



PARO: Supporting Dynamic Power Controlled Routing in Wireless Ad Hoc Networks

JAVIER GOMEZ* and ANDREW T. CAMPBELL

Department of Electrical Engineering and Center for Telecommunications Research, 500 W. 120th Str., Mail Code 4712, Columbia University, NY 10027, USA

MAHMOUD NAGHSHINEH and CHATSCHIK BISDIKIAN

IBM T.J. Watson Research Center, 30 Saw Mill River Road, Hawthorne, NY 10953, USA

Abstract. This paper introduces PARO, a dynamic power controlled routing scheme 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 wireless 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 energy conserving point-to-point on-demand design. We discuss our experiences from an implementation of the protocol in an experimental wireless testbed using off-the-shelf radio technology. We also evaluate the impact of dynamic power controlled routing on traditional network performance metrics such as end-to-end delay and throughput.

Keywords: power optimization, power control, ad hoc networks, routing protocols

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., sensors) 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.

* Corresponding author.

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 [5,13,15,20,24,27]. 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 uses redirector nodes to shorten the length of individual hops, thereby reducing the overall power consumption. This approach is in direct contrast to MANET routing

protocols (e.g., AODV, DSR and TORA) [11], which attempt to minimize the number of hops between source–destination pairs. One common property of these routing protocols [11] 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 this paper, 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 overhearing, 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 sections 5 and 6, respectively. Section 7 discusses our experiences from an implementation of the protocol in an experimental wireless testbed using IEEE 802.11 technology. In section 8, we study the impact that dynamic power control routing schemes such as PARO have on traditional network performance metrics such as throughput and

delay. Finally, we present related work in section 9 and some concluding remarks in section 10.

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). Note that even in CSMA access protocols, packets are subject to interference (collisions) during the sensing period, as a result of hidden terminals. 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 acknowledgment 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,

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

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 sections 3 and 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

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

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 them 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

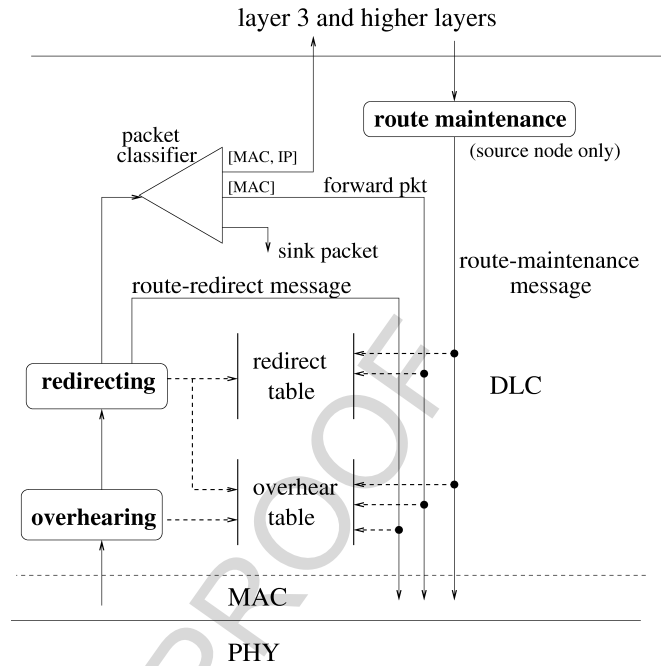


Figure 1. PARO model.

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 addresses match. 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 layers it searches the redirect table to determine if a route toward the

destination node exists. If this is not the case, PARO searches the overhear table to determine if there is any transmission power information related to the destination node 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 a 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 , $T_{i,j}^{\min}$, is such that $R_{j,i} = R_i^{\min}$.) It is important to note that overhearing does not add energy consumption because signal strength measurements can be derived from L2 headers.

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 traditional propagation model that considers the strength of the received signal to be $\sim T/d^\gamma$. It is important to note, however, that other propagation models that best match a particular operating environment should replace the simple model presented

here. We first compute the distance separating the source and destination nodes by

$$d^\gamma = \omega \frac{T_{i,j}}{R_{j,i}}, \quad (3)$$

where d is the distance separating the transmitter and the overhearing node, γ is the attenuation factor of the environment typically in the range 2–4 (e.g., for indoor and outdoor environments) and ω is a proportionality constant that typically depends on factors such as antenna gain and antenna height of the transmitter and overhearing nodes. Initially a transmitter use $T_{j,i} = P_{\max}$ if no previous information about the intended receiver is known. After this $T_{j,i}^{\min}$ can be approximate by

$$T_{i,j}^{\min} = \frac{R_i^{\min} d^\gamma}{\omega}. \quad (4)$$

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, $\overline{T}_{i,j}^{\min}$, 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 define the minimum transmission range between nodes i and j , $\overline{D}_{i,j}^{\min}$, as

$$\overline{D}_{i,j}^\gamma = \frac{\omega \overline{T}_{i,j}^{\min}}{R_{j,i}}. \quad (6)$$

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.

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

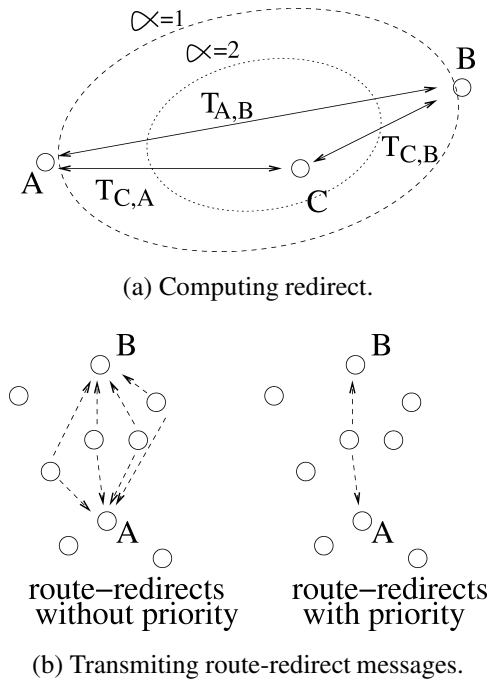


Figure 2. Redirect operation.

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 nodes 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. In figure 2(a), we show the equivalent region where a potential redirector can be located for $\alpha = 1$ and $\alpha = 2$. The size and shape of these regions for finding potential redirectors depend mainly on the propagation loss parameter. For networks where nodes are static and saving battery power is important (e.g., a sensor network) α can be set to approximately 1.1–1.2, meaning that even a small improvement in transmission power savings is worth the effort of adding an extra redirector 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 contend 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 candidate 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 candidate 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 refrains from transmitting its own route-redirect request (see figure 2(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 bounds MAC dependent. In this paper, we use the IEEE 802.11 MAC protocol and compute the waiting interval as

$$interval = Opt \cdot 100 \text{ msec}. \quad (9)$$

In the unlikely scenario that more than one route-redirect request is transmitted, the target node will choose the one providing the lowest 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 a specific source–destination route.

3.3. Route convergence

Previously we discussed the case where only one intermediate redirector node is 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

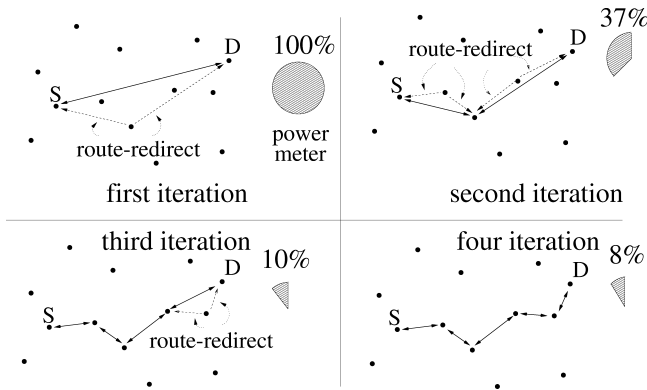


Figure 3. PARO convergence.

time. As a result PARO will converge as fast as the transmission rate 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 the source and destination nodes after each iteration of PARO. During the first iteration, the source node communicates directly with the destination node. Let us consider that the transmission power $\bar{T}_{S,D}^{\min}$ 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 $\bar{T}_{S,D}^{\min}$ value. Note that the third and fourth 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 to 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 or on. However, in many cases nodes are mobile (e.g., MANETs). 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 are 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, $\overline{D}_{i,j}^{\text{new}}$, as

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

where Δ represents how much the transmitting node over estimates 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 that 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 new 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 are within the same transmission range of 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.

5. Performance evaluation

In this section, we present an evaluation of PARO and discuss a number of performance issues associated with route convergence, power optimization and route maintenance.

5.1. Simulation environment

We used the *ns* network simulator with the CMU wireless extension [26] to evaluate PARO. The simulator supports phys-




Figure 4. An example of removing a suboptimal redirector from an existing route.

ical, link and routing layers for single/multi hop ad-hoc networks. The propagation model is based on a two-ray model, which is appropriate for outdoor environments where a strong line of sight signal exists between the transmitter and receiver nodes and where 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. This model computes the strength of the received signal source and destination nodes by

$$R_{j,i} = \frac{T_{i,j} G_t G_r h_t^2 h_r^2}{d^4}, \quad (11)$$

where d is the distance separating transmitter from the over-hearing node, and $G_t h_t^2$ and $G_r h_r^2$ are the antenna gain and antenna height of the transmitter and over-hearing node, respectively. After receiving a packet each node invokes the 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 power 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-

¹ For 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. An example of the problem of transmitting RTS-CTS packets using dynamic transmission power control.

CTS packets in the case of a system with dynamic transmission power control, we always transmit RTS-CTS packets at maximum transmission power. Figure 5 illustrates this problem. In the figure node A communicates with node B while at the same time node C communicates with node D. In this scenario nodes C and D transmit RTS-CTS packets using minimum transmission power. Under such conditions nodes A and B may not be able to overhear (dashed circle) or sense (dotted circle) the RTS-CTS packet exchange between nodes C and D and may attempt to transmit their own RTS-CTS thereby interfering and disrupting the on-going communication between nodes C and D.

Clearly transmitting RTS-CTS packets at maximum transmission power does not exploit the spectral reuse potential in the network. A node transmitting a packet to another node in close proximity at the minimum transmission range uses RTS/CTS at full transmission range. This inhibits other nodes in the entire RTS/CTS region from transmitting even if deferring transmission for the nodes is unnecessary. There are a number of new MAC proposals that address such limitations. In [15] the authors present the Power Controlled Media Access Protocol (PCMAP) that operates within the framework of collision avoidance protocols such as CSMA/CS that use RTS/CTS. In PCMAP, active receivers advertise a periodic busy tone on a separate frequency band to other potential transmitters including their maximum tolerance to admit extra noise (e.g., interference). A node intending to transmit a packet first senses the busy tone signal. If a busy tone exists, then the node adjusts its transmission power such that it does not disrupt ongoing transmissions prior to communication with its intended receiver. We believe MAC protocols such as PCMAP can efficiently support the necessary power-controlled operations required by PARO in comparison to off-the-shelf radios such as IEEE 802.11. We discuss these limitations further in section 7 and section 8.

As a general methodology comment each point in the graphs shown in the following sections on route convergence, power optimization and route maintenance represents an average of five different simulation runs. Each simulation run uses




Figure 6. Route convergence time.

a different seed number affecting both the traffic and mobility behavior.

5.2. Route convergence

Figure 6 shows simulation results concerning the convergence of PARO versus different packet inter-arrival rates. Twenty static nodes are randomly positioned in a 100×100 network. We conducted two separate experiments for UDP/CBR and TCP/FTP applications. In each experiment each node is the source and recipient of a flow. In the case of UDP/CBR applications, each source node transmits a 512-byte packet with different inter-packet intervals times ranging from 30 msec to 1.5 seconds. In the case of the TCP/FTP applications, each source node transmits 512-byte packets as fast as the link layer permits. As anticipated the results show that PARO converges in the same proportion as the inter-packet interval times. Thus, the faster nodes transmit packets the faster routes converge. In the case of TCP/FTP applications, this time represents a few dozens milliseconds (the corresponding points in figure 6 are so close to each other that they appear to be overlapping). As discussed in section 3.3, PARO requires several iterations to converge to an optimum route with minimum power. The number of iterations per session is dependent on the node density and the specific position of nodes with respect to each other. Because different sessions may require a different number of iterations to converge, the session needing more iterations will take the longest time to converge assuming all sessions have similar traffic patterns and start at the same time. Figure 6 also contrasts the convergence of PARO for different number of iterations for the same network size and number of nodes. As expected PARO converges linearly with respect to the number of iterations required.

5.3. Power optimization

As discussed in section 3.3, the more densely populated the network the higher the average number of potential redirector




Figure 7. Transmission power versus number of nodes.

nodes, and the lower the average transmission power between source–destination pairs. The simulation topology consists of a 100×100 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 7 shows that the aggregate power necessary to transmit all data packets versus the number of nodes in the network. Figure 7 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 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. 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. We observe the aggregate transmission power decreases as the number of nodes increases from 10 to 30. This is a consequence of the availability of additional appropriately located redirectors.

Figure 7 shows that in terms of transmission power alone, it does not pay to have more than three redirectors per source–destination pair in networks where nodes are distributed homogeneously.² Having more than three redirectors may increase end-to-end delay and likelihood of network partitions. Figure 7 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 7, we clearly observe the benefit (i.e., power savings) of adding intermediate redirector nodes. However, even if no intermediate nodes are found between source–

² When nodes are not distributed homogeneously in the network it may occur that having 4 or even 5 redirectors per route on the average provides a noticeably power savings improvement.

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.

5.4. Route maintenance

In this section, we analyze the performance of PARO in support of mobile nodes. Figure 8 shows the transmission 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 100×100 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 8, 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 8 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,




Figure 8. Transmission success performance.

as discussed in section 5.3. 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 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 (where the 'M' in MLSR stands for Modified LSR) 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 transmission range of

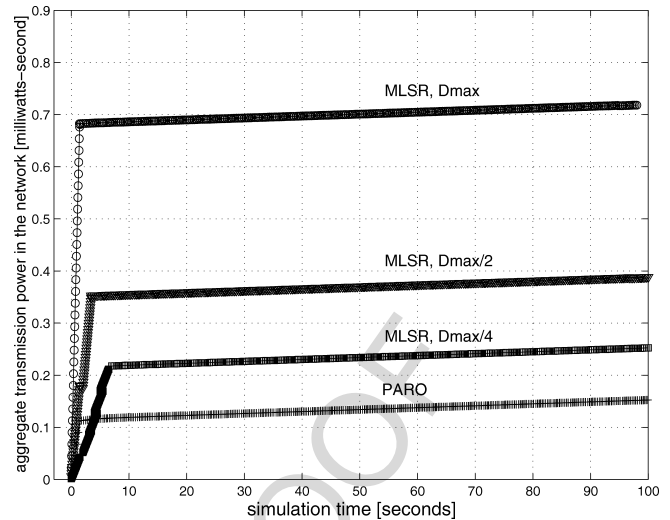


Figure 9. Aggregated transmission power consumed by data and signaling for PARO and MLSR.

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 flood 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 and size of the network.

Figure 9 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 100×100 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 9 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 9 that relative to the power consumed by the first data and




Figure 10. Signal coverage at different transmission powers for an indoor experiment.

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 our experiences implementing PARO in an experimental wireless ad hoc testbed. 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. We conducted experiments with PARO operating in both indoor and outdoor settings.

7.1. Propagation model

Figure 10 shows the area covered by a transmitter used for our indoor experimentation. This represents an indoor laboratory environment for 1, 5, 20, 50 and 100 milliwatts transmission levels. There has been considerable work on propagation models for indoor environments [10,22]. The main purpose of this experiment is to illustrate what can be expected for this particular IEEE 802.11 radio in an indoor laboratory setting as a basis for understanding PARO’s approach to dynamic power control. In this experiment, we kept one radio in a fixed position while we moved a second radio around the corridors of the floor. Both radios use the same transmission power level and transmit five small UDP packets to each other every second.

We define the coverage area for a given transmission power level as the area for which both radios did not observe packet loss. As we can observe from figure 10, the path attenuation factor for this setting is quite strong, especially around the corridor corners. The strong attenuation is mostly due to radio signals going through walls, floor, ceilings and metal obstacles. Strong attenuation factors emphasize the advantages of performing PARO-style route optimization. In the extreme case, where node A communicates with node B through node C, an aggregate transmission power of two milliwatts is required compared with the original 100 milliwatts when node A communicates directly with node B. In those environments where direct communications between two nodes is not possible due to signal obstacles, PARO can improve connectivity by adding redirectors in other locations where signals can travel more freely.

In the implementation of PARO we used two different path attenuation models in place of equation (3) depending of the type of environment. For the outdoor environment we used a typical path attenuation of $n = 2$ [10]. In contrast, for the indoor environment we used a propagation model presented in [22] with a path attenuation of $n = 3.25$ with a standard deviation $\sigma = 16.3$ [dB]. The model in [22] was obtained from an office building with a large layout area divided into several smaller cubicles-style offices. The path attenuation for a transmitter and receiver separated by d meters is defined by [22]:

$$PL(d)[\text{dB}] = \overline{PL}(d)[\text{dB}] + X_{\sigma}[\text{dB}]. \quad (12)$$

With $\overline{PL}(d)[\text{dB}]$ the mean path attenuation between transmitter and receiver separated by d meters and $X_{\sigma}[\text{dB}]$ is a zero mean log-normally distributed random variable with standard deviation σ in decibels. The parameter $\overline{PL}(d)[\text{dB}]$ is computed as follows:

$$\overline{PL}(d)[\text{dB}] = PL(d_0[\text{dB}]) + 10n \log_{10} \left(\frac{d}{d_0} \right), \quad (13)$$

where $PL(d_0[\text{dB}])$ is the free-space propagation from the transmitter to a 1 meter reference distance and n is the path attenuation factor. For complete details see [22].

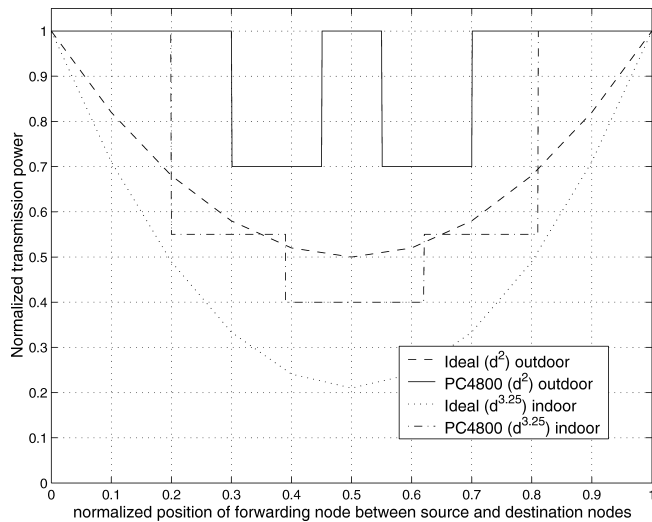


Figure 11. Experimental results for transmission power versus the position of a redirector between a source–destination pair for indoor and outdoor environments.

7.2. Power optimization

One initial drawback of using the Aironet PC4800 radio as a basis to implement PARO is that it could only approximate the minimum transmission power much of the time. This is a product of only offering a small set of transmission power levels. PARO software is designed to always round 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 using the Aironet radio. 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 11 shows the aggregate transmission power necessary to transmit one packet between a source–destination pair using a single redirector; that is, a single packet between nodes A and B using packet forwarding by redirector node C. Node C is positioned at different locations along a line between nodes A and B. Figure 11 shows the power optimization results for an “ideal” transceiver (determined by equation (1)) against results obtained from the Aironet radio. Figure 11 confirms that the Aironet PC4800 transceiver can only approximate the performance of the ideal transceiver.

Table 1 shows the aggregate transmission power needed to transmit a single packet between nodes A and B using not one but several intermediate redirector nodes. In all cases we evenly distributed the forwarding nodes between source and destination pairs. For both outdoor and indoor environments we separated the source and destination nodes to the maximum distance allowed while transmitting at 100 milliwatts. From table 1, we observe that power optimization is better for stronger path attenuation conditions (i.e., indoor versus outdoor). This result is expected since the strength of radio waves decay faster under strong path attenuation settings. Thus, nodes in indoor environments benefit much more of the presence of redirector nodes.

Table 1

Aggregate transmission power versus number of nodes for outdoor and indoor settings.

	Number of forwarding nodes					
	0	1	2	3	4	5
$n = 2$ (outdoor)	100	100	60	80	25	30
$n = 3.25$ (indoor)	100	40	15	20	5	6

7.3. Discussion

The experimental results show that PARO can be partially implemented using off-the-shelf radio technology providing transmission power savings. However, due to a number of limitations with existing radio and software support technology the full power savings of PARO are difficult to attain today. In what follows, we discuss our implementation experiences and the impact of these limitations on the potential gains of PARO. Much of these comments are driving our future work.

Some radio anomalies are highlighted in figure 11. For example, when the redirector is positioned at the mid-point between the source and destination nodes the ideal transceiver offers significant savings. However, in the case of the outdoor experiment using the Aironet PC4800 radio, positioning the redirector at the mid-point provides no power savings. Such anomalies are mainly the product of the operational granularity (i.e., the number of transmission power levels available) of the radio used.

Almost as important as having a larger set of transmission power levels is the manner in which these different levels are spaced with respect to each other. Transmission power levels for the Aironet PC4800 radio are exponentially spaced at 1, 5, 20, 50 and 100 mW. Therefore, the Aironet PC4800 is capable of using 1 mW for destinations that are a short distance from the transmitter, and 100 mW when the receiver is far away for example. Separating transmission levels in such an exponential fashion allows for a better approximation of the minimum transmission power at both near and far distances within the maximum transmission range of the transceiver. In contrast, a linear spacing of transmission power levels provides good accuracy only at either near or far distances but not at both near and far distances.

Regarding the delay involved in switching between transmission powers, ideally the transceiver should be capable of switching transmission power at the RTS–CTS time-scale. Whether or not this is possible in the future strongly depends on how much transmission power savings would improve the overall power consumption of a device, thus, motivating transceiver designers to improve this switching speed. For the PC4800 radio this delay is approximately 7 milliseconds. During this period the radio transceiver is neither capable of receiving nor transmitting packets. As we discussed in section 5, RTS–CTS packets need to be transmitted at maximum transmission power to guarantee the operation of the IEEE 802.11 MAC. This constraint means that PARO cannot fully operate using the Aironet radio because it is not possible to switch the transmission power between RTS–DATA

packets (transmitting node) nor CTS-ACK packets (receiver node) given the slow switching time. The only scenario where PARO could be deployed using current IEEE 802.11 Aironet PC4800 radios is in the case where the network operates at low traffic loads. In this case, the probability of a node finding hidden terminals while transmitting a packet would be small, thus RTS-CTS is unnecessary.

The Aironet PC4800 radios permit switching RTS-CTS mode ON and OFF depending on an RTS threshold. This threshold determines the minimum size of a transmitted data packet that requires the use of RTS-CTS. When the transmitted packet is equal to, or larger than the RTS threshold, an RTS packet is sent. This threshold ranges from 0 to 2400 bytes with a default value of 2048 bytes. The rationale behind this threshold is that the presence of hidden terminals is more disruptive for larger, rather than smaller data packets. This is because a transmitting node does not learn that a collision (due to hidden terminals) has occurred until the end of transmitting a data packet. Therefore, for larger data packets a transmitting node waits longer before it retransmits a packet. RTS-CTS packets are smaller in size and, if lost due to collision, can be retransmitted quickly with little overhead in comparison to data packets. The disadvantage of using RTS-CTS is that for each data packet transmitted that is larger than the threshold size, another packet must be transmitted and received, thereby reducing throughput.

The current RTS threshold does not relate to the network load and, therefore, it cannot be used to PARO's advantage. What is needed is a threshold that switches RTS-CTS on, when the traffic load is high, and off, when the traffic load is low; this operation is equivalent to switching PARO's operation off and on, respectively. A module implementing this functionality could take advantage of the number of collisions that data packets experience (which relates to network load) in order to switch the RTS-CTS mechanism on or off. It is important to note that because packets being forwarded between a source and destination may interfere with each other, switching RTS-CTS off may work only if the inter-packet delay is longer than the end-to-end delay of the path. Such a restriction is very limiting. As we discussed earlier the introduction of new MAC protocols such as Power Controlled Media Access Protocol (PCMAP) [15] can help overcome many of these limitations. We are studying how new MAC protocols can best offer the necessary dynamic power control support for PARO as part of our future work.

PARO proposes a cost function that makes the assumption that power consumption during the transmission mode is dominant and outweighs the collective power consumption during reception, idle and sleep modes. Therefore, in this work we only consider transmission power during data communication. We refer to a radio with these characteristics as an *ideal* radio. The full realization of an ideal radio is not possible because devices consume power during other radio operations. For example, some IEEE 802.11 radios have a power consumption of 1400 mW in the transmission mode, 1000 mW in the reception mode, 830 mW in the idle mode, and 130 mW during sleep mode [25]. There-

fore, when IEEE 802.11 radios are used, applying PARO route optimization has little impact on the resulting power savings of the network interface. PARO introduces redirectors between source-destination nodes that otherwise can communicate with each other directly. Introducing redirectors increases the number of times that a packet is received and transmitted before reaching its final destination. Until new radios are developed where the power consumption during reception is significantly smaller in comparison to power consumption during transmission then power optimization protocols such as PARO will show limited benefit.

8. Quality of Service

The main goal of PARO is to reduce the overall transmission power in the network in a simple, scalable manner. Using dynamic power control to accomplish this goal, however, impacts traditional QoS metrics such as throughput and end-to-end delay. While PARO is not designed to provide QoS assurances it is important to understand its impact on these performance metrics. Clearly, the introduction of one or more redirectors will have a negative impact on some of these metrics (e.g., end-to-end delays). In what follow, we discuss this impact. Many of the observations we make are more generally applicable to multihop radio systems.


8.1. Power control

The impact of transmission power control on network throughput has been widely studied in the context of cellular networks [5,20], and more recently for shared medium wireless ad hoc networks [1,16]. The later analysis typically focuses on the maximum capacity of the network as a function of the transmission range, node density, and the average distance between source-destination pairs. In [6] the authors show that the end-to-end throughput available to each node is $O(1/\sqrt{n})$ for random traffic patterns where n is the number of nodes. The main QoS tradeoff involved in a power controlled multihop ad hoc network has to do with the average number of times a packet is forwarded versus the average number of interfering nodes per attempted transmission. Increasing the transmission range reduces the number of times a packet needs to be forwarded en route to its final destination. However, increasing the transmission range increases the interference and therefore the channel contention every time a node attempts to transmit, thus increasing transmission delays. An opposite tradeoff applies when the transmission range is reduced. In [6] it is shown that reducing the transmission range is a better solution in terms of increasing the traffic carrying capacity of the network. The analysis presented in [6] only considers the physical capacity of the network, and not, the inefficiencies of the MAC protocol being used to transport data on top of the physical network. Unfortunately MAC protocols designed for cellular WLAN access may not be appropriate for multihop ad hoc operation.

In order to analyze how transmission power control impacts PARO we performed some experiments on a simple PARO “chain” network topology using CBR/UDP traffic. A chain network refers to a network where all the forwarding nodes are located over a straight line connecting source and destination nodes, as illustrated in figure 12(a). We use a chain network because it is easy to analyze its behavior and offers what we would expect to be the better results.

Figure 12(a) shows the simulation scenario of a simple chain network with the source (node 1) and destination (node 5) nodes set 200 meters apart and with three redirectors set 50 meters apart between them. The dashed line in the figure corresponds to the transmission range and the dotted line to the sensing range. Figures 12(b), (c) show simulation results for a varying number of redirectors between the source–destination nodes and packet sizes (viz. 64, 512 and 1500 bytes). In each case we locate redirectors at equal distances between the source–destination nodes. As figures 12(b) and (c) show the channel utilization drops sharply and the end-to-end delay increases as the number of redirectors increase. In the experiments illustrated in figure 12(a) all redirectors are separated by the same minimum transmission range from other adjacent redirectors in the chain network. This set up allows PARO to transmit both RTS–CTS and DATA packets using the same minimum transmission power level. In the more general case where redirectors are not located at the same transmission range, nodes are required to transmit RTS–CTS packets with the maximum transmission range to maintain MAC operations, thus reducing spectral reuse in the network. Therefore, the results shown in figures 12(b), (c) can be considered as the best case scenario, and in general, the performance of IEEE 802.11 based PARO networks would be worse. The number of redirectors used is shown in brackets (e.g., (3) indicates that 3 redirectors are used) in figures 12(b), (c).

Several factors contribute toward the observed degraded performance. It is widely known that IEEE 802.11 is not the best MAC for multihop ad hoc networks and results in lower throughput and increases in end-to-end delays seen by applications [2,14]. Decreasing the transmission range not only reduces the number of transmission opportunities that redirectors can use for their own transmissions, but every time a node attempts a transmission and senses the medium busy, it backs off for an exponentially increasing period of time after each failed attempt before trying again to transmit using the CSMA/CA access protocols. Referring to figure 12(a), when node 4 transmits to node 5 no other node in the network can transmit during that time. This is because when node 2 senses an ongoing transmission between node 4 and node 5, inhibiting node 1 from transmitting to node 2; that is, node 2 will not send a CTS after receiving an RTS from node 1. A similar situation occurs when nodes 1, 2 and 3 transmit. As a result the theoretical channel utilization of this simple chain network is 1/4 of the maximum capacity. In fact figure 12(b) shows that only 1/5 of the maximum throughput is achieved. Similar results are discussed in the literature [14].



An Example of a Simple Chain Network Showing 3 Redirectors for a Source (1) and Destination (5) Pair.

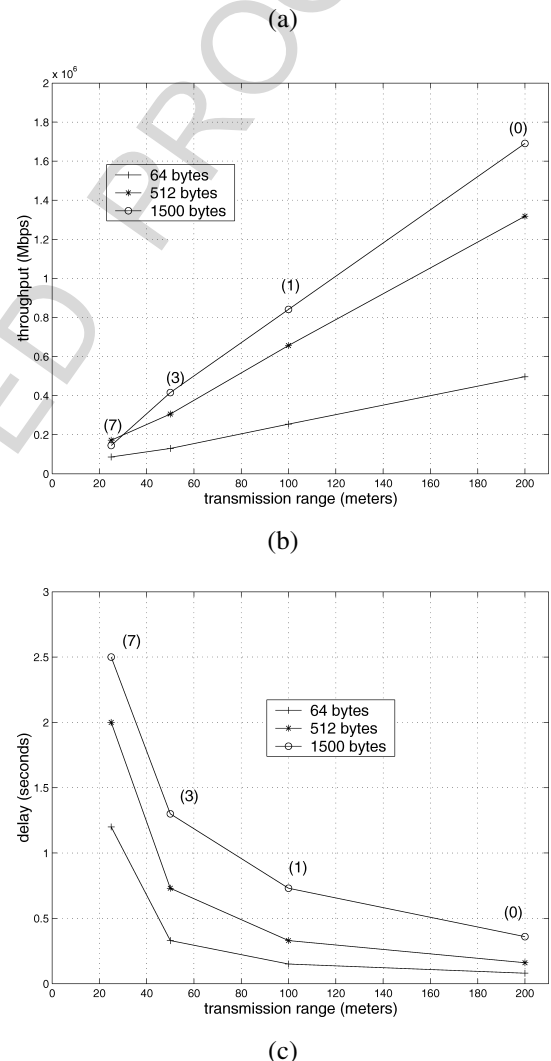


Figure 12. Throughput and delay performance of PARO.

8.2. Sensing and reception ranges

Sensing and transmitting ranges impact the performance of CSMA/CA based (e.g., IEEE 802.11) multihop ad hoc networks. The transmission range in CSMA/CA is defined as the maximum distance from the transmitter where an over-

Table 2
Sensing range (S_x)/reception range (R_x) ratio for the IEEE 802.11.

# redirectors	R_x (meters)	S_x (meters)	S_x/R_x
0	200	550	2.75
1	100	220	2.2
3	50	144	2.88
7	25	102	4.08

hearing node can still decode the received signal correctly. The sensing range, on the other hand, is the maximum distance from the transmitter where an overhearing node considers the channel busy, independent of whether or not this node can decode the received signal correctly. For WaveLAN IEEE 802.11 radios the transmitting and sensing range thresholds corresponds to $3.652e-10$ and $1.559e-11$ watts, respectively. This corresponds to transmitting and sensing range of 200 and 550 meters for a normal power transmission of 0.2818 watts, respectively.

An ongoing transmission by a node inhibits any other node transmitting within its sensing range. Table 2 shows the equivalent sensing range for a given transmission range using the carrier sense and reception thresholds for the WaveLAN IEEE 802.11 radio. Table 2 shows this for different numbers of redirector. In each case the transmission power is the minimum transmission power between adjacent redirectors. Table 2 also shows the ratio between sensing and transmission ranges (S_x/R_x). This ratio is a very important parameter not only for the performance of PARO but also for the performance of any multihop routing protocol (e.g., MANET). This is because whenever a forwarding node is actively transmitting it inhibits S_x/R_x other forwarding nodes from transmitting at the same time. A high S_x/R_x ratio limits the number of simultaneous transmissions along a given route, thus reducing the overall channel utilization. The S_x/R_x ratio does not remain constant but increases as the number of redirectors increases, as shown in table 2, thus reducing the overall channel utilization as the number of redirectors increase.

8.3. Spectral reuse

As discussed in section 3.1 most propagation models assume the strength of the received signal to be $\sim 1/d^\gamma$ fraction of the strength of the transmitted signal. The higher the value of γ the faster the signal strength decays with distance, and therefore the closer two transmitting nodes can be to each other without interfering with each other's transmissions. For indoor environments most propagation models assume an attenuation proportional to $\sim 1/d^4$, thus contributing toward a higher spectral reuse in the network. In the case of outdoor environments some propagations models consider an attenuation of $\sim 1/d^2$ or consider a two path loss model. The later model considers two regions: a first region where the signal attenuation is proportional to $\sim 1/d^2$ (inside the Fresnel zone), and a second region outside the Fresnel zone where the signal attenuation is proportional to $\sim 1/d^4$. Because the distances between redirectors in PARO are mostly within the

$\sim 1/d^2$ zone instead of the $\sim 1/d^4$ zone there is less spectral reuse when using PARO in outdoor environments.

Ideally, it is desirable to have the sensing range closer to the receiver range in order to increase spectral reuse in the network. This goal can be achieved by lowering the minimum signal to interference ratio (SIR) that a node can tolerate when receiving a packet correctly. However, such a change would increase the complexity of the hardware and similarly its cost.

8.4. Toward QoS-aware PARO

One interesting area of future work will consider building QoS support into PARO for applications that want to trade-off better QoS for suboptimal power savings. This could be achieved by simply limiting the number of redirectors introduced between a source-destination pair thereby achieving a certain throughput/delay objective.

The first QoS enhancement would allow sessions to control the number of redirectors used in a greedy fashion. In this scenario each source-destination session would introduce one redirector at a time, measure the end-to-end performance, and then decide based on the measured performance metrics, whether or not another redirector should be added to the route. This process would continue until the desired balance between QoS and power savings is achieved. A new flag could be inserted into the packet header to advertise whether or not each session is willing to add more redirectors along the path. A redirector introducing itself in the route could turn this flag OFF so no other node along the route attempts to become a redirector. A technical problem that would need to be solved is how to add redirectors in each route in a way such that they appear (as much as possible) to be located equal distances from each other. This condition is important because having redirectors too close to each other adversely impacts the performance, as discussed previously. Because each session makes local decisions about the introduction of redirectors this may result in an overall unstable solution where other sessions try to optimize their own constraints at the same time. For example, a certain session that introduces a new redirector along its path may impact the QoS performance of other active sessions in the vicinity, possibly triggering the continuous addition and removal of redirectors by other active sessions. There is a need to best understand these QoS and PARO tradeoffs and develop a suitable set of algorithms.

Another approach that looks promising is to study the concept of power controlled differentiate QoS. This approach as the name suggest provides different QoS performance to flows belonging to different service classes. This is very similar to the concept of differentiated services being discussed in the IETF. The simple idea is that a flow with fewer redirectors will obtain better QoS performance compared to flows having more redirectors in the path. Best effort traffic, for example, could introduce as many redirectors as the basic PARO protocol would support. In contrast, real-time traffic may limit the number of redirectors or add no redirectors at all.

9. Related work

Previous work in the area of power optimization in wireless networks has mainly focused on reducing the power of devices at the hardware level [8,9] or at the MAC level [3,23]. This goal is generally achieved by allowing devices to operate in low-power modes, sleeping during periods when no packets are destined for reception at a particular device.

Transmission power control in wireless networks has mainly addressed the control of the amount of interference that wireless devices operate in. In [20] work on joint power control between the base station and mobile devices determines the minimum transmission power for each mobile device for the uplink in a manner where the SNR thresholds for each communication link is met. In [5] microeconomics concepts and game theory are applied to power control in a distributed CDMA wireless system. In [15] transmission power control is used to improve the throughput capacity in a wireless packet network. In [19] power control is used to shape the topology of a multi-hop wireless network in a way that balances network-partitioning resilience versus spatial reuse.

In [13] a wireless ad-hoc network is divided into several clusters with a cluster-head responsible for handling most of the routing load in a power-efficient manner. In [7] micro sensor nodes use signal attenuation information to route packets towards a fixed destination known to all nodes in an energy-efficient way. In [27] different algorithms to discover energy-efficient broadcast and multicast trees are presented. Work presented in [24] uses a shortest-hop routing algorithm to discover the route with the lowest total cost among alternative paths from a source to a destination. The cost of each segment of the path is determined by the remaining lifetime of each forwarding node. The energy consumed in transmitting and receiving one packet over one hop is assumed to be constant in this work. In [4] an energy efficient routing protocol balances the traffic load in the network in order to maximize the lifetime of forwarding nodes.

A routing protocol addressing a similar problem space as PARO is discussed in [21]. In [21] wireless-enabled nodes discover energy-efficient routes to neighboring nodes and then use the shortest path Bellman-Ford algorithm to discover routes to any other node in the network. PARO differs from [21] in several ways. PARO devices do not rely on the availability of GPS to track the location of mobile nodes but uses signal attenuation to discover energy-efficient routes to neighboring nodes (i.e., those nodes located within the maximum transmission range). In addition, PARO does not only target finding energy-efficient routes as a goal. Rather, PARO attempts to achieve this goal using the minimum energy. Finally, PARO is designed to operate below standard layer 3 ad-hoc routing protocols to provide wide area coverage support in mobile environments.

Development of routing protocols capable of operating in wireless ad-hoc networks is the goal of the MANET working group in the IETF [11]. Little attention, however, has been placed on power conservation by the group. Rather, MANET routing protocols attempt to “minimize” the number

of intermediate hops (thereby minimizing delay) between any source-destination pair in the network [12,17,18]. MANET protocols are based on broadcast flooding schemes and, therefore, suffer of the same drawbacks as MLSR in order to discover power-efficient routes.

10. Conclusion

In this paper, we have presented PARO, a dynamic power controlled routing scheme 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 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 further 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.


Acknowledgements

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] S. Agarwal, S.V. Krishnamurthy, R.H. Katz and S. Dao, Distributed power control in ad hoc wireless networks, in: *Proceedings of PIMRC*, San Diego, CA (2001).
- [2] I. Ali, R. Gupta, S. Bansal, A. Misra, A. Razdan and R. Shorey, Energy efficiency and throughput for TCP traffic in multi-hop wireless networks, in: *Proceedings of InfoCom*, New York, NY (2002).
- [3] L. Bononi, M. Conty and L. Donatiello, PS-DCC: A distributed contention control mechanism leading to the power saving optimality in random access wireless local area networks, Technical Report (1999).
- [4] J.-H. Chang and L. Tassiulas, Energy conserving routing in wireless ad-hoc networks, in: *Proceedings of IEEE InfoCom*, Tel-Aviv, Israel (2000).
- [5] D. Goodman and N. Mandayam, Power control for wireless data, in: *Proceedings of IEEE International Workshop on Mobile Multimedia Communications (MOMUC)* (1999).
- [6] P. Gupta and P.R. Kumar, The capacity of wireless networks, *IEEE Transactions on Information Theory* 46(2) (2000).
- [7] W.R. Heinzelman, A. Chandrakasan and H. Balakrishnan, Energy-efficient routing protocols for wireless microsensor networks, in: *Proceedings of Hawaii International Conference on System and Sciences*, Hawaii (January 2000).
- [8] <http://www.bluetooth.com>
- [9] <http://www.homerf.org>
- [10] IEEE Vehicular Technology Society Committee on Radio Propagation, Coverage prediction for mobile radio systems operating in the 800/900 MHz frequency range 914 MHz, *IEEE Transactions on Vehicular Technology* 37(1) (1988).

- [11] IETF Mobile Ad-Hoc Network Working Group (MANET), <http://www.ietf.org/html.charters/manet-charter.html>
- [12] D.B. Johnson and D.A. Maltz, The dynamic source routing protocol for mobile ad hoc networks, draft-ietf-manet-dsr-02.txt (1999), work in progress.
- [13] T.J. Kwon and M. Gerla, Clustering with power control, in: *Proceedings of IEEE MILCOM*, Atlantic City, NJ (November 1999).
- [14] J. Li et al., Capacity of ad hoc wireless networks, in: *Proceedings of ACM MobiCom* (July 2002).
- [15] J. Monks, V. Bharghavan and W. Hwu, A Power Controlled Multiple Access Protocol (PCMAP) for wireless packet networks, in: *Proceedings of InfoCom*, Anchorage, AK (April 2001).
- [16] S. Narayanaswamy, V. Kawadia, R.S. Sreenivas and P.R. Kumar, Power control in ad-hoc networks: theory, architecture, algorithm and implementation of the COMPOW protocol, in: *Proceedings of European Wireless 2002*, Florence, Italy (February 2002).
- [17] C.E. Perkins and P. Bhagwat, Highly dynamic Destination-Sequenced Distance-Vector routing (DSDV) for mobile computers, *Proceedings of ACM SigComm 24(4)* (1994).
- [18] C.E. Perkins, E.M. Royer and S.R. Das, Ad hoc On-demand Distance Vector (AODV) routing, draft-ietf-manet-aodv-03.txt (1999), work in progress.
- [19] R. Ramanathan and R. Hain, Topology control of multihop wireless networks using transmit power adjustment, in: *Proceedings of IEEE InfoCom*, Tel-Aviv, Israel (2000).
- [20] F. Rashid-Farrokhi, K.J. Ray Liu and L. Tassiulas, Transmit beamforming and power control for cellular wireless systems, *IEEE Journal on Selected Areas in Communications* 16(8) (1996).
- [21] V. Rodoplu and T.H. Meng, Minimum energy mobile wireless networks, *IEEE Journal on Selected Areas in Communications* 17(8) (1999).
- [22] S.Y. Seidel and T.S. Rappaport, 914 MHz path loss prediction models for indoor wireless communications in multifloored buildings, *IEEE Transactions on Antennas and Propagation* 2(2) (1992).
- [23] S. Singh and C.S. Raghavendra, PAMAS: Power Aware Multi-Access protocol with Signalling for ad hoc networks, *ACM SigComm Computer Communication Review* (CCR) (July 1998).
- [24] S. Singh, M. Woo and C.S. Raghavendra, Power-aware routing in mobile ad hoc networks, in: *Proceedings of ACM/IEEE MobiCom*, Dallas, TX (October 1998).
- [25] M. Stemm and R.H. Katz, Measuring and reducing energy consumption of network interfaces in hand-held devices, *E80 B(8)* (August 2002).
- [26] The CMU Monarch Project, <http://www.monarch.cs.cmu.edu/>
- [27] J. Wieselthier, G. Nguyen and A. Ephremides, On the construction of energy-efficient broadcast and multicast trees in wireless networks, in: *Proceedings of IEEE InfoCom*, Tel-Aviv, Israel (2000).



Javier Gomez Castellanos obtained a B.S. degree with honors in electrical engineering in 1993 from National Autonomous University of Mexico (UNAM), Mexico, and M.S. degree in electrical engineering in 1996 from Columbia University, New York. Since 1996 he has been a Ph.D. student in the COMET Group at the Center for Telecommunications Research, Columbia University, New York. His research interests cover routing, QoS, and power-aware design for cellular and ad-hoc wireless networks.

E-mail: javierg@comet.columbia.edu



Andrew T. Campbell is an Associate Professor in the Department of Electrical Engineering, Columbia University, and a member of the COMET Group and the Columbia Networking Research Center (CNRC). His research interests encompass mobile computing and networking, open programmable networks, and QoS research. Andrew currently serves as technical program co-chair for the 8th ACM International Conference on Mobile Computing and Networking (ACM MobiCom 2002), and technical chair of the special track on networking technologies, services and protocols for IFIP Networking 2002. Andrew received his Ph.D. in computer science in 1996 and the NSF CAREER Award for his research in programmable mobile networking in 1999.

E-mail: campbell@comet.columbia.edu



Mahmoud Naghshineh is Director of emerging markets at IBM Technology Group. He is responsible for technology roadmap in the embedded space and defining software, services and solutions in support of Technology Groups offering. Prior to his current position, he was a Senior Manager at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, where he managed the Pervasive Computing Infrastructure Department and had world wide responsibility for IBM research projects in Pervasive Computing and directs the work of more than 50 researchers at the T.J. Watson Research Center. He manages projects in software and services infrastructure pervasive computing, mobile and wireless Internet, embedded software, secure platforms, network processor technologies, telecommunications services, and web-based payment infrastructure. He joined IBM in 1988. Prior to his current position, he has worked on communication and networking protocols, fast packet-switched/broadband IP networks, wireless and mobile networking optical networking, QoS provisioning, call admission, routing, and resource allocation, and network security. He has had several main technical contributions to IBM products in areas of networking technologies and software. He has contributed to IETF mobility protocols, IrDA, Bluetooth and IEEE 802.11 and 802.15 standards. He received his doctoral degree from Columbia University, New York, in 1994. He is a senior member of the IEEE. Aside from IBM, he has been very active in the external research and industrial for a. He is currently the Editor-in-Chief of IEEE Wireless Communications Magazine a leading magazine in mobile and wireless Internet. He has served as Program co-chair of MobiCom 2001 and as a technical editorial board member of many wireless and mobile networking/computing journals, as a member of technical program committee, session organizer and chairperson for many IEEE/ACM, NSF and Government conferences and workshops. Currently, he is an adjunct faculty member of the Department of Electrical Engineering at Columbia University teaching a graduate course on wireless and mobile networking. He has published over 100 technical papers and holds a number of IBM awards and patents.

E-mail: mahmoud@us.ibm.com



Chatschik Bisdikian was born in Thessaloniki, Greece. He holds a Ph.D. degree in electrical engineering from the University of Connecticut, Storrs, CT, USA. He is currently a Research Staff Member at IBM T.J. Watson Research Center, in Hawthorne, NY, USA. His interests include short-range wireless networks, wireless LAN deployment, service discovery, spontaneous networking, peer networking and other related areas, including playing the drums. He has been involved with the development of the Bluetooth specification from its early stages. He has served as vice-chair of the IEEE 802.15.1 task group that developed a standard for personal area networks. Dr. Bisdikian is a 1995 finalist for the Eta Kappa Nu Outstanding Young Electrical Engineer. He served in the editorial boards of several tech-

nical publications. He has guest edited various special issues in technical journals, including an issue on Wireless Personal Area Net-Works (IEEE Network Magazine, September/October 2001) and an issue on Ubiquitous Computing and Communications (IEEE Personal Communications Magazine, February 2000). He has authored over 85 technical, peer-reviewed papers, and holds six patents. He has the 2002 best tutorial award from IEEE

Communications Society for his paper titled An Overview of the Bluetooth Wireless Technology. He has appeared in several listings including in Whos Who in Science and Technology. He is a senior member of the IEEE, and a lifelong member of the Eta Kappa Nu, and Phi Kappa Phi Honor Societies. E-mail: bisdik@us.ibm.com

UNCORRECTED PROOF