

Conserving Transmission Power in Wireless Ad Hoc Networks

Javier Gomez, Andrew T. Campbell
Dept. of Electrical Engineering and
Center for Telecommunications Research
500 W. 120th Str., Mail Code 4712
Columbia University, NY 10027, USA
{javierg,campbell}@comet.columbia.edu

Mahmoud Naghshineh, Chatschik Bisdikian
IBM T.J. Watson Research Center
30 Saw Mill River Road
Hawthorne, NY 10953, USA
{mahmoud,bisdik}@us.ibm.com

Abstract

This paper introduces PARO, a power-aware routing optimization that helps to minimize the transmission power needed to forward packets between wireless devices in ad hoc networks. Using PARO, one or more intermediate nodes called “redirectors” elects to forward packets on behalf of source-destination pairs thus reducing the aggregate transmission power consumed by wireless devices. PARO is applicable to a number of networking environments including sensor networks, home networks and mobile ad hoc networks. In this paper, we present the detailed design of PARO and evaluate the protocol using simulation and experimentation. We show through simulation that PARO is capable of outperforming traditional broadcast-based routing protocols (e.g., MANET routing protocols) due to its power conserving point-to-point on-demand design. We discuss some initial experiences from an early implementation of the protocol in an experimental wireless testbed using off-the-shelf radio technology.

1. Introduction

A critical design issue for future wireless ad hoc networks is the development of suitable communication architectures, protocols and services that efficiently reduce power consumption thereby increasing the operational lifetime of network enabled wireless devices. Transmission power control used for communications impacts the operational lifetime of devices in different ways. For devices where the transmission power accounts only for a small percentage of the overall power consumed (e.g., a wireless LAN radio attached to a notebook computer) reducing the transmission power may not significantly impact the device’s operational lifetime. In contrast, for small computing/communication devices with built-in or attached radios (e.g., cellular phones, PDAs, sensors, etc.) reducing

the transmission power may significantly extend the operational lifetime of a device, thus, enhancing the overall user experience.

The design of routing protocols for wireless ad hoc networks is challenging. Bandwidth and power resources available in wireless networks represent scarce resources. The signaling overhead of routing protocols may consume a significant percentage of the available resources reducing the end user’s bandwidth and power availability. This is compounded by the fact that topology changes in wireless and mobile networks occur at a much faster time scale in comparison to wired networks. Thus, routing protocols should be capable of rapidly responding to these changes using minimum signaling and taking into account the power reserves distributed in wireless networks.

To address these challenges, we propose PARO, a power-aware routing technique for wireless ad hoc networks where all nodes are located within the maximum transmission range of each other. PARO uses a packet forwarding technique where immediate nodes can elect to be *redirectors* on behalf of source- destination pairs with the goal of reducing the overall transmission power needed to deliver packets in the network, thus, increasing the operational lifetime of networked devices.

Optimization of transmission power as a means to improve the lifetime of wireless-enabled devices and reduce interference in wireless networks is beginning to gain attention in the literature [7] [8] [3] [5] [9] [6]. Typically, more power is consumed during the transmission of packets than the reception or during “listening” periods. Transmission to a distant device at higher power may consume a disproportionate amount of power in comparison to transmission to a node in closer proximity. PARO is based on the principle that adding additional forwarding (i.e., redirectors) nodes between source- destination pairs significantly reduces the transmission power necessary to deliver packets in wireless ad hoc networks. We propose that intermediate redirector nodes forward packets between source-destination pairs

even if the source and destination are located within direct transmission range of each other. Therefore, PARO assumes that radios are capable of dynamically adjusting their transmission power on a per-packet basis.

PARO attempts to maximize the number of redirector nodes between source-destination pairs thereby minimizing the transmission power. This is in direct contrast to MANET routing protocols (e.g., AODV, DSR and TORA) [4] which attempt to minimize the number of hops between source-destination pairs. One common property of these routing protocols [4] is that they discover routes using a variety of broadcast flooding protocols by transmitting at maximum power in order to minimize the number of forwarding nodes between any source-destination pair. Wide-area routing protocols discover unknown routes using high power to both reduce the signaling overhead and to make sure routing information is entirely flooded in the network. Delivering data packets in wireless ad hoc networks using minimum-hop routes, however, requires more transmission power to reach destinations in comparison to alternative approaches such as PARO that uses more intermediate nodes. In [2], we show that MANET routing based on broadcast flooding techniques are either inefficient, because they generate too many signaling packets at lower transmission power, or are incapable of discovering routes that “maximize” the number of intermediate forwarding nodes between source-destination nodes. Because of these characteristics, MANET routing protocols do not provide a suitable foundation for discovering optimal power-aware routes in wireless ad hoc networks. As a result, there is a need to develop new power-aware routing approaches.

The design of a power-efficient routing protocol should consider both data transmission and route discovery. In terms of power transmission, these protocols should be capable of efficiently discovering routes involving multiple hops, thus minimizing the transmission power in comparison to standard flooding based ad hoc routing designs. PARO departs from broadcast-based designs and supports a node-to-node based routing approach that is more suited to the efficient discovery of power-aware routes. PARO is not only applicable as a local area routing technology where all nodes are within direct transmission range of each other (e.g., personal area networks, home networks, sensor networks, WLANs) but it can also perform power optimization as a layer 2.5 routing technology operating below wide-area MANET routing protocols. In this case, PARO provides wide-area routing protocols with local energy-conserving routes and wide-area routing is used to forward packets when the source and destination nodes are outside the maximum transmission range of each other.

The structure of this paper is as follows. Section 2 presents the PARO model and Section 3 discusses the detail design of the core algorithms that include the overhear-

ing, redirecting, route convergence and route maintenance mechanisms. Following this, enhancements to the core algorithms to support mobility are presented in Section 4. A performance evaluation of PARO, and comparison to a broadcast-based link state routing protocol that uses transmission power as the link cost unit are presented in Section 5 and Section 6, respectively. Section 7 provides some initial experiences from an early implementation of the protocol in an experimental wireless testbed using IEEE 802.11 technology. Finally, we present some concluding remarks in Section 8.

2. PARO Model

2.1. Link Assumptions

PARO requires that radios are capable of dynamically adjusting the transmission power used to communicate with other nodes. Commercial radios that support IEEE 802.11 and Bluetooth include a provision for power control. PARO assumes that the transmission power required to transmit a packet between nodes A and B is somewhat similar to the transmission power between nodes B and A. This assumption may be reasonable only if the interference/fading conditions in both directions are similar in space and time, which is not always the case. Because of this constraint PARO requires an interference-free Media Access Control (MAC) found in frequency band radios such as Channel Sense Multiple Access (CSMA). In addition, PARO requires that every data packet successfully received is acknowledged at the link layer and that nodes in the network are capable of overhearing any transmissions by other nodes as long as the received signal to noise ratio (SNR) is above a certain minimum value. Any node should be capable of measuring the received SNR of overheard packets. This includes listening to any broadcast, unicast and control (e.g., acknowledgment) packets.

2.2. Cost Function

The goal of PARO is to minimize the transmission power consumed in the network. A node keeps its transmitter “on” to transmit one data packet to another node for L/C seconds where L is the size of the transmitted frame in bits (e.g., data plus layer 2 headers) and C is the raw speed of the wireless channel in bits/second. Similarly, the receiver node keeps its transmitter on to acknowledge a successful data transmission for a combined period of l/C seconds where l is the size of the acknowledgement frame including layer 2 headers.

Now consider a network composed of several static nodes. Lets assume there are several alternative routes between a given source-destination pair in the network and

that each route involves a different set and number of forwarding nodes. Then the aggregate transmission power to forward one packet along an alternative route k , P_k , is defined as follows:

$$P_k = \sum_{i=0}^{N_k} (T_{i,i+1}L + T_{i+1,i}l)/C \quad (1)$$

The factor $T_{i,j}$ in Equation 1 is the *minimum transmission power* at node i such that the receiver node j along route k is still able to receive the packet correctly ($T_{i,j}$ will be defined formally in Section 3.1), while N_k is the number of times a data packet is forwarded along route k including the source node. Equation 1 considers transmission power only, thus, it neglects the cost of processing overheard packets and the cost of keeping the radio in a listening mode. PARO is suitable for devices for which adjusting the transmission power benefits the overall power consumption. The power consumption during the transmission mode of such devices is higher than the power consumption during reception and listening modes, as is the case with a number of commercial radios. In this case, Equation 1 represents an “idealized” communication device.

PARO mainly uses data packets for route discovery. However, in some cases the protocol uses explicit signaling to discover routes in the network, as discussed in Section 3 and Section 4. The goal of any power-efficient routing protocol should be to reduce the signaling overhead to a minimum in order to save power. PARO tries to find the route k for which the transmission power, P_k , is minimized, and furthermore, it tries to do discover this route using as little transmission power as possible. Let R_k be the transmission power consumed by the routing protocol to discover the route for which P_k is a minimum, then the cost function for transmitting Q packets between a given source-destination pair along the best route, k , is:

$$C_k = R_k + Q \sum_{i=0}^{N_k} (T_{i,i+1}L + T_{i+1,i}l)/C \quad (2)$$

PARO accommodates both static (e.g., sensor networks) and mobile (e.g., MANETs) environments. In the case of static networks, once a route has been found there is no need for route maintenance unless some nodes are turned on or off. In a static network, transmitting a large amount of data traffic (e.g., a large Q) clearly outweighs the cost of finding the best power-efficient route (R_k). In this case, PARO may not need to be as efficient while discovering such a route. In mobile environments, however, there is a need for route maintenance.

2.3. Protocol Operations

Prior to transmitting a packet, a node updates its packet header to indicate the power required to transmit the packet.

A node overhearing another node’s transmission can then use this information plus, a localized measure of the received power, to compute (using a propagation model) the minimum transmission power necessary to reach the overheard node. In this simple manner, nodes can learn the minimum transmission power toward neighboring nodes. PARO does not, however, maintain routes to other nodes in the network in advance but discover routes on a per-node on-demand basis. This approach has the benefit that signaling packets, if any, are transmitted only when an unknown route to another node is required prior to data transmission, thus reducing the overall power consumption in the network.

At first the operation of PARO may seem counter-intuitive because in the first iteration of PARO the source node communicates with the destination node directly without involving any packet forwarding by intermediate nodes (i.e., redirectors). Any node capable of overhearing both source and destination nodes can compute whether packet forwarding can reduce the transmission power in comparison to the original direct exchange between source and destination nodes. When this is the case an intermediate node may elect to become a redirector and send a *route-redirect* message to the source and destination nodes to inform them about the existence of a more power efficient route to communicate with each other. This optimization can also be applied to any pair of communicating nodes; thus, more redirectors can be added to a route after each iteration of PARO with the result of further reducing the end-to-end transmission power. PARO requires several iterations to converge toward a final route that achieves the minimum transmission power, as defined in Equation 1.

The PARO model comprises three core algorithms that support *overhearing*, *redirecting* and *route-maintenance*, as shown in Figure 1. The overhearing algorithm receives packets overheard by the MAC and creates information about the current range of neighboring nodes. Overheard packets are then passed to the redirecting algorithm, which computes whether route optimization through the intermediate node would result in power savings. If this is the case, the node elects to become a potential redirector, transmits *route-redirect* messages to the communicating nodes involved and creates appropriate entries in its redirect table. The overheard packet is then processed by the packet classifier module with the result that one of the following actions is taken: (i) the packet is passed to the higher layers if both MAC and IP addresses match; (ii) the packet is dropped if neither MAC nor IP addresses match; or (iii) the packet is forwarded to another node when only the MAC address matches. In the latter case, PARO searches the redirect table to find the next node en route and then searches the overhear table to adjust the transmission power to reach that node.

When PARO receives a data packet from the higher lay-

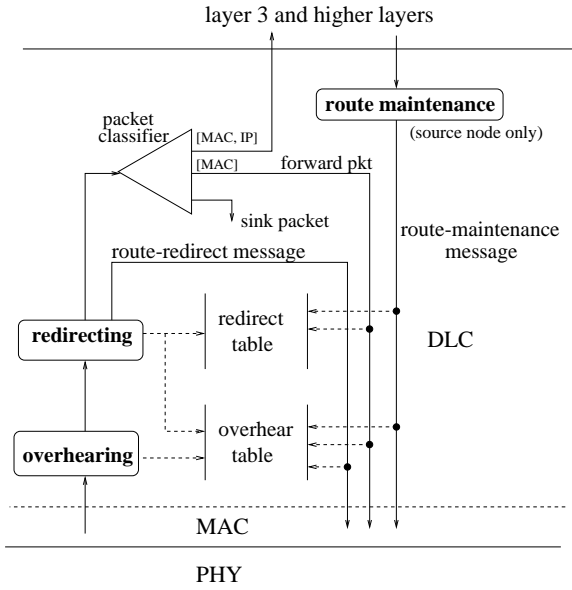


Figure 1. PARO Model

ers it searches the redirect table to see if a route toward the destination node exists. If this is not the case, PARO searches the overhear table to see if transmission power information regarding the destination node is available. If this is not the case, PARO transmits the packet using the maximum transmission power anticipating that the receiving node is located somewhere in the neighborhood. Once the destination node replies with a packet of its own then PARO's route optimization follows as described previously. PARO relies on data packets as the main source of routing information in the network. When nodes are mobile and no data packets are available for transmission, a source node may be required to transmit explicit signaling packets to maintain a route. The role of the route maintenance algorithm is to make sure that a minimum flow of packets is transmitted in order to maintain the route when there are no data packets available to send at the transmitter.

3. Protocol Design

In what follows, we first describe the necessary core algorithms for overhearing, redirecting and route-maintenance. These core algorithms provide support for static environments (e.g., sensor networks) and serve as a set of foundation algorithms for mobile environments. In Section 4, we discuss the detailed enhancements to the core algorithms to support mobility.

3.1. Overhearing

The overhearing algorithm processes packets that are successfully received by the MAC, and creates a cache entry in the overhear table or refreshes an entry in the case that information about the overheard node already exists. This cache entry contains the triple $[ID, time, T^{min}]$, where the ID is a unique identifier of the overheard node (e.g., MAC or IP address), $time$ is the time at which the overheard event occurred, and T^{min} is the *minimum transmission power* necessary to communicate with the overheard node. Definition: Let R_i^{min} be the minimum signal sensitivity level at node i at which a packet can still be received properly. If $R_{j,i}$ is the measured received signal power at node i from a packet transmitted by node j at power T_j , then the minimum transmission power for node i to communicate with node j , $T_{j,i}^{min}$, is such that $R_{j,i} = R_i^{min}$.

The computation of $T_{j,i}^{min}$ is difficult because of the time-varying characteristics of wireless channels. In our analysis and simulation results discussed later we use a two-ray propagation model. It is important to note, however, that we use the two-ray model in this paper to illustrate how a simple propagation model could be used in the operation of the protocol. As a general rule, the appropriate propagation model that best matches the operating environment should replace the simple two-ray model presented included here. The two-ray propagation model is appropriate for outdoor environments where a strong line of sight signal exists between the transmitter and receiver nodes and when the antennas are omnidirectional.

The two-ray propagation model assumes there are two main signal components. The first component is the signal traveling on the line of sight and the second component is a reflection wave from a flat ground surface. We first compute the distance separating the source and destination nodes using the two-ray model by:

$$d^4 = \frac{T_{i,j} G_t G_r h_t^2 h_r^2}{R_{j,i}}, \quad (3)$$

where d is the distance separating transmitter and the overhearing node, and $G_t h_t^2$ and $G_r h_r^2$ are the antenna gain and antenna height of the transmitter and overhearing node, respectively. It is possible to approximate $T_{j,i}^{min}$ by:

$$T_{i,j}^{min} = \frac{R_i^{min} d^4}{G_t G_r h_t^2 h_r^2} \quad (4)$$

Because of fading and other channel impairments it is not recommended to compute $T_{i,j}^{min}$ using only a single overheard packet. Rather, a better approximation for $T_{i,j}^{min}$ is to take a moving worst-case approach, $\bar{T}_{i,j}^{min}$, where the overhearing node buffers up to M previous measurements of $T_{i,j}^{min}$ and then chooses the one with the highest value. If

$T_{i,j}^{min}[k]$ is the value of $T_{i,j}^{min}$ computed for the last overheard packet then we can compute the value of $\overline{T}_{i,j}^{min}$ as:

$$\overline{T}_{i,j}^{min} = \max[T_{i,j}^{min}[k], T_{i,j}^{min}[k-1], \dots, T_{i,j}^{min}[k-M]], \quad (5)$$

where M is the number of previous measurements of $T_{i,j}^{min}$. The actual value of M can be tuned for each particular environment depending on the observed variations of the measured path attenuation. Depending on the statistical nature of these variations in time of $T_{i,j}^{min}$ a more complex computation of $\overline{T}_{i,j}^{min}$ can be provided. Similarly, we can use the two-ray model to define the minimum transmission range between nodes i and j , $\overline{D}_{i,j}^{min}$, as:

$$\overline{D}_{i,j}^4 = \frac{\overline{T}_{i,j}^{min} G_t G_r h_t^2 h_r^2}{\overline{R}_{j,i}} \quad (6)$$

3.2. Redirecting

The redirecting algorithm is responsible for performing the route optimization operation that may lead to the discovery of new routes that require less transmission power. The redirecting algorithm performs two basic operations: *compute-redirect*, which computes whether a route optimization between two nodes is feasible; and *transmit-redirect*, which determines when to transmit route-redirect messages.




Figure 2. Redirect Operation

Compute Redirect. Figure 2(a) illustrates how *compute-redirect* operates. In this example, nodes A , B and C are located within maximum transmission range of each other and, initially, node A communicates directly with node B . Because node C is capable of overhearing packets from both A and B nodes, it can compute whether the new route $A \leftrightarrow C \leftrightarrow B$ has a lower transmission power than the original route $A \leftrightarrow B$. More precisely, node C computes that a route optimization between nodes A and B is feasible if:

$$\overline{T}_{A,B}^{min} > \alpha(\overline{T}_{C,A}^{min} + \overline{T}_{C,B}^{min}) \quad (7)$$

Similarly, we define the optimization percentage of adding a redirector between two other communicating node in a

route, Opt , as:

$$Opt = \frac{(\overline{T}_{C,A}^{min} + \overline{T}_{C,B}^{min})}{\overline{T}_{A,B}^{min}} \quad (8)$$

The factor α in Equation 7 restricts the area between two communicating nodes where a potential redirector node can be selected from. For networks where nodes are static and saving battery power is important (e.g., a sensor network) α can be set around 1.1-1.2, meaning that even a small improvement in transmission power is worth the effort of adding an extra redirector (e.g., hop) to the route. Once a node computes that route optimization is feasible, it creates an entry in its redirect table that contains the IDs of the source and destination nodes, the time when the table entry is created, the IDs of the previous hop and next node en route, and the total transmission power for single packet to traverse the route. The items contained in a route-redirect message include the IDs of the source and destination nodes, optimization percentage, ID of the target node that sent the route-redirect message, ID of node transmitting route-redirect message, and the transmission power to reach the node transmitting the route-redirect message.

Transmit Redirect. Using PARO several intermediate nodes may simultaneously vie to become redirectors on behalf of a transmitting node with the result that multiple route-redirect messages are sent to a single transmitting node. Because only one intermediate node between two communicating nodes can be added as a redirector node at a time the transmission of multiple route-redirect messages (with the exception of the one transmitted by the node computing the lowest Opt percentage) represents wasted bandwidth and power resources. For sparsely populated networks, this may not be a problem. However, this is clearly an issue in the case of densely populated networks where several potential redirector nodes would be anticipated. The transmit-redirect algorithm addresses this issue by giving priority for the transmission of a route-redirect message to the potential redirector that computes lowest route optimization values first. In this manner, a potential redirector that overhears a route-redirect request from another potential redirector with a lower Opt value would refrain from transmitting its own route-redirect request (see Figure 5 (b)).

There are several ways to give preferential access to certain messages in a distributed manner. We used a simple approach that consists of applying a different time-window before transmitting a route-redirect message after the triggering event takes place (e.g., the lower the Opt value computed, the shorter the intermediate node waits to transmit its route-redirect request). The lower and upper bound of the waiting interval are set such that they do not interfere with predefined timers used by the MAC protocol, making these

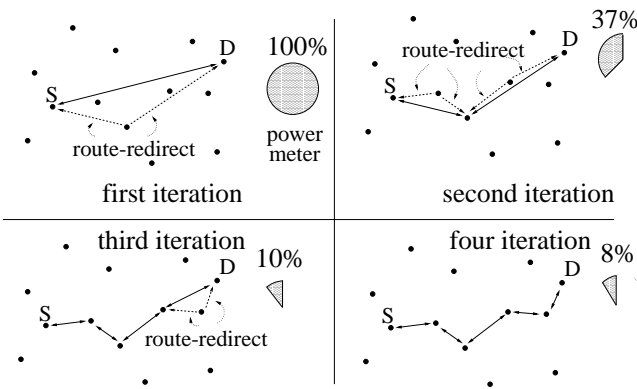


Figure 3. PARO Convergence

bounds MAC dependent. In this paper, we use the IEEE 802.11 MAC protocol and compute the waiting interval as:

$$interval = Opt * 100msec \quad (9)$$

In the unlikely scenario that more than one route-redirect request is transmitted, the target node will choose the one providing a lower Opt value. After receiving a route-redirect message, a node modifies its redirect-table putting the source of the redirect message as the next hop node (i.e., redirector) for the specific source-destination route.

3.3. Route Convergence

Previously we discussed the case where only one intermediate redirector node was added to a route between a source-destination pair. The same procedure can be applied repeatedly to further optimize a route into smaller steps with the result of adding more redirectors between source-destination nodes. Figure 3 illustrates an example of a source-destination route comprised of five segments with four redirectors requiring four iterations for route convergence. Figure 3 shows the route taken by data packets after each iteration and the intermediate nodes selected as redirectors after transmitting successful route-redirect requests.

PARO optimizes routes one step at a time, thus it requires several iterations to converge to an optimum route. The word “iteration” refers to the event in which a data packet triggers a node to transmit a route-redirect request for the first time. As a result PARO will converge as fast as the transmission speed of data (e.g., a flow measured in packets per second) transmitted by a source. Applications based on TCP (e.g., FTP, HTTP, etc.) transmit packets in bursts, potentially providing faster convergence. Applications based on UDP, on the other hand, are suitable for transmission of real-time media where the periodicity of packets transmitted depends on each specific application, thus the convergence of a route is application specific.

Figure 3 illustrates the transmission power (see “power meter”) used to transmit one packet between source and destination nodes after each iteration of PARO. During the first iteration, the source node communicates directly with the destination node. Lets consider the transmission power $\overline{T}_{S,D}^{min}$ corresponds to 100% when no redirector is presented. During the second iteration, adding one redirector in the route reduces the transmission power by 63% compared to the original $\overline{T}_{S,D}^{min}$ value. Note that the third and four iterations represent less impressive reductions in transmission power, especially the last iteration which only provides a 2% improvement. A nice property of PARO is that even after the first iteration of the protocol, considerable savings in transmission power is achieved. This means that nodes do not have to wait for the protocol to converge on the best/final route before obtaining significant power saving benefits. It can be observed from Figure 3 that each iteration simply adds one more redirector between adjacent forwarding nodes found in the previous iteration. In this respect, the new redirectors added to a route during an iteration are very much dependent on the redirectors found in the previous iteration. It is possible that the first iteration, which seemed optimal (e.g., it optimized the route better than any other intermediate node), can lead to a final route which is not the route achieving the minimum transmission power. In fact, PARO cannot avoid this from a practical point of view unless an exhaustive search is applied which works against saving power in the network. Therefore, the use of terms such as “optimum” and “minimum” assume this caveat when used in the context of PARO.

4. Mobility Support

In static networks (e.g., sensor networks) there is no need for route maintenance once the initial route between source-destination pairs has been found, other than when nodes are turned off and on. Adding support for mobile nodes to the core algorithms is challenging because of the uncertainty concerning the current range of neighboring nodes as they move in the network. In what follows, we discuss the necessary enhancements to the core algorithms to support mobility.

4.1. Route Maintenance

PARO relies on data packets as the main source of routing information. In the case of mobile nodes, data traffic alone may not be sufficient to maintain routes. Consider the extreme case of a source node transmitting packets once every second to a destination where every node moves at 10 meters/second on average. In this example, information about the range of the next redirector en route would be

outdated as a basis for the transmission of the next packet. Depending on node density and mobility there is a need to maintain a minimum rate of packets between source and destination pairs in order to discover and maintain routes as redirectors move in and out of existing routes.

A natural solution to this problem is to let the source node transmit explicit signaling packets when there is no data packets available to send. Transmitting signaling packets, however, consumes bandwidth and power resources even if those signaling packets are only a few bytes in length. Under fast mobility conditions signaling packets could potentially consume more power resources than the case where a source communicates directly with a destination node assuming certain traffic patterns. In what follows, we discuss a number of enhancements to the overhearing and redirecting algorithms to resolve these issues in support of mobile nodes.

4.2. Overhearing

Any node transmitting a packet to the next hop redirector in the route has to determine the next hop's current range, which may be different from its last recorded position. Clearly, the preferable transmission estimate is the one that transmits a packet using the minimum transmission range. If a node transmits a packet assuming that the next hop's current range is the same as the last recorded range, then three scenarios may occur: (i) The current position of the next redirector is within the current transmission range. In this case, the transmitting node finds the next redirector but some power is wasted because more power is used than necessary for this operation. (ii) The current position of the next redirector is at the same transmission range thus the transmission is optimum. (iii) The current position of the next redirector is outside the current transmission range. In this case, the transmitting node fails to find the next redirector and has to attempt a new transmission using more power than the current level.

Scenario 3 is more inefficient than Scenario 1 because not only is more power used, but also longer delays are experienced in reaching the next hop. An intuitive solution to this problem is to transmit a packet with a higher transmission range than previously recorded, increasing the probability of reaching the next hop node on the first attempt. We define a new minimum transmission range, $\overline{D}_{i,j}^{new}$, as:

$$\overline{D}_{i,j}^{new} = \overline{D}_{i,j}^{old} + \Delta, \quad (10)$$

where Δ represents how much the transmitting node overestimates the transmission range of the next node en route. The value of Δ depends on the average speed of nodes and the time interval between the last time the next redirector en route was overheard and the current time; we refer to this interval as the *silence-interval*. The longer the

silence-interval the greater the uncertainty about the current range of the next node and therefore the larger the value of Δ . We resolve this problem by requiring the source nodes transmit *route-maintenance* packets toward destination nodes whenever no data packets are available for transmission for a specific interval called *route-timeout*. Transmission of route-maintenance messages only occurs whenever a node, which is actively communicating with another node, stops transmitting data messages for a route-timeout period. The transmission of route-maintenance messages puts an upper bound on the silence-interval, thus, an upper bound on Δ .

4.3. Redirecting

Because of mobility, a redirector node may move to a location where it no longer helps to optimize the transmission power between two communicating nodes. In this case, it is necessary to remove such a node from the path using a route-redirect message. Figure 4 illustrates this scenario. Node A communicates with node D using nodes B and C as redirector nodes, as shown in Figure 4(a). Figure 4(b), shows the position of nodes after some time has elapsed. In Figure 4(b) node B moves to a position where both nodes B and C have the same transmission range from node A. When node A sends a packet to node B, it is also overheard by node C. Because node B is the previous hop to node C along the route between nodes A and D, then node C can determine that node B has moved out of the optimum route. In this case, node C transmits a route-redirect message toward node A requesting node A to re-route its data packets directly to node C. Figure 4 (c) shows the new route after node A re-routes new packets to node C.




Figure 4. An Example of Removing a Suboptimal Redirector from an Existing Route

5. Performance Evaluation

In this section, we present an evaluation of PARO and discuss a number of performance issues associated with power optimization and route maintenance. We used the *ns* network simulator with the CMU wireless extension [1] to evaluate PARO. The simulator supports physical, link and routing layers for single/multi hop ad-hoc networks. The propagation model is based on a two-ray model discussed in Section 3.1. After receiving a packet each node invokes a propagation model to determine the power at which the packet was received. If the node determines that the packet was successfully received (e.g., the received power was above a certain threshold) it passes the packet to the MAC layer. If the MAC layer receives an error-free packet it passes the packet to the link layer and so on. The simulation uses the standard *ns*/CMU mobility model.

We use the IEEE 802.11 MAC protocol which uses Channel Sense Multiple Access with Collision Avoidance (CSMA/CA) also referred to in IEEE 802.11 as the Distributed Coordination Function (DCF). In IEEE 802.11 a packet is successfully captured by a node's network interface if the sensed SNR of the received packet is above a certain minimum value¹ otherwise the packet cannot be distinguished from background noise/interference. Communication between two nodes in IEEE 802.11 uses RTS-CTS signaling before the actual data transmission takes place. Due to the potential problem of nodes not being able to listen to RTS-CTS packets in the case of a system with dynamic transmission power control, we always transmit RTS-CTS packets at maximum transmission power.

5.1. Power Optimization

As discussed in Section 3.3, the more densely populated the network the higher the average number of potential redirector nodes, and the lower the average transmission power between source-destination pairs. The simulation topology consists of a 100x100 network with 10, 30 and 100 randomly positioned static nodes for each experiment. The simulation trace lasts for a duration of 100 seconds with ten UDP/CBR flows transmitting 512 bytes packet every three seconds. The simulation uses a value for $\alpha = 1$ which configures PARO to find the best power-efficient route. Figure 5 shows that the aggregate power necessary to transmit all data packets versus the number of nodes in the network. Figure 5 also indicates (between parenthesis) the average number of times a packet is forwarded before reaching its destination node (i.e., average number of redirectors en route). This number is dependent on the number of nodes

¹For Wavelan, this values corresponds to 0.2818 watts for normal power transmission; 1.559e-11 watts for carrier sense threshold to detect a collision; and 3.652e-10 watts for the sensitivity of receiver.




Figure 5. Transmission Power versus Average Number of Redirectors

and node density, as mentioned previously. The higher the number of nodes in the network the higher the probability of having more redirectors between communicating nodes. We observe that the aggregate transmission power decreases as the number of redirector nodes increases. At first the aggregate transmission power decreases rapidly when there are between an average of 0.5 and 2.9 redirectors present. The aggregate transmission power then decreases slowly up to an average of 5.4 intermediate redirector nodes, as shown in the simulation plot.

Figure 5 shows that in terms of transmission power alone, it does not pay to have more than three redirectors per source-destination pair. Having more than three redirectors may increase end-to-end delay and likelihood of network partitions. Figure 5 also indicates the transmission power needed if no redirectors were added between source-destination pairs. Comparing the two scenarios (i.e., with and without redirectors) in Figure 5, we clearly observe the benefit (i.e., power savings) of adding intermediate redirector nodes. However, even if no intermediate nodes are found between source-destination pairs, by default PARO will use the minimum transmission power information (if available) to communicate with a destination node. This operation is in contrast with traditional wireless LAN systems, which always use the maximum transmission power to communicate with a destination node even if the destination node is in very close proximity to the transmitter.

5.2. Route Maintenance

In this section, we analyze the performance of PARO in support of mobile nodes. Figure 6 shows the transmis-

sion success ratio versus the speed of nodes and the packet inter-arrival interval. We define the “transmission success ratio” as the number of packets that are correctly received by the corresponding destination nodes divided by the total number of packets transmitted. The simulation includes 30 nodes in a 100x100 network. Ten randomly chosen nodes transmit a UDP/CBR flow to 10 randomly chosen destination nodes. Each flow consists of 100 byte packets transmitted using different time intervals. In Figure 6, we highlight three separate regions on the graph which are of interest because of the different network dynamics operating in those regions; these are as follows. Region (I): Nodes operating in this region move slowly. As a result, redirectors remain in the path of a route for longer intervals which translates into fewer route/updates per second. This condition results in a high transmission success ratio, even in the case of a slow flow of packets traversing between source-destination pairs. Region (II): Nodes operating in this region transmit packets with small inter-arrival intervals. The faster data packets are transmitted the faster PARO can discover, for example, that a redirector has moved to a different location and to take appropriate measures. As a result, the transmission success ratio is high even for the case where nodes move fast. Region (III): Nodes operating in this region move fast and transmit packets slowly. Because of high mobility several route changes per second occur. However, packets are not transmitted at a fast enough rate to maintain routes in the network due to the long silence-intervals between packets. Data packets transmitted by nodes operating in this region are likely to be lost. This is because transmitting nodes may not have accurate range information concerning the next hop redirectors en route. As a result, the transmission success ratio is low. Figure 6 also shows the importance of transmitting route-maintenance packets to maintain a route in the case where a source node transmits packets too slowly.

Determining the optimum value of the silence-interval (introduced in Section 4.2) to overcome node mobility (in order to guarantee a certain success ratio) is a complex issue. This value is dependent on the size of the network and the node density as well as mobility and data packet inter-arrival rate. Larger areas with high nodal density will likely support routes with several redirectors. Maintaining a route with fewer redirectors requires less signaling packets both in terms of route-redirect and route-maintenance messaging. A route reduces the transmission power by a significant amount simply by limiting the number of redirectors to 2-3 forwarding nodes, as discussed in Section 5.1. The benefit of adding additional redirectors beyond this point may be undermined by the signaling overhead required to maintain longer multi-hop routes. Two complementary methods can be used to reduce the number of redirectors along a route. Choosing a higher value for α (see Section 3.2) restricts the




Figure 6. Transmission Success Performance

area where a redirector can be located between two communicating nodes. Such an approach would reduce the number of redirectors compared to the case where a parameter value of $\alpha = 1$ is adopted. Second, packets could carry a counter similar to the IP packet TTL field that would be decremented by each redirector visited en route toward the destination. After reaching zero, no other redirectors would be added to further optimize the route. This enhancement is currently being studied.

6. Comparison

PARO discovers routes on-demand on a node-to-node basis. An alternative approach would generate full routing tables in advance where, for example, all nodes would be aware of power-efficient routes to all other nodes in the network. Such protocol behavior is similar to Link State Routing (LSR) using transmission power as the link cost unit. We refer to this modification to LSR as MLSR in the remainder of this section. The basic LSR operation requires each node in the network to broadcast a routing packet (or PROP message using link state terminology). The PROP packet contents contains information about the transmission cost of all known destinations. After collecting PROP messages from all parts of the network, any node should be capable of computing optimum routes to any other node in the network.

Because of the fundamental difference in these two approaches, we compare PARO and MLSR to best understand the various tradeoffs and limitations of our design. In what follows, we describe an MLSR implementation that supports transmission power as in the case of PARO. We then compare the performance of MLSR to PARO. Consider a network composed of N nodes located within transmis-




Figure 7. Aggregated Transmission Power Consumed by Data and Signaling for PARO and MLSR

sion range of each other. MLSR nodes can compute the minimum transmission power T^{min} to a transmitting node by listening to a PROP signaling packet transmitted by the node. The PROP message includes the transmission power T^{PROP} used to transmit the packet. Depending on the value of T^{PROP} , the content of a PROP message may require to be forwarded by other nodes to propagate the entire network. Each node computes routes to any other node in the network using a standard link-state Dijkstra algorithm. In a network of N nodes, it takes K iterations (i.e., K PROP packets transmitted by each node) for the content of a PROP message to be entirely flooded in the network. The value K mainly depends on the parameter T^{PROP} and the density of nodes in the network.

Figure 7 shows a simulation trace of the aggregate transmission power consumed by both signaling and data packets for both PARO and MLSR. The network simulation consists of 30 static nodes a 100x100 in size with ten UDP/CBR flows transmitting a 100-byte packet every 3 seconds. In the case of MLSR, signaling packets are first transmitted at different transmission ranges to generate full routing tables. Once routing information is available MLSR data packets are transmitted using power-efficient routes. In the case of PARO, data packets are first transmitted at high power because the range of destination nodes is unknown to source nodes. Figure 7 shows the transmission “power offset” (shown in the figure as the initial fast increase in power consumption) while the routing protocol converge to optimum routes for both PARO and MLSR. In the case of MLSR, this offset is independent of the number of active sessions and dependent on the number of nodes in the network and

the number of iterations required for the content of a PROP message to flood the network. This means that if there is double the number of nodes in the network then the value of the offset would roughly double. In contrast, the routing offset for PARO depends on the number of active sessions. Therefore, PARO is less sensitive to the number of nodes in the network. We observe from Figure 7 that relative to the power consumed by the first data and signaling packets, the contribution of data transmission to the overall power consumption is less significant. This result suggests an important design principle for future power-aware routing protocols is the avoidance of “blind” (e.g., broadcast) transmissions at high power.

In the case of the MLSR simulations, a transmission range of $D_{max}/4$ represented the lowest transmission range observed before route partitions appeared in the network. As discussed previously, route partitions appear because broadcast messages do not completely flood the network. When we consider a transmission range of $D_{max}/5$ for PROP messages (not shown in Figure 7), we observe that network partitions consistently appear leaving nodes with routes to only a subset of destination nodes. This result emphasize the fact that even if the performance of MLSR at $D_{max}/4$ is somewhat similar to PARO (i.e., being able to reduce its transmission range), this operation results in non-stable performance. In addition, it is unlikely that MLSR could find such a transmission range in a practical setting.

7. Implementation

In what follows, we discuss some experiences implementing PARO in an experimental wireless ad hoc testbed. For a more detailed discussion see [2]. We implemented PARO using the Linux Redhat 6.2 software platform on 700 MHz Pentium III notebooks equipped with Aironet PC4800 series radios. The Aironet PC4800 supports the IEEE 802.11 standard and provides five different transmission power levels (viz. 1, 5, 20, 50 and 100 milliwatts). The overhearing, redirecting, and route-maintenance algorithms are implemented in user space using the Berkeley Packet Filter’s Packet Capture Library (PCAP) for processing and forwarding of IP packets.

One initial drawback of using the Aironet PC4800 radio as a basis to implement PARO was that it could only approximate the minimum transmission power much of the time. This was a product of only offering a small set of transmission power levels. PARO software is designed to always rounded up to the next available power level. For example, if PARO computed the minimum transmission power to be 10 milliwatts then the packet would be transmitted at 20 milliwatts. This has the impact of using more power than necessary but the extra margin is useful in the case of mobility and stability of routes. Figure 8 shows the aggre-




Figure 8. Experimental Results for Transmission Power versus the Position of a Redirector between a Source-Destination Pair for Indoor and Outdoor Environments

gate transmission power necessary to transmit one packet between a source-destination pair using a single redirector. We positioned the redirector node at different locations between a source and destination. Figure 8 shows the power optimization results for an “ideal” transceiver (determined by Equation 1) against results obtained from the Aironet radio. We conducted experiments in both indoor and outdoor settings. The indoor path attenuation model implemented in the testbed used a propagation model with a path attenuation of $n = 3.25$ with a standard deviation $\sigma = 16.3$ [dB]. The outdoor model used a typical path attenuation of $n = 2$. Figure 8 confirms that the Aironet PC4800 transceiver can only approximate the performance of the ideal transceiver. Some anomalies are highlighted in the graph. For example, when the redirector is positioned at the mid point between the source and destination the ideal transceiver offers significant savings, as discussed earlier. In the case of the outdoor experiment, however, positioning the redirector at the mid point provides no power savings. Such anomalies are mainly the product of the operational granularity (i.e., transmission power levels available) of the radio. With the exception of these limitations, the experimental results indicate that PARO can be implemented using existing technology and that the protocol delivers transmission power savings.

8. Conclusion

In this paper, we have presented PARO, a power-aware routing optimization for wireless ad hoc networks. We evaluated PARO and compared its performance to MLSR. We

found that PARO consumed less power in order to find power-efficient routes compared to MLSR due to its point-to-point on-demand design. An early implementation of the PARO system using a commercial IEEE 802.11 radio showed a basic proof of concept even though some inefficiencies and anomalies were identified. Currently, we are studying the performance of Internet applications and transport protocols operating over PARO. We are particularly interested in studying quality of service issues such as delay, “goodput” and packet error rates under such a regime. Furthermore, we are investigating complementary techniques that help save reception and idle power in PARO-based wireless ad hoc networks.

9. Acknowledgement

This work is supported in part by the National Science Foundation (NSF) under the WIRELESS TECHNOLOGY Award ANI-9979439 and with support from IBM Research.

References

- [1] The cmu monarch project. <http://www.monarch.cs.cmu.edu/>, 2000.
- [2] J. Gomez, A.T. Campbell, M. Naghshineh, and C. Bisdikian. Supporting transmission power control for routing in wireless ad hoc networks. *COMET Group Technical Report*, September 2001.
- [3] David Goodman and Nayaran Mandayam. Power control for wireless data. *Proceedings, IEEE International Workshop on Mobile Multimedia Communications (MOMUC)*, 1999.
- [4] IETF MANET Working Group. <http://www.ietf.org/html.charters/manet-charter.html>. 2001.
- [5] T. J. Kwon and Mario Gerla. Clustering with Power Control. *In Proceedings of IEEE MILCOM, Atlantic City, NJ*, November 1999.
- [6] J. Monks, V. Bharghavan, and W. Hwu. Transmission power control for multiple access wireless packet networks. *proceedings of IEEE Conference on Local Computer Networks, Tampa, FL*, November 2000.
- [7] F. Rashid-Farrokhi, K. J. Ray Liu, and L. Tassiulas. Transmit Beamforming and Power Control for Cellular Wireless Systems. *IEEE Selected Areas in Communications Journal (JSAC)*, Vol. 16. No. 8, October 1996.
- [8] S. Singh, M. Woo, and C.S. Raghavendra. Power-aware Routing in mobile ad hoc Networks. *Proceedings, ACM/IEEE MOBICOM, Dallas, Texas*, October 1998.
- [9] Jeffrey Wieselthier, Gam Nguyen, and Anthony Ephremides. On the construction of energy-efficient broadcast and multicast trees in wireless networks. *proceedings of IEEE INFOCOM, Tel-Aviv, Israel*, 2000.