
Transmission Power Control in Wireless Ad Hoc Networks: Challenges, Solutions, and Open Issues

Marwan Krunz and Alaa Muqattash, The University of Arizona
Sung-Ju Lee, Hewlett-Packard Laboratories

Abstract

Recently, power control in mobile ad hoc networks has been the focus of extensive research. Its main objectives are to reduce the total energy consumed in packet delivery and/or increase network throughput by increasing the channel's spatial reuse. In this article we give an overview of various power control approaches that have been proposed in the literature. We discuss the factors that influence the selection of the transmission power, including the important interplay between the routing (network) and the medium access control (MAC) layers. Protocols that account for such interplay are presented.

Mobile ad hoc networks (MANETs) have recently been the topic of extensive research. The interest in such networks stems from their ability to provide temporary and instant wireless networking solutions in situations where cellular infrastructures are lacking and are expensive or infeasible to deploy (e.g., disaster relief efforts, battlefields, etc.). Due to their inherently distributed nature, MANETs are more robust than their cellular counterparts against single-point failures, and have the flexibility to reroute around congested nodes. Furthermore, MANETs can conserve battery energy by delivering a packet over a multihop path that consists of short hop-by-hop links. While wide-scale deployment of MANETs is yet to be realized, several efforts are currently underway to standardize protocols for the operation and management of such networks.

The ad hoc mode of the IEEE 802.11 standard is, by far, the most dominant MAC protocol for ad hoc networks. This protocol generally follows the CSMA/CA (carrier sense multiple access with collision avoidance) paradigm, with extensions to allow for the exchange of RTS/CTS (request-to-send/clear-to-send) handshake packets between the transmitter and the receiver. These control packets are used to reserve a *transmission floor* for the subsequent data and ack packets. 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 over.

Although the RTS/CTS exchange (also known as *virtual channel sensing*) is fundamentally needed to reduce the likelihood of collisions due to the hidden terminal problem,¹ it has two severe drawbacks. First, it negatively impacts channel utilization by not allowing concurrent transmissions to take place over the reserved floor. This situation is shown in Fig. 1, where node *A* uses its maximum transmission power to send its packets to node *B*. (For simplicity we assume omnidirectional antennas, so a node's reserved floor is represented by a

circle in the 2D space.) Nodes *C* and *D* hear *B*'s CTS message and, therefore, refrain from transmitting. It is easy to see that both transmissions, $A \rightarrow B$ and $C \rightarrow D$, can in principle take place at the same time if nodes are able to select their transmission powers appropriately. The second drawback of the fixed-power approach is that the received power may be far more than necessary to achieve the required signal-to-interference-and-noise ratio (SINR), thus wasting the node's energy and shortening its lifetime. Therefore, there is a need for a solution, possibly a multi-layer solution, that allows concurrent transmissions to take place in the same vicinity and simultaneously conserves energy.

The main objective of this article is to review the main approaches for transmission power control (TPC) that have been proposed in the literature. We start by discussing the tradeoffs involved in selecting the power level. A class of energy-oriented power control schemes is then discussed. This class is mainly aimed at reducing energy consumption, with throughput being a secondary factor. It includes network-layer solutions (i.e., power-aware routing). Power control schemes that incorporate the MAC perspective into their design are then presented. These schemes include a class of algorithms that use TPC primarily to control the topological properties of the network. In the same section we also discuss a class of interference-aware TPC schemes that use broadcasted interference information to bound the power levels of subsequent transmissions. Other protocols that are based on clustering or that combine scheduling and TPC are presented. The article concludes with a discussion of open research issues.

Tradeoffs in Selecting the Transmission Power

The transmission power determines the range over which the signal can be coherently received, and is therefore crucial in determining the performance of the network (throughput, delay, and energy consumption). The selection of the "best" transmission range has been investigated extensively in the literature. It has been shown that a higher network capacity can




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

be achieved by transmitting packets to the nearest neighbor in the forward progress direction.

The intuition behind this result is that halving the transmission range increases the number of hops by two but decreases the area of the reserved floor to one fourth of its original value, hence allowing for more concurrent transmissions to take place in the same neighborhood.

In addition to improving network throughput, reducing the transmission range plays a significant role in reducing the energy required to deliver a packet in a multihop fashion. The power consumed by the radio frequency (RF) power amplifier of the network interface card (NIC) is directly proportional to the power of the transmitted signal, and thus it is of great interest to control the signal transmission power to increase the lifetime of mobile nodes. Presently, the RF power amplifier consumes almost half (or more in the case of sensor nodes) of the total energy consumed by the NIC. This ratio is expected to increase in future NICs, as the processing components become more power-efficient. Therefore, there is potential for a significant energy saving by reducing the signal transmission power (range) and increasing the number of hops to the destination.

On the other hand, the transmission power determines who can hear the signal, so reducing it can adversely impact the connectivity of the network by reducing the number of active links and, potentially, partitioning the network (see the example in Fig. 2). Thus, to maintain connectivity, power control should be carried out while accounting for its impact on network topology. Furthermore, since route discovery in MANETs is often *reactive* (i.e., the path is acquired on demand), power control can be used to influence the decisions made at the routing layer by controlling the power of the *route-request* (RREQ) packets (discussed in more detail in a later section).

The above discussion provides sufficient motivation to dynamically adjust the transmission power for data packets. However, there are many open questions at this point, perhaps the most interesting being whether TPC is a network-layer or MAC-layer issue. The interaction between the network layer and MAC layer is fundamental to power control in MANETs. On the one hand, the power level determines who can hear the transmission, and hence directly impacts the selection of the next hop. Obviously this is a net-

work-layer issue. On the other hand the power level also determines the floor that the node reserves exclusively for its transmission through an access scheme. Obviously this is a MAC-layer issue. Hence we have to introduce power control from the perspectives of both layers. Other important questions are: How can a node find an energy-efficient route to the destination? What are the implications of adjusting the transmission powers of data and control packets? How can multiple transmissions take place simultaneously in the same vicinity? We address these questions in the subsequent sections.

Energy-Oriented Power Control Approaches

In this section we present power control approaches that aim at reducing energy consumption of nodes and prolonging the lifetime of the network. Throughput and delay are secondary objectives in such approaches.

TPC for Data Packets Only

One possible way to reduce energy consumption is for the communicating nodes to exchange their RTS/CTS packets at maximum power (P_{max}), but send their DATA/ACK packets at the minimum power (P_{min}) needed for reliable communication. The value of P_{min} is determined based on the receiver's power sensitivity, the SINR threshold, the interference level at the receiver, the antenna configuration (omni or directional), and the channel gain between the transmitter and the receiver. We refer to this basic protocol as SIMPLE. Note that SIMPLE and the IEEE 802.11 scheme have the same *forward progress rate per hop*, that is, the distance traversed by a packet in the direction of the destination is the same for both protocols. Thus, the two protocols achieve comparable throughputs. However, energy consumption in SIMPLE is expectedly less. The problem with SIMPLE, however, is when a min-hop routing protocol (MHRP) (which is the *de facto* routing approach in MANETs) is used at the network layer. In selecting the next hop (NH), a MHRP favors nodes in the direction of the destination that are *farthest* from the source node, but still within its maximum transmission range. When network density is high the distance between the source node and the NH is very close to the maximum transmission range; thus, SIMPLE would be preserving very little energy. The problem lies in the poor selection of the NH (i.e., links are long), and so a more "intelligent" routing protocol that finds an energy-efficient route to the destination is required. In other words, *for SIMPLE to provide good energy saving, a power-aware protocol on top of SIMPLE is needed*, which is the topic of the next section.





Figure 2. *Effect of power level on network connectivity: (a) low transmission power; (b) high transmission power.*



■ Figure 3. Drawbacks of the PARP/SIMPLE approach. Nodes E and D have to defer their transmissions when the data packets from A to B are routed via node C.

Power-Aware Routing Protocols (PARPs)

The first generation of routing protocols for MANETs [1] are essentially MHRPs that do not consider power efficiency as the main goal. Several recent routing protocols propose energy-efficient schemes. Singh *et al.* [2] first raised the power-awareness issue in ad hoc routing and introduced new metrics for path selection, which include the energy consumed per packet, network connectivity duration (i.e., the time before network partitions), node power variance, cost per packet, and maximum node cost. PARPs discussed in the remainder of this section use one or more of these metrics in path selection.

The first wave of PARPs was based on proactive shortest path algorithms. Instead of using delay or hop count as the link weight, these protocols use energy-related metrics such as signal strength, battery level at each node, and power consumption per transmission. The link condition and power status of each node are obtained via a periodic route table exchange, as is done in proactive routing protocols. It has been argued that the sole minimization of the total consumed energy per end-to-end packet delivery drains out the power of certain nodes in the network. Instead, energy consumption must be balanced among nodes to increase network lifetime.

Proactivity implies that each node must periodically exchange local routing and power information with neighboring nodes, which incurs significant control overhead. For this reason, proactive shortest path algorithms are mainly suitable for networks with little (or no) mobility, such as sensor networks. These schemes are shown to consume more power than on-demand routing protocols, as transmitting more control packets results in more energy consumption. Power-Aware Routing Optimization (PARO) [3] also utilizes power consumption as the route metric, but it is an on-demand protocol and, therefore, does not have the problems associated with proactive routing in MANETs. However, as its sole focus is on minimizing the transmission power consumed in the network, it does not account for balancing the energy consumption among nodes.

In [4] the authors proposed a scheme to conserve energy and increase network lifetime based on the use of directional antennas. This scheme first builds “minimum energy consumed per packet” routes using Dijkstra-like algorithms, and then schedules node transmissions by executing a series of

maximum weight matchings. The scheme is shown to be energy-efficient when compared with shortest-path routing under omni-directional antennas. However, since each node is assumed to have a single-beam directional antenna, the sender and the receiver must redirect their antenna beams toward each other before transmission and reception can take place. Moreover, it is preferred that each node participate in only one session at a time, as redirecting antennas requires a large amount of energy. These restrictions cause large delays, and hence the scheme is not adequate for time-sensitive data transmission.

Limitations of the PARP/SIMPLE Approach

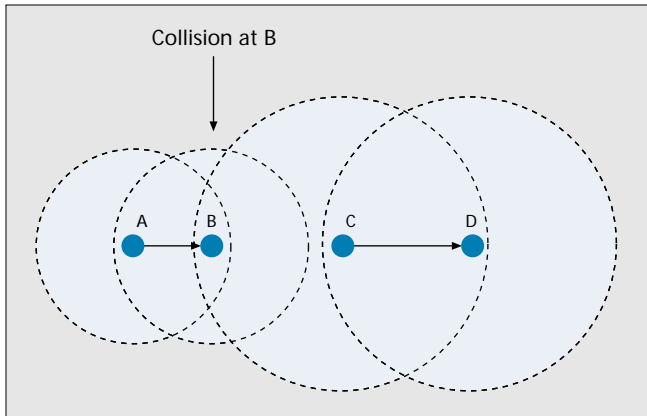
In the previous section we showed how a PARP/SIMPLE combination can significantly reduce energy consumption in a MANET. This reduction, however, comes at the expense of a decrease in network throughput and an increase in packet delays. To illustrate these drawbacks consider the example in Fig. 3. Nodes A, B, and C are within each other’s maximum transmission range. Node A wants to send packets to node B. According to a MHRP/802.11 solution, node A sends its packets directly to B. Thus, nodes E and D, who are unaware of the transmission $A \rightarrow B$, are able to communicate concurrently. On the other hand, according to a PARP/SIMPLE approach, data packets from A to B must be routed via node C, and thus, nodes E and D have to defer their transmissions for two data packet transmission periods. More generally, all nodes within C’s range but outside B’s or A’s range are not allowed to transmit, for they are first silenced by C’s CTS to A, and then again by C’s RTS to B. This shows that a PARP/SIMPLE approach forces more nodes to defer their transmissions, resulting in lower network throughput than that of the MHRP/802.11 approach.

TPC: The MAC Perspective

The throughput degradation in PARP/SIMPLE has to do with the fixed-power *exclusive-reservation* mechanism at the MAC layer. Hence it is natural to consider a medium access solution that allows for the adjustment of the reserved floor depending on the *data* transmission power. A power controlled MAC protocol reserves different floors for different packet destinations. In such a protocol both the channel bandwidth and the reserved floor constitute network resources for which nodes contend. For systems with a shared data channel (i.e., one node uses all the bandwidth for transmission) the floor becomes the single critical resource. This is in contrast to cellular systems and the IEEE 802.11 scheme, where the reserved floor is always fixed.

Topology Control Algorithms

We now present a family of protocols that use TPC as a means of controlling network topology (e.g., reducing node degree while maintaining a connected network). The size of the reserved floor in these protocols varies in time and among nodes, depending on the network topology. In [5] the authors proposed a distributed position-based topology control algorithm that consists of two phases. Phase one is used for link setup and configuration, and is performed as follows. Each node broadcasts its position to its neighbors and uses the position information of its neighbors to build a sparse graph called the *enclosure graph*. In phase two, nodes find the “optimal” links on the enclosure graph by applying the distributed Bellman-Ford shortest path algorithm with power consumption as the cost metric. Each node i broadcasts its cost to its neighbors, where the cost of node i is defined as the minimum power necessary for i to establish a path to a destination. The



■ Figure 4. *Challenge in implementing power control in a distributed fashion. Node C is unaware of the ongoing transmission $A \rightarrow B$, and hence it starts transmitting to node D at a power that destroys B's reception.*

protocol requires nodes to be equipped with GPS receivers. In [6] a cone-based solution that guarantees network connectivity was proposed. Each node i gradually increases its transmission power until it finds at least one neighbor in every cone of angle $\alpha = 2\pi/3$ centered at i (a $5\pi/6$ angle was later proven to guarantee network connectivity). Node i starts the algorithm by broadcasting a "Hello" message at low transmission power and collecting replies. It gradually increases the transmission power to discover more neighbors and continuously caches the direction in which replies are received. It then checks whether each cone of angle α contains a node. The protocol assumes the availability of directional information (angle-of-arrival), which requires extra hardware. Some researchers proposed the use of a synchronized global signaling channel to build a global network topology database, where each node communicates only with its nearest N neighbors (N is a design parameter). This approach, however, requires a signaling channel in which each node is assigned a dedicated slot.

One common limitation of the above protocols is their sole reliance on CSMA for accessing/reserving the shared wireless channel. It is known that using CSMA alone for accessing the channel can significantly degrade network performance (throughput, delay, and power consumption) because of the well known hidden terminal problem. Unfortunately, this problem cannot be overcome using a standard RTS/CTS-like channel reservation approach, as explained in the example in Fig. 4. Here, node A has just started a transmission to node B at a power level that is just enough to ensure coherent reception at B . Suppose that node B uses the same power level to communicate with A . Nodes C and D are outside the floors of A and B , so they do not hear the RTS/CTS exchange between A and B . (For simplicity we assume in this example that the carrier-sensing and the reception ranges are the same.) For nodes C and D to be able to communicate they have to use a power level that is reflected by the transmission floors in Fig. 4 (the two circles centered at C and D). However, the transmission $C \rightarrow D$ will interfere with transmission $A \rightarrow B$, causing a collision at B . In essence, the problem is caused by the asymmetry in the transmission floors (i.e., B can hear C 's transmission to D but C cannot hear B 's transmission to A).

Interference-aware MAC Protocols

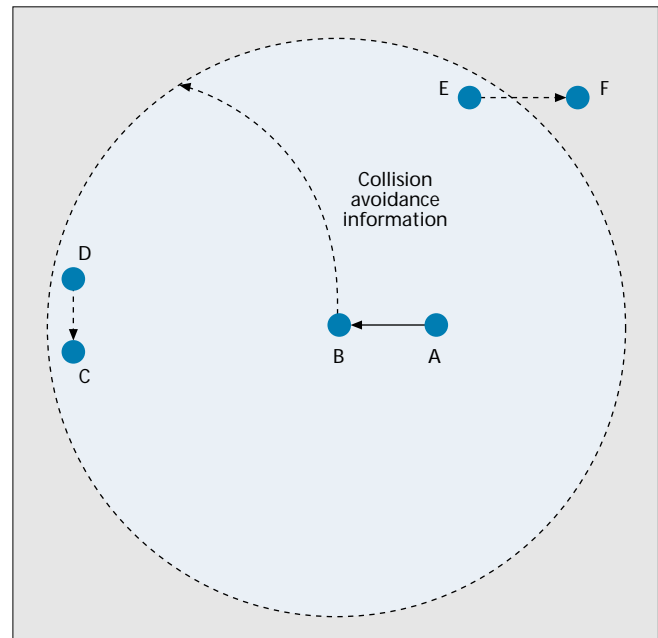
Topology control protocols discussed above lack a proper channel reservation mechanism (e.g., RTS/CTS like), which negatively impacts the achievable throughput under these protocols. To address this issue more sophisticated MAC protocols are needed, in which information about an ongoing transmission is made known to all possible interferers. Figure

5 illustrates the intuition behind such protocols. Node A intends to send its data to B . Before this transmission can take place, node B broadcasts some "collision avoidance information" (CAI) to all possible interfering neighbors, which include C , D , and E . Unlike the RTS/CTS packets used in the 802.11 scheme, this CAI does not prevent interfering nodes from accessing the channel. Instead, it *bounds the transmission powers* of future packets generated by these nodes. Thus, in Fig. 5 future transmitters (D and E in this example) can proceed only if the powers of their signals are not high enough to collide with the ongoing reception at node B .

To understand what this CAI is and how nodes can make use of it, consider the transmission of a packet from some node i to some node j . Let $SINR(i, j)$ be the signal-to-interference-and-noise ratio at node j for the desired signal from node i . Then,

$$SINR(i, j) = P(i, j) / \left(\sum_{k \neq i} P(k, j) + \eta_j \right),$$

where $P(i, j)$ is the received power at node j for a transmission from node i and η_j is the thermal noise at node j . A packet is correctly received if the SINR is above a certain threshold (say, $SINR_{th}$) that reflects the QoS of the link. By allowing nearby nodes to transmit concurrently, the interference power at receiver j increases, and so $SINR(i, j)$ decreases. Therefore, to be able to correctly receive the intended packet at node j , the transmission power at node i must be computed while taking into account potential future transmissions in the neighborhood of receiver j . This is achieved by incorporating an interference margin in the computation of $SINR(i, j)$. This margin represents the additional interference power that receiver j can tolerate while ensuring coherent reception of the upcoming packet from node i . Nodes at some interfering distance from j can now start new transmissions while the transmission $i \rightarrow j$ is taking place. The interference margin is incorporated by scaling up the transmission power at node i beyond what is minimally needed to overcome the current interference at node j . Due to the distributed nature of the



■ Figure 5. *Broadcasting collision avoidance information in interference-aware MAC protocols.*

TPC problem it makes sense that the computation of the appropriate transmission power level is made by the *intended receiver*, which is more capable of determining the potential interferers in its neighborhood than the transmitter. Note that the power level is determined for *each* data packet separately (possibly via an RTS/CTS handshake) just before the transmission of that packet. This is in contrast to cellular networks in which the power is determined not only at the start of the transmission but also while the packet is being transmitted (e.g., the transmission power is updated every 125 μ s in the IS-95 standard for cellular systems).

Now a node with a packet to transmit is allowed to proceed with its transmission if the transmission power will not disturb the ongoing receptions in the node's neighborhood *beyond the allowed interference margin*. Allowing for concurrent transmissions increases network throughput and decreases contention delay.

Proposed interference-aware MAC protocols differ mainly in how they compute the CAI and how they distribute it to neighboring nodes. In [7] the authors proposed the *power controlled multiple access* (PCMA) protocol, in which each receiver sends busy-tone pulses to advertise its interference margin. The signal strength of the received pulses is used to bound the transmission power of the (interfering) neighboring nodes. A potential transmitter i first senses the busy-tone channel to determine an upper bound on its transmission power for all of its control and data packets, adhering to the most sensitive receiver in its neighborhood. After that node i sends its RTS at the determined upper bound and waits for a CTS. If the receiver, say j , is within the RTS range of node i , and the power needed to send back the CTS is below the power bound at j , node j sends back a CTS allowing the transmission to begin. The simulation results in [7] show significant throughput gain (more than twice) over the 802.11 scheme. However, the choice of energy-efficient links is left to the upper layer (e.g., a PARP). Furthermore, the interference margin is fixed and it is not clear how it can be determined. Contention among busy-tones is also not addressed. Finally, according to PCMA a node may send many RTS packets without getting any reply, thus wasting the node's energy and the channel bandwidth.

The use of a separate control channel in conjunction with a busy-tone scheme was proposed in [8]. The sender transmits data packets and busy-tones at reduced power, while the receiver transmits its busy-tones at the maximum power. A node estimates the channel gain from the busy-tones and is allowed to transmit if its transmission is not expected to add more than a fixed interference to the ongoing receptions. The protocol is shown to achieve considerable throughput improvement over the original *dual busy-tone multiple access* (DBTMA) protocol. The authors, however, make strong assumptions about 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) and that there is no need for any interference margin. Also, the power consumption of the busy-tones was not addressed. Furthermore, as in PCMA the choice of energy-efficient links is left to the upper layer.

The power controlled dual channel (PCDC) protocol [9] emphasizes the interplay between the MAC and network layers, whereby the MAC layer indirectly influences the selection of the next-hop by properly adjusting the power of the RREQ packets. According to PCDC the available bandwidth is divided into two frequency-separated channels for data and control. Each data packet is sent at a power level that accounts for a receiver-dependent interference margin. This margin allows for concurrent transmissions to take place in the neigh-

borhood of the receiver, provided that these transmissions do not individually interfere with the ongoing reception by more than a fraction of the total interference margin. The CAI is inserted into the CTS packet, which is sent at maximum power over the control channel, thus informing all possible interferers about the ensuing data packet and allowing for interference-limited simultaneous transmissions to take place in the neighborhood of a receiving node. Furthermore, each node continuously caches the estimated channel gain and angle of arrival of every signal it receives over the control channel, regardless of the intended destination of this signal. This information is used to construct an energy-efficient subset of neighboring nodes, called the *connectivity set* (CS). The intuition behind the algorithm is that the CS must contain only neighboring nodes with which direct communication requires less power than the indirect (two-hop) communication via any other node that is already in the CS. Let $P_{\text{conn}}^{(i)}$ denote the minimum power required for node i to reach the *farthest* node in its CS. Node i uses this power level to broadcast its RREQ packets. This results in two significant improvements. First, any simple MHRP can now be used to produce routes that are very power efficient and that increase network throughput (i.e., reduce the total reserved floor). Hence, no intelligence is needed at the network layer and no link information (e.g., power) has to be exchanged or included in the RREQ packets in order to find power-efficient routes. Clearly, this reduces complexity and overhead. Second, considering how RREQ packets are flooded throughout the network, significant improvements in throughput and power consumption can be achieved by limiting the broadcasting of these packets to nodes that are within the connectivity range $P_{\text{conn}}^{(i)}$. It was shown in [9] that if the network is connected under a fixed-power strategy (i.e., RREQ packets are broadcasted using power P_{max}), then it must also be connected under a CS-based strategy.

PCDC was shown to achieve considerable throughput improvement over the 802.11 scheme *and* significant reduction in energy consumption. The authors, however, did not account for the processing and reception powers, which increase with the number of hops along the path (note that PCDC results in longer paths than the 802.11 scheme when both are implemented below a MHRP). Furthermore, there is an additional signaling overhead in PCDC due to the introduction of new fields in the RTS and CTS packets.

Other TPC Approaches

In this section we describe two additional TPC approaches that adopt completely different philosophies to the problem than what has been discussed so far.

The first approach is clustering [10] in which an elected cluster head (CH) performs the function of a base station in a cellular system. It uses closed-loop power control to adjust the transmission powers of nodes in the cluster. Communications between different clusters occur via gateways, which are nodes that belong to more than one cluster. This approach simplifies the forwarding function for most nodes, but at the expense of reducing network utilization since all communications have to go through the CHs. This can also lead to the creation of bottlenecks. A joint clustering/TPC protocol was proposed in [11], in which each node runs several routing-layer agents that correspond to different power levels. These agents build their own routing tables by communicating with their peer routing agents at other nodes (i.e., the protocol is distributed with no CHs). Each node along the packet route determines the lowest-power routing table in which the destination is reachable. The routing overhead in this protocol grows in proportion to

the number of routing agents, and can be significant even for simple mobility patterns. (Recall that for DSR, RREQ packets account for a large fraction of the total received bytes.)

Another novel approach for TPC is based on joint scheduling and power control [12], and consists of scheduling and power control phases. The purpose of the scheduling phase is to eliminate strong interference that cannot be overcome by TPC. It also makes the TPC problem similar to that of cellular systems. In the scheduling phase the algorithm searches for the largest subset of nodes that satisfy "valid scenario constraints." A node satisfies such constraints if it does not transmit and receive simultaneously, it does not receive from more than one neighbor at the same time, and when receiving from a neighbor the node is spatially separated from other interferers by at least a distance D . This D is set to the "frequency reuse distance" parameter used in cellular systems. In the TPC phase the algorithm searches for the largest subset of users generated from the first phase that satisfy admissibility (SINR) constraints. The complexity of both phases is exponential in the number of nodes. Because the algorithm is invoked on a slot-by-slot basis, it is computationally expensive for real-time operation. The authors in [12] proposed heuristics to reduce the computational burden. A simple heuristic for the scheduling phase is to examine the set of valid scenarios *sequentially* and defer transmissions accordingly. There is still a need for a centralized controller to execute the scheduling algorithm (i.e., the solution is not fully distributed). For the TPC phase the authors examined a cellular-like solution that involves deferring the user with the minimum SINR in an attempt to lower the level of multiple access interference. It is assumed here that the measured SINR at each receiver is known to *all* transmitters (e.g., via flooding). The case of TPC for multicast transmission was addressed in [13], where the authors proposed a distributed joint scheduling and power control scheme for multicast transmissions.

Summary and Open Issues

TPC has great potential to improve the throughput performance of a MANET and simultaneously decrease energy consumption. In this article we surveyed several TPC approaches. Some of these approaches (e.g., PARP/SIMPLE) are successful in achieving the second goal, but sometimes at the expense of a reduction (or at least no improvement) in throughput performance. By locally broadcasting "collision avoidance information" some protocols are able to achieve both goals of TPC simultaneously. These protocols, however, are designed based on assumptions (e.g., channel stationarity and reciprocity) that are valid only for certain ranges of speeds and packet sizes. Furthermore, they generally require additional hardware support (e.g., duplexers). The key message in the design of efficient TPC schemes is to account for the interplay between the routing layer (network) and the MAC layer.

Many interesting open problems remain to be addressed. Interference-aware TPC schemes are promising, but their feasibility and design assumptions need to be evaluated. For instance, PCDC assumes that the channel gain is the same for the control channel and the data channel. This holds only when the control channel is within the coherence bandwidth of the data channel, which places an upper bound on the allowable frequency separation between the two channels. Ideally, one would like to have a single-channel solution for the TPC problem. Interoperability with existing standards and hardware is another important issue. Currently most wireless devices implement the IEEE 802.11b standard. TPC schemes proposed in the literature are often not backward-compatible with the IEEE 802.11 standard, which makes it difficult to

deploy such schemes in real networks. Another important issue is the incorporation of a sleep mode in the design of TPC protocols. A significant amount of energy is consumed by unintended receivers. In many cases it makes sense to turn off the radio interfaces of some of these receivers to prolong their battery lives. The effect of this on the TPC design has not been explored.

The schemes presented in this article assume that nodes are equipped with omnidirectional antennas. Directional antennas have recently been proposed as a means of increasing network capacity under a fixed-power strategy (e.g., [14]). The use of TPC in MANETs with directional antennas can provide significant energy saving. However, the access problem is now more difficult due to the resurfacing of various problems such as the hidden terminal, deafness, etc., which need to be addressed. Power control for CDMA-based MANETs is another interesting topic that has not received enough attention. 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 consider the use of CDMA in MANETs. The situation, however, is more complicated in the case of MANETs due to the presence of nonnegligible cross-correlations between different CDMA codes, which can induce multi-access interference at receivers and cause "secondary" packet collisions (collisions between two or more transmissions that use different CDMA codes). This problem, known in the literature as the *near-far problem*, is both an access problem and a TPC problem. An initial attempt at addressing this combined problem is given in [15], but more work is still needed to better understand the capacity of a CDMA-based MANET and the optimal design of TPC for such a network.

Variable-rate support is another optimization that TPC protocols have not considered yet. It is known that adapting the transmit power, data rate, and coding scheme increases spectral efficiency. The IEEE 802.11b scheme allows nodes to increase their information rate up to 11 Mb/s, depending on the SINR at the receiver. The performance achieved through TPC can be further improved by allowing for dynamic adjustment of the information rate, increasing this rate when the interference is low and vice versa. The "mechanics" of such an approach are yet to be explored.

References

- [1] E. M. Royer and C.-K. Toh, "A Review of Current Routing Protocols for Ad Hoc Mobile Wireless Networks," *IEEE Pers. Commun. Mag.*, vol. 6, no. 2, Apr. 1999, pp. 46–55.
- [2] S. Singh, M. Woo, and C. S. Raghavendra, "Power Aware Routing in Mobile Ad Hoc Networks," *Proc. ACM MobiCom Conf.*, 1998, pp. 181–90.
- [3] J. Gomez *et al.*, "PARO: Supporting Dynamic Power Controlled Routing in Wireless Ad Hoc Networks," *ACM/Kluwer J. Wireless Networks*, vol. 9, no. 5, 2003, pp. 443–60.
- [4] A. Spyropoulos and C. S. Raghavendra, "Energy Efficient Communications in Ad Hoc Networks using Directional Antennas," *Proc. IEEE INFOCOM Conf.*, Apr. 2003.
- [5] V. Rodoplu and T. Meng, "Minimum Energy Mobile Wireless Networks," *IEEE JSAC*, vol. 17, no. 8, Aug. 1999, pp. 1333–44.
- [6] R. Wattenhofer *et al.*, "Distributed Topology Control for Power Efficient Operation in Multihop Wireless Ad Hoc Networks," *Proc. IEEE INFOCOM Conf.*, 2001, pp. 1388–97.
- [7] J. Monks, V. Bharghavan, and W.-M. Hwu, "A Power Controlled Multiple Access Protocol for Wireless Packet Networks," *Proc. IEEE INFOCOM Conf.*, 2001, pp. 219–28.
- [8] 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 JSAC*, vol. 18, no. 9, 2000, pp. 1647–57.
- [9] M. Muqattash and M. Krunz, "Power Controlled Dual Channel (PCDC) Medium Access Protocol for Wireless Ad Hoc Networks," *Proc. IEEE INFOCOM Conf.*, 2003, pp. 470–80.
- [10] T.-J. Kwon and M. Gerla, "Clustering with Power Control," *Proc. IEEE MILCOM Conf.*, 1999, pp. 1424–28.

-
- [11] V. Kawadia and P. R. Kumar, "Power Control and Clustering in Ad Hoc Networks," *Proc. IEEE INFOCOM Conf.*, 2003, pp. 459–69.
 - [12] T. El Batt and A. Ephremides, "Joint Scheduling and Power Control for Wireless Ad Hoc Networks," *Proc. IEEE INFOCOM Conf.*, 2002, pp. 976–84.
 - [13] K. Wang *et al.*, "A Distributed Joint Scheduling and Power Control Algorithm for Multicasting in Wireless Ad Hoc Networks," *Proc. IEEE ICC Conf.*, vol. 1, 2003, pp. 725–31.
 - [14] R. R. Choudhury *et al.*, "Using Directional Antennas for Medium Access Control in Ad Hoc Networks," *Proc. ACM MobiCom Conf.*, 2002, pp. 59–70.
 - [15] A. Muqattash and M. Krunz, "CDMA-based MAC Protocol for Wireless Ad Hoc Networks," *Proc. ACM MobiHoc Conf.*, June 2003.

Biographies

MARWAN KRUNZ [SM'04] (Krunz@ece.arizona.edu) is an associate professor of electrical and computer engineering at the University of Arizona. He received the Ph.D. degree in electrical engineering from Michigan State University in 1995. From 1995 to 1997 he was a postdoctoral research associate with the department of computer science, University of Maryland, College Park. He also held visiting research positions at INRIA, Sophia Antipolis, France; HP Labs, Palo Alto; and US West Advanced Technologies, Boulder, Colorado. His recent research interests include medium access and routing protocols for mobile ad hoc networks, quality of service provisioning over wireless links, constraint-based routing, WWW traffic modeling, and media streaming. He has published more than 80 journal articles and refereed conference papers in these areas. He received the National Science Foundation CAREER Award (1998-2002). He currently serves on the editorial board for the *IEEE/ACM Transactions on Networking* and the *Computer Communications Journal*. He was a guest co-editor for special issues in *IEEE Micro* and *IEEE Communications Magazines*. He served as the technical program co-chair for the IEEE INFOCOM 2004 Conference and the 2001 Hot Interconnects Symposium (Stanford University, August 2001). He has served and continues to serve on the executive and technical program committees of several international conferences. He consults for a number of corporations in the telecommunications industry.

ALAA MUQATTASH (Alaa@ece.arizona.edu) is a Ph.D. candidate and a research assistant in the department of electrical and computer engineering at the University of Arizona. He received the B.S.E.E. degree from The University of Jordan in 2000, and the M.S.E.C.E. degree from the University of Arizona in 2002. His current research interests are in system architecture and communication protocols for wireless ad hoc networks with emphasize on transmission power control.

SUNG-JU LEE (sjlee@hpl.hp.com) is a research scientist at the Mobile & Media Systems Lab (MMSL) of HP Labs. Lee received the Ph.D. in computer science from the University of California, Los Angeles. He has published more than 35 papers in the field of mobile networking and ad hoc networks. He is serving on the editorial board of several journals and is a technical program committee member for leading networking conferences. His research interests include mobile networking and computing, ad hoc networks, energy-efficient protocols, streaming media distribution networks, public WLAN hotspot networks, overlay networks, and adaptive service infrastructure. He is a member of IEEE and ACM.