efficient solution when centralized control is infeasible (e.g., emergency and rescue operations, disaster-relief efforts, etc.). A key design objective in MANETs is achieving high network throughput while maintaining energy-efficient wireless communications for mobile terminals. Several attempts have been made to address these objectives [1], [17], [14].

The "ad hoc" mode of the IEEE 802.11 standard [7] has so far been used as the *de facto* MAC protocol for MANETs. This standard optionally uses an RTS/CTS (request-to-send/clear-to-send) handshake to coordinate channel access and resolve contention. Several studies have documented the poor performance of this protocol, which is attributed to its conservative treatment of potential interferers (the RTS/CTS packets are used to silence all overhearing terminals), its use of fixed transmission powers (TPs), and its inefficient handling of ACK packets (the exposed terminal problem). These problems are particularly acute in dense networks, where the transmitter-receiver distances are relatively small. To overcome these problems, researchers considered the use of power control at the MAC layer as a way to improve the spatial reuse (e.g., [12], [20], [14]) and/or reduce energy consumption (e.g., [3], [9]). Our emphasis in this paper is on the former class. In this class, terminals broadcast some *collision avoidance information* (CAI) to neighboring terminals. This information is used to *bound* the transmission powers of potential future transmitters in the neighborhood. For example, in POWMAC [14], CAI includes information about the maximum tolerable interference (MTI), defined as the amount of interference power that a receiver of a data or ACK packet is able to tolerate from one future interferer. Future transmitters use the overheard MTI values along with other information (e.g., channel gains, load tolerance, etc.) to determine their TPs. This way, multiple interference-limited transmissions can take place concurrently in the vicinity of the same receiver.

Despite their demonstrated performance gains, previously proposed TPC schemes suffer from problems such as backward incompatibility with the 802.11 architecture, extra hardware cost, and lack of ACK protection (see [14] for details). For instance, many of them offer dual-channel solutions, which often require two transceivers per terminal. Recently, the authors in [14] proposed POWMAC, a single-channel, single-transceiver solution that can achieve up to 40% throughput improvement over the 802.11 scheme and provide ACK protection. Although this is an impressive gain, we argue that POWMAC does not fully exploit the potential of power control. Specifically, the MTI (and hence the TPs) in POWMAC is determined heuristically, often unnecessarily silencing some possible transmissions.

To illustrate, consider a MANET of four terminals: *A*, *B*, *C*, and *D*, as depicted in Fig. 1. Let G_{ij} and d_{ij} be, respectively, the channel gain and distance between any two terminals i and j. Suppose that terminals A and C wish to transmit to terminals B and D, respectively, and suppose that A succeeds in transmitting its RTS before C sends its RTS. Assume that $G_{AB} = G_{CD}$, and that A and C are within the maximum transmission range of each other. Let SNR_{th} (required signal-to-noise ratio at the receiver) be 3 dB and assume a two-ray channel propagation model with a path loss factor of

4[1]. For simplicity, we ignore the thermal noise. Clearly, the spatial reuse is maximized when both transmissions $A \rightarrow B$ and $C \rightarrow D$ are allowed to proceed concurrently. In principle, this can be achieved with appropriate TPs as long as $d_{AC} \geq 0.3 d_{AB}$ is satisfied. In the 802.11 scheme, the successful transmission of one terminal's RTS will simply prevent any other terminal within the same neighborhood from transmitting. According to POWMAC [14], terminal A calculates the minimum power needed for correct reception, *heuristically* inflates this power by adding to it some "interference margin," and evenly distributes this margin among multiple future interfering terminals (whose number is not known in advance). Depending on the value of the interference margin, terminal C may not be able to transmit concurrently with A. In our example, POWMAC permits both transmissions to take place simultaneously as long as $d_{AC} > 0.7 d_{AB}$. If $0.7 d_{AB} \geq d_{AC} \geq 0.3 d_{AB}$, then POWMAC allows only one transmission to proceed at the same time. This inefficiency is due to the fixing of the MTI. Setting the value for the interference margin is non-trivial. A value that is too large adds unnecessary interference in the network and wastes energy, while a value that is too small prevents some feasible concurrent transmissions.

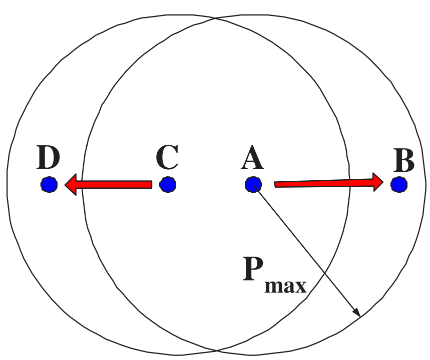

Fig. 1: Example of two transmissions that can proceed simultaneously if transmission powers are computed appropriately.

To select the appropriate TPs that maximize the spatial reuse, we formulate the channel contention problem as a non-cooperative power-control game and present a novel MAC protocol (GMAC) that implements this game in a distributed fashion. Game theory provides a powerful mathematical tool for decision-making among contending transmissions. It has been applied in the context of infrastructure-based (cellular) wireless networks (e.g, [19], [4], [8], [21]). In these approaches, each user attempts to adjust his individual power in order to maximize a given utility function that may incorporate conflicting goals (e.g., signal quality and energy consumption). Extending game theory to MANETs is challenging, due to the lack of a centralized infrastructure in these networks.

We use game theory to compute the TPs that achieve the Nash equilibrium (NE) for a number of contending transmitters within some neighborhood. GMAC supports a CSMA/CA-like access mechanism for distributing these powers to contending transmitters, hence enabling more concurrent transmissions than previously proposed techniques. GMAC can also accommodate adaptive-rate transmissions. The protocol is asynchronous, completely distributed, and uses only a single channel for both data and control packets. This ensures hardware compatibility with the 802.11 scheme. We also introduce a linear pricing function to obtain Pareto improvements in the achieved NE solution (this will be explained in Section III). To the best of our knowledge, GMAC is the first to exploit game theory in a distributed MAC protocol design for MANETs.

The rest of the paper is organized as follows. In Section II, we briefly discuss TPC schemes for MANETs and outline previous work on the application of game-theoretic power control analysis. Section III introduces our game-theoretic formulation for resolving channel contention in MANETs. Based on the above analysis, we present our GMAC protocol in Section IV. In Section V, we provide simulation results and insight on the performance of GMAC compared to previous approaches. Finally, we draw our conclusions in Section VI and propose some extensions for future work.

II. Related Work

TPC protocols that have been proposed in the literature focus on either reducing energy consumption (e.g., [3], [9]), or increasing network throughput (e.g., [12], [20], [14]). To reduce energy consumption, Agarwal et al. [3] assumed that each terminal possesses a fixed number of power levels, and proposed a mechanism that uses the minimum required power level for transmission. A similar approach was proposed in [9] to reduce energy consumption in networks where terminals are not uniformly distributed. These approaches can reduce the overall energy consumption, but achieve comparable throughput performance to that of the 802.11 standard.

Several *throughput-oriented* TPC MAC protocols were proposed in the literature, including PCMA [12], PCDC [13], and POWMAC [14] (see [10] for a complete overview of various TPC MAC protocols in MANETs). Our work in this paper belongs to this class of protocols. In PCMA [12], the receiver advertises its interference margin by sending busytone pulses over a separate control channel. The use of a control channel along with a busy tone scheme was proposed in [20], where the sender transmits the data packets and busy tones at a reduced power, while the receiver transmits its busy tones at the maximum possible power. The PCDC protocol [13] uses two frequency-separated channels for data and control packets, allowing for interference-limited concurrent transmissions in the vicinity. In contrast to the protocols in [13], [12], [20], the POWMAC protocol [14] uses a single channel for both data and control packets. Data packets are transmitted after several RTS/CTS exchanges take place. This enables the scheduling of multiple concurrent transmissions in the same vicinity, provided that a certain interference margin is not exceeded at each transmitter-receiver pair.

[1]It was shown in [18] that the two-ray ground reflection model gives more accurate prediction at long distances than the free-space model.

In [19], Saraydar et al. proposed a game-theory-based power control algorithm for data transmissions in cellular networks. According to this algorithm, the utility function is defined as the number of efficient bits transmitted per unit of energy. By using a linear pricing function for the transmission power, Pareto improvement of the game is achieved. Ji and Huang [8] formulated a framework for uplink power control in cellular networks using the function of signal-to-interference ratio and the transmitting power as the utility function. An iterative algorithm that searches for the equilibrium solution was introduced and analyzed for different scenarios. The same framework was adapted in [21] by using a different sigmoid utility function. Alpcan et al. [4] adopted a similar utility function to the one used in our paper. They also proposed two update algorithms and proved that these algorithms converge under certain conditions.

All the game-theoretic protocols described above consider infrastructure-based cellular networks. Our work, though driven by similar game theoretic analysis, is proposed for MANETs and is quite different from these protocols in several aspects. First, in MANETs, multi-hop communication is common due to the limited transmission range of the wireless terminals. Second, for uplink transmissions in cellular networks, interference occurs at the BS, which can be easily estimated. On the other hand, in MANETs, each reception will endure interference from all active transmissions in the network, particularly those in the receiver's proximity. Accounting for all possible interferences is very difficult. Another important difference is that in a single cell, terminals make their power decisions in parallel. However, in our MAC design, each receiver must have knowledge of all previously scheduled transmissions in its vicinity in order to make its decision. This will be explained in Section IV.

III. Formulation of the Power-Control Game

We first define an appropriate utility function for the power-controlled channel access game. Our goal is to maximize the overall network throughput while preventing terminals from using unnecessarily high power. We use the capacity (maximum achievable rate) as the figure of merit, and let the physical layer decide on the appropriate coding/modulation scheme that achieves this capacity[2]. We approximate the achievable rate by using Shannon's capacity [16]. Accordingly, the utility function for an active link i is defined as:

$$u_i(p_i, \mathbf{p}_{-i}) = \log(1+\gamma_i) - \alpha_i p_i \tag{1}$$

where p_i is the transmission power of i's transmitter (to be computed) and $\mathbf{p}_{-i}$ is a vector that represents the transmission powers of all links other than i. γ_i denotes the SINR at the receiver and is given by:

$$\gamma_i = \frac{h_{ii} p_i}{\sum_{j\neq i} h_{ji} p_j + \sigma^2} \tag{2}$$

[2]Some approaches incorporate the modulation scheme in the utility function [19].

where h_{ii} denotes the channel gain between the two terminals of link i, h_{ji} is the channel gain between the transmitter of link j and the receiver of link i, and σ^2 is the thermal noise power.

The second term in the utility function is a linear pricing function that represents the "price" for consuming a specific amount of power. For simplicity, the pricing factor α_i is set to a constant value for all links. Each terminal selects its transmission power such that its own utility function is maximized. This results in a standard non-cooperative game theory problem [6] of the following form:

$$\max_{p_i} \quad u_i(p_i, \mathbf{p}_{-i}), \text{for all } i = 1, 2, \ldots, n \tag{3}$$

subject to the constraints:

$$C_1 \quad : \quad p_i \in S_i \triangleq [0, P_{max}] \tag{4}$$

$$\text{and} \quad C_2 \quad : \quad \gamma_i \geq \text{SNR}_{th}. \tag{5}$$

We assume that all terminals use the same coding/ modulation scheme and that SNR_{th} and P_{max} are the same for all terminals. The solution to the above game, if feasible, is the one that achieves the NE. Previous work [19] states that the NE exists if the following two conditions are satisfied:

1) S_i is a nonempty and convex subset of some Euclidean space.
2) u_i is a continuous and quasi-concave function in p_i.

The first condition is readily satisfied. To show that the second condition is also satisfied, we take the second-order derivative of u_i:

$$\frac{\partial^2 u_i}{\partial p_i^2} = -\frac{h_{ii}^2}{(h_{ii}p_i + \sum_{j\neq i} h_{ji}p_j + \sigma^2)^2}. \tag{6}$$

Since $\frac{\partial^2 u_i}{\partial p_i^2} < 0$, the second condition is satisfied. Therefore, the NE exists.

To find the NE of the game, we construct and analyze the players' best response functions [15]. The best response of link i is the transmission power which maximizes its utility function and satisfies the constraints C_1 and C_2. The power that maximizes i's utility function can be obtained by equating the first-order derivative of u_i to zero, which yields:

$$\frac{\partial u_i}{\partial p_i} = \frac{h_{ii}}{h_{ii}p_i + \sum_{j\neq i} h_{ji}p_j + \sigma^2} - \alpha_i = 0. \tag{7}$$

Therefore,

$$p_i = \frac{1}{\alpha_i} - \frac{\sum_{j\neq i} h_{ji}p_j + \sigma^2}{h_{ii}}. \tag{8}$$

If the p_i computed above satisfies C_1 and C_2, then it represents the best response of link i.

We tentatively fix all the α_i's to a constant value, $1/P_{max}$. This guarantees that the computed p_i is upper bounded by P_{max}. By rearranging the terms in (8) and writing n simultaneous equations, (8) can be reduced into the following matrix-form expression:

$$\mathbf{HP}^* = \mathbf{G} \tag{9}$$

where $\mathbf{H} = [h_{ij}]_{i,j}$ is an $n \times n$ matrix representing the channel gains between each of the transmitter/receiver pairs; $\mathbf{G} = [g_1, g_2, ..., g_n]^T$ is an $n \times 1$ vector, where $g_i \triangleq \frac{h_{ii}}{\alpha_i} - \sigma^2$. If there exists a feasible and unique solution for (9), then $\mathbf{P}^* = [p_1^*, p_2^*, ..., p_n^*]$ is the unique NE solution, and can be calculated by:

$$\mathbf{P}^* = \mathbf{H}^{-1}\mathbf{G}. \tag{10}$$

To ensure the feasibility of P^*, the constraints C_1 and C_2 should both be satisfied. We use (8) to reformulate the constraint C_2 as a lower bound constraint on p_i as follows:

$$p_i \geq \frac{1}{\alpha} \frac{\text{SNR}_{th}}{1 + \text{SNR}_{th}}. \tag{11}$$

For example, if $\text{SNR}_{th} = 3dB$, we require $p_i \geq \frac{2}{3} P_{max}$. Therefore, the constraint on p_i becomes:

$$\frac{\text{SNR}_{th}}{1 + \text{SNR}_{th}} P_{max} \leq p_i \leq P_{max}. \tag{12}$$

If the computed p_i does not satisfy (12) for some i, the transmissions should not proceed concurrently at the computed powers $p_1^*, p_2^*, ..., p_n^*$. The power levels of some of these users need to be fixed to the boundary condition (either $\frac{SNR_{th}}{1+SNR_{th}} P_{max}$ or P_{max}), and re-computations are needed until the power levels of all the other users are feasible. There may exist multiple NE solutions in this case. As we will explain later in our protocol design, we handle this problem at the admission control phase. A terminal competing for the channel will compute its power level, and based on the feasibility of the computed power decides whether to transmit concurrently with the previously scheduled transmissions. This decision-making process is made serially by terminals during a contention period known as the *access window* (AW). Therefore, we can always get a unique feasible power solution after all the control packets are exchanged.

IV. Proposed GMAC Protocol

A. Assumptions

We assume that the channel gain is stationary for the transmission durations of a few control packets and one data packet. This assumption holds for typical mobility patterns and transmission rates. We also assume channel-gain reciprocity, i.e., the channel gain between two terminals is the same for both directions of the transmission. This is the underlying assumption in any RTS/CTS-based protocol, including the IEEE 802.11 standard. We focus on single-hop transmissions in this work.

B. Protocol Overview

Our proposed GMAC protocol has the following key features:

- Unlike the IEEE 802.11 approach, GMAC does not use RTS/CTS control packets to silence neighboring terminals. Instead, these packets are used to broadcast interference and channel-gain information that can be used by terminals to decide the feasibility of concurrent transmissions and their corresponding transmission power values.
- To assert their intentions to transmit, terminals exchange control packets over a certain time duration, referred to as the access window (AW). The AW allows several pairs of neighboring terminals to exchange their control packets so that data transmissions can proceed concurrently (see Fig. 2). Using an AW for contention was originally proposed in the MACA-P protocol [2] and was later integrated in the design of POWMAC [14]. However, we exploit the AW differently in GMAC, as will be explained in Section IV-C. The number of slots in the AW is adjusted dynamically according to the network load.
- Contending transmitter powers are not assigned until contention in the AW is over. This avoids the conservative presetting of tolerable interference (as in POWMAC) and results in better use of the network capacity.
- GMAC protects ACK packets by sending them sequentially at power P_{max} after data transmissions are completed. The order in which an ACK is transmitted corresponds to the order in which its terminals appear in the AW.

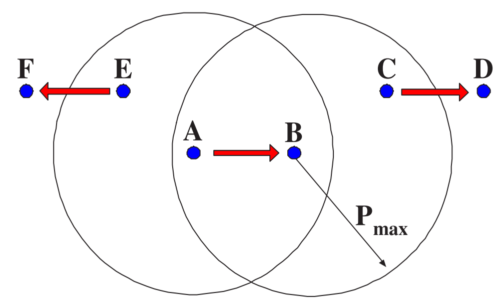

Fig. 2: An example of three concurrent transmissions.

GMAC allows several *clusters* (regions where multiple links contend for channel access) to be formed dynamically in a multi-hop network. The transmitters in the cluster will not be aware of their data packet transmission powers until being announced by the cluster head (see Section IV-C for more details). The purpose of such a design is to maximize the spatial throughput in each cluster in order to improve the throughput over the entire network.

C. Operational Details

In GMAC, a potential transmitter that senses a free channel and is not aware of any scheduled transmissions is considered a *master terminal*. The intended receiver of a master terminal is called a *master receiver*. We refer to the link that involves the master transmitter and receiver as the *master link*. All other communicating terminals in the same vicinity of a master link are *slave terminals*. The region formed by the transmission range of the master receiver constitutes a cluster in which the master receiver is its head. All terminals inside this cluster, other than the master terminal, are called *in-cluster* slave terminals. Such terminals can hear and correctly decode the master receiver's CTS packet. Terminals that are outside the master receiver's cluster but that are within the master transmitter's cluster are called *out-cluster* slave terminals.

Below, we describe the operation of each type of terminals and provide details of the GMAC protocol operation.

1) Master terminals: Consider a master terminal, say A, that has a data packet to transmit to another terminal, say B. If A does not sense a carrier for a random duration of time, it sends an RTS message at maximum power P_{max}, and includes in this RTS the remaining number of slots (N_{AW}) in the current AW (how A decides N_{AW} will be explained later). Upon receiving the RTS packet, receiver B uses the predetermined P_{max} value and the power of the received signal to estimate the channel gain h_{AB} (since we assume channel reciprocity, $h_{AB} = h_{BA}$). Terminal B then calculates the power p_A^* that maximizes the utility function in (1) assuming that only one transmission ($A \rightarrow B$) will take place in B's cluster. Accordingly,

$$p_A^* = \frac{1}{\alpha} - \frac{\sigma^2}{h_{AB}}. \quad (13)$$

If $p_A^* < p_{min} \triangleq \frac{\text{SNR}_{th}}{1+\text{SNR}_{th}} P_{max}$, then the existing interference at receiver B is too high and the transmission should not be allowed to proceed. In this case, terminal B will respond with a negative CTS (NCTS), informing A that it can not proceed with its transmission. On the other hand, if $p_A^* \geq p_{min}$, B will send back a positive CTS containing the channel gain h_{AB} and N_{AW}. It should be noted that the computed p_A^* at this point is not necessarily the transmission power of the data packet from terminal A. The final transmission power used by A will not be decided until the end of all negotiations in the AW.

Upon receiving B's CTS, terminal A replies back with a DTS (decide-to-send) packet that includes the channel gain h_{AB}. The DTS is needed to inform out-cluster slave terminals about the success of the RTS/CTS exchange between A and B. The 3-way (RTS/CTS/DTS) handshake was also used in POWMAC [14] and is depicted in Fig. 3.

2) In-cluster slave terminals: We now describe how in-cluster slave terminals operate after hearing the master terminals' RTS/CTS/DTS exchange. We use the scenario in Fig. 2 as an example.

Consider an in-cluster slave terminal, say C. Suppose that C overhears B's CTS packet and has a data packet to send. It first backs off for a random duration of time and if no carrier is sensed, it sends an RTS packet at power P_{max}. Terminal C's RTS will include h_{AB}, obtained from the previous RTS/CTS/DTS exchange between A and B. It will also include h_{CB} to be used by receiver D. When D receives the RTS, it first calculates h_{CD}. If D previously overheard A's DTS, it computes h_{AD}; otherwise, h_{AD} is set to zero (that means terminal D is out of the maximum transmission range of A). Terminal D then calculates the NE power vector $\mathbf{P}^* = [p_{AB}^* \; p_{CD}^*]$, assuming that two transmissions, $A \rightarrow B$ and $C \rightarrow D$, will take place simultaneously. If the computed powers are feasible, i.e., satisfy the feasibility condition in (12), D sends back a CTS that includes h_{CD} and h_{AD}. Otherwise, if either p_{AB}^* or p_{CD}^* is infeasible, D sends back a negative CTS. Upon receiving a CTS from D, terminal C sends a DTS that includes h_{CD} and h_{AD}. This channel gain information will be later used by the master receiver to compute the final transmission powers for all transmitters within its cluster. If more transmissions are to be scheduled following the RTS/CTS/DTS exchange between C and D, the same procedure is repeated.

In general, the RTS of any transmitter contains the channel gains between that transmitter and all previously-scheduled receivers in the same AW. Each receiver needs to make an admission decision by calculating the NE power vector $\mathbf{P}^*$, and will send back a CTS if a feasible solution exists. The CTS should contain all the channel gains between the current receiver and the transmitters of all previously scheduled transmissions. Finally, The transmitter sends a DTS that announces the channel gains included in the CTS to be used by the master receiver.

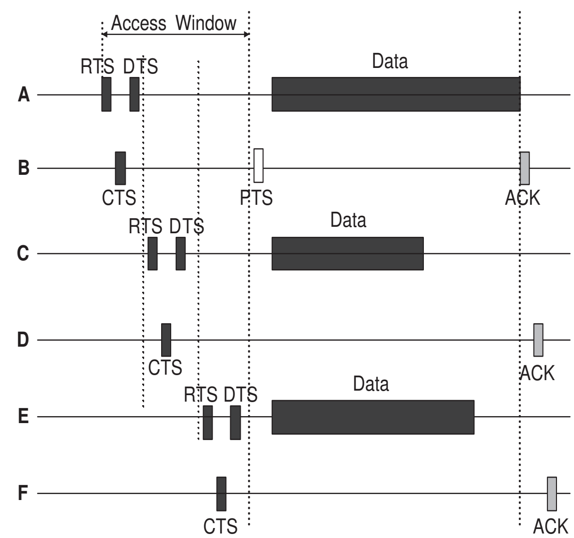

Fig. 3: Exchange of control and data packets in GMAC.

3) Final computation of transmission powers: After the master receiver (cluster head) receives information about all in-cluster transmissions, it computes the NE power values for all transmitters. The cluster head will then broadcast a power-to-send (PTS) packet, informing all transmitters of the power to be used for transmitting their data packets. The coordination of control and data packet transmissions is demonstrated in Fig. 3.

4) Out-cluster slave terminals: The data transmissions of out-cluster slave transmitters do not add significant interference at the master receiver. Therefore, the strategy that we propose for out-cluster slave terminals is different from the one for in-cluster slave terminals. An out-cluster transmitter, say E, only needs to increase its own utility without worrying about the interference it causes to the master receiver. Terminal E assumes that transmitters of all previously scheduled links will use the maximum transmission power (P_{max}) for their transmissions. It computes its transmission power accordingly.

Terminal E first sends an RTS packet if no carrier is sensed for a short duration. The RTS does not need to include the channel gain information as before. However, it needs to include A's address to the receiver to specify that A's transmission has been scheduled. Receiver F assumes that transmitter A will transmit at power P_{max} and the received power p_{AF} is considered background noise. Thus, the transmission power that E should use can be computed as follows:

$$p^*_{EF} = \frac{1}{\alpha} - \frac{\sigma^2 + h_{AF}P_{max}}{h_{EF}}. \quad (14)$$

If the computed p^*_{EF} is feasible, terminal F will send back a CTS. The power p^*_{EF} should be included in this CTS. After receiving this CTS, terminal E will send a DTS that includes p^*_{EF}. This p^*_{EF} is the final data transmission power of terminal E in this scenario.

5) ACK transmissions: Each terminal in the vicinity will have knowledge of the time the data transmissions are completed. ACK packets will be serially transmitted using P_{max} after all data transmissions are finished. As previously mentioned, the order of the ACK transmission will conform to the order of its terminals in the AW, so ACK transmissions do not interfere with each other.

6) Computing the AW size: Each terminal i will store an initial value of the AW. This value should be used if i acts as a master terminal. As described below, the AW size of i will be updated (increased or decreased) according to the load in its vicinity. The choice of AW size should aim at maximizing the number of concurrent transmissions, while not wasting the network resource.

After each data transmission duration, every terminal checks the number of concurrent transmissions in its vicinity. If this number is less than a pre-specified fraction (say $\delta\%$) of the AW size, then the current AW size is too large compared to the load. In this case, the AW size is decreased by one. Similarly, if this number is larger than $\delta\%$ of the current AW size, then AW size will be increased by one to allow for possibly more concurrent transmissions.

7) Terminals within two master clusters: In a multi-hop network, there may be cases where a single terminal, say C, resides within two different master clusters. This situation is exemplified in Fig. 4, where transmissions $A \to B$ and $A' \to B'$ are two master links that have been scheduled but have not started their data transmissions yet. The AWs of these two master receivers will not be synchronized. If terminal C has a packet to transmit, it has two options. One option is to keep silent and wait until these two master links finish their transmissions. Another option is to compute whether the misalignment between the two AWs is less than the maximum backoff window. If so, then terminal C can send its RTS and compete for the channel. We adopt the first option in this work and leave the other option in future work.

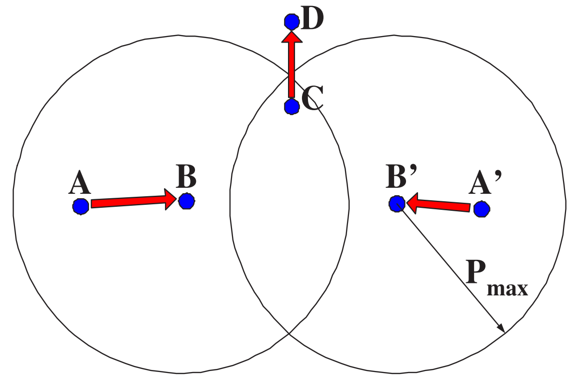

Fig. 4: A terminal (C) within 2 clusters.

TABLE I: Simulation parameters

Data Packet Size	2 KB
Data rate	1 Mbps
SINR threshold	6 dB
Maximum transmission range	750 meters
Maximum carrier-sense range	1500 meters
Path loss factor	4
Number of terminals	100

V. Performance Evaluation

A. Simulation Setup

We now evaluate the performance of the GMAC protocol and compare it with two single-channel protocols: POW-MAC [14] and IEEE 802.11 (ad-hoc mode). Our performance metrics are the network goodput (number of successful data transmissions per time unit), the number of concurrent transmissions, and the total energy consumption per delivered packet. We conduct simulations using the CSIM simulation package [11]. (CSIM is a C-based process-oriented discrete-event simulation package). The simulation parameters are provided in Table I. These parameters correspond to the Cisco Aironet Series 350 hardware specifications [5]. We use the two-ray propagation model with a path loss factor of 4.

The network area is divided into 10×10 equal-sized squares, one for each terminal. The location of each terminal is randomly assigned within its square. We conduct experiments for single-neighborhood and multi-neighborhood configurations, as described below.

B. Single-Neighborhood Network Configuration

In some scenarios, such as computers in a conference room, each terminal is within the maximum transmission range of all other terminals. In other words, all the exchanged control packets, which are sent using P_{max}, can be received by each terminal. We model this scenario in our simulations by limiting the field size to be $500m \times 500m$. Each terminal generates packets according to a Poisson process with rate λ. For each generated packet, the destination is randomly selected.

Fig. 5 depicts the network throughput versus the packet generation rate for the three examined protocols. This figure shows that GMAC achieves about 80% improvement in network goodput over the 802.11 standard, and about 40% improvement over POWMAC. Fig. 6 depicts a histogram of the number of concurrent transmissions (m) for both POWMAC and GMAC. It is clear that GMAC achieves more than twice the frequency of having $m \geq 2$ concurrent transmissions as POWMAC. This explains the goodput differences between the two protocols. Note that, in this configuration, the 802.11 standard only allows one transmission to proceed at one time.

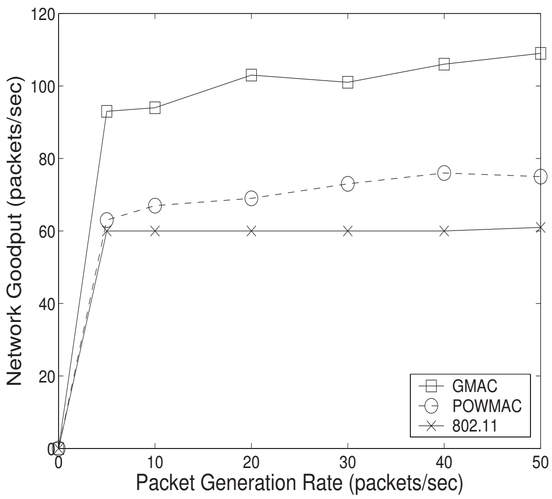

Fig. 5: Network goodput vs. load for a single-neighborhood network configuration.

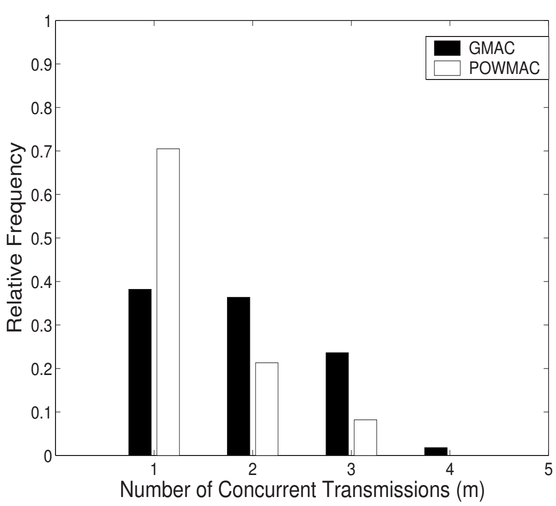

Fig. 6: Relative frequency of having m concurrent transmissions during network operation.

C. Multi-Neighborhood Network Configuration with a Fixed Field Size

We now examine a more general ad-hoc network scenario, whereby terminals can be out of range from each other, leading to hidden terminal problems. Specifically, we place 100 terminals within a square area of length 1500 meters. The square is split into 100 smaller squares, one for each terminal. The location of a terminal within each small square is randomly selected. For each generated packet, the destination is randomly selected from the one-hop neighbors of the source. Fig. 7 depicts the network goodput versus λ and shows that GMAC can achieve up to 70% increase in goodput over 802.11 and up to 25% increase over POWMAC. This improvement is due to the increase in the number of concurrent transmissions. Fig. 8 depicts the energy consumption versus λ for the three protocols. Energy consumption in this figure accounts for both data and control packet transmissions. It is clear that the energy consumption associated with GMAC is comparable to that of POWMAC and 802.11.

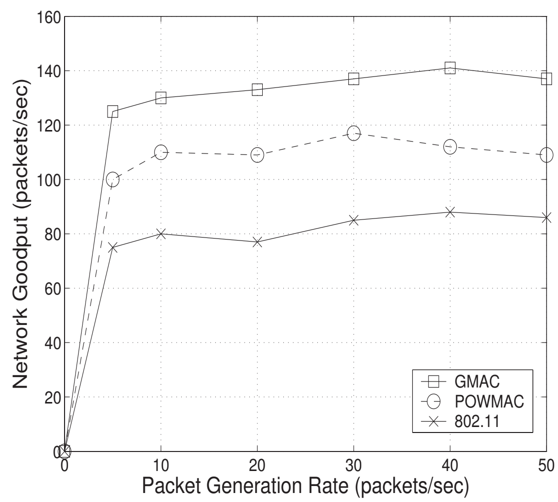

Fig. 7: Network goodput vs. load for a multi-neighborhood network configuration with a fixed field size.

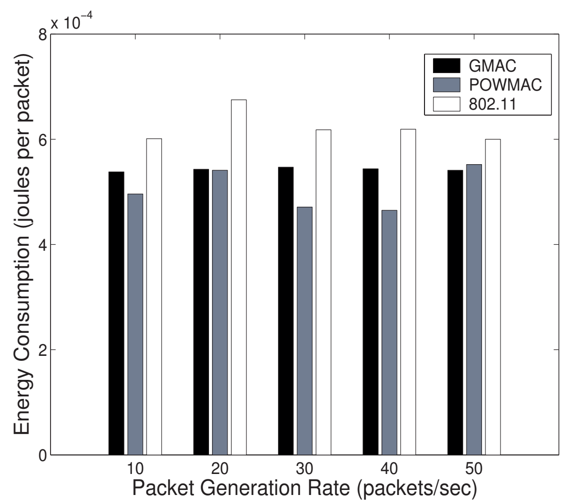

Fig. 8: Energy consumption vs. load for a multi-neighborhood network configuration with a fixed field size.

D. Multi-Neighborhood Network Configuration with a Variable Field Size

In this scenario, we vary the length of the square field while fixing the number of terminals (i.e., vary the node density). The packet generation rate is fixed at 40 packets/sec. The achieved goodput is shown in Fig. 9. In this scenario, GMAC shows consistent goodput improvement over both 802.11 scheme and POWMAC, especially under high densities. In the mean time, the energy consumption of GMAC is comparable to both the 802.11 scheme and POWMAC, as shown in Fig. 10.

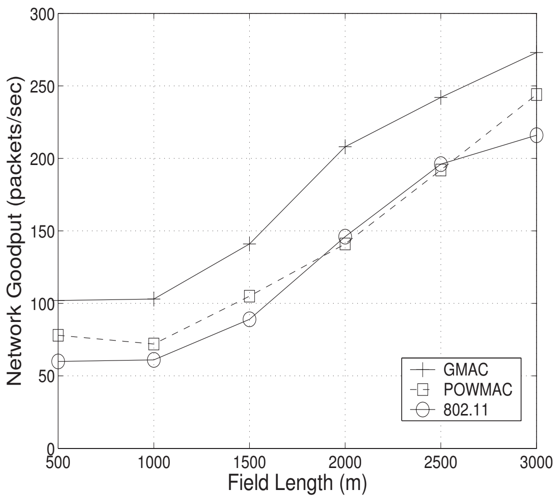

Fig. 9: Network goodput vs. field size for a network configuration with a fixed number of nodes.

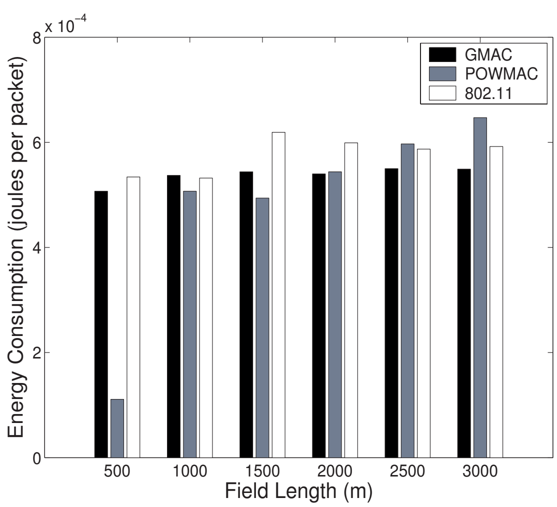

Fig. 10: Energy consumption vs. field size for a network configuration with a fixed number of nodes.

VI. Conclusions

In this paper, we proposed a game-theoretic power control MAC protocol (GMAC) for maximizing throughput in MANETs. GMAC uses a single channel for both data and control packet transmissions. It allows each user to determine whether or not it is feasible to transmit concurrently with previously scheduled transmissions. GMAC enables multiple transmissions to proceed concurrently by computing NE powers for all contending transmitters.

We compared the performance of GMAC with the IEEE 802.11 standard and the POWMAC scheme. Our simulation results show that GMAC significantly improves the network goodput over both schemes. In some scenarios, the network goodput in the case of GMAC was 80% (40%) larger than that of the 802.11 (POWMAC) scheme. GMAC also maintains comparable energy consumption levels.

For future extensions, we plan to optimize our protocol for multi-neighborhood networks. We will examine other pricing functions than linear pricing and incorporate other QoS requirements. We also plan to extend our protocol to support variable transmission rate for each link using similar game-theoretic techniques.

Acknowledgments

This work was supported by NSF under grants ANI-0095626, ANI-0313234, and ANI-0325979, and by the Center for Low Power Electronics (CLPE) at the University of Arizona. CLPE is supported by NSF (grant No. EEC-9523338), the State of Arizona, and a consortium of industrial partners. Any opinions, findings, conclusions, and recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

References

[1] http://www.ietf.org/html.charters/manet-charter.html.

[2] A. Acharya, A. Misra, and S. Bansal. MACA-P: A MAC for concurrent transmissions in multi-hop wireless networks. In *Proceeding of the First IEEE PerCom 2003 Conference*, volume 2, pages 59–66, March 2003.

[3] S. Agarwal, S. Krishnamurthy, R. Katz, and S. Dao. Distributed power control in ad-hoc wireless networks. In *Proceedings of the International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC)*, pages 59–66, 2001.

[4] T. Alpcan, T. Basar, R. Srikant, and E. Altman. CDMA uplink power control as a noncooperative game. *Wireless Networks*, 8:659–670, November 2002.

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

[6] D. Fudenberg and J. Tirole. *Game Theory*. The MIT Press, 1991.

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

[8] H. Ji and C.-Y. Huang. Non-cooperative uplink power control in cellular radio systems. *Wireless Networks*, 4(3):233–240, 1998.

[9] V. Kawadia and P. R. Kumar. Power control and clustering in ad hoc networks. April 2003.

[10] M. Krunz, A. Muqattash, and S.-J. Lee. Transmission power control in wireless ad hoc networks: Challenges, solutions, and open issues. *IEEE Network*, 18(5):8–14, September-October 2004.

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

[12] J. Monks, V. Bharghavan, and W.-M. Hwu. A power controlled multiple access protocol for wireless packet networks. pages 219–228, 2001.

[13] A. Muqattash and M. Krunz. A distributed transmission power control protocol for mobile ad hoc networks. *IEEE Transaction on Mobile Computing*, 3(2):113–128, April-June 2004.

[14] A. Muqattash and M. Krunz. POWMAC: A single-channel power-control protocol for throughput enhancement in wireless ad hoc networks. *IEEE Journal on Selected Areas in Communications*, 23(5):1067–1084, May 2005.

[15] M. J. Osborne. *An Introduction to Game Theory*. Oxford University Press, 2004.
[16] J. G. Proakis. *Digital Communication*. McGraw-Hill, Inc., 2001.
[17] R. Ramanathan and J. Redi. A brief overview of ad hoc networks: Challenges and directions. *IEEE Communications Magazine*, 40(5):20–22, May 2002.
[18] T. S. Rappaport. *Wireless Communications, Principles and Practice*. Prentice Hall, 1996.
[19] C. U. Saraydar, N. B. Mandayam, and D. J. Goodman. Efficient power control via pricing in wireless data networks. *IEEE Transactions on Communications*, 50(2):291 – 303, February 2002.
[20] 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.
[21] M. Xiao, N. B. Shroff, and E. K. P. Chong. Utility-based power control in cellular wireless systems. volume 1, pages 412 – 421, April 2001.