

Energy Efficient Reliable Communication for Multi-hop Wireless Networks

Suman Banerjee
 Department of Computer Science
 University of Maryland
 College Park, MD 20742, USA
 Email: suman@cs.umd.edu

Archan Misra
 IBM T.J. Watson Research Center
 30 Saw Mill River Road
 Hawthorne, NY 10532, USA
 Email: archan@us.ibm.com

Abstract

Current algorithms for minimum-energy routing in wireless networks typically select minimum-cost multi-hop paths. In scenarios where the transmission power is fixed, each link has the same cost and the minimum-hop path is selected. In situations where the transmission power can be varied with the distance of the link, the link cost is higher for longer hops; the energy-aware routing algorithms select a path with a large number of small-distance hops. In this paper, we argue that such a formulation based solely on the energy spent in a single transmission is not able to find minimum energy paths for end-to-end reliable packet transmissions. The proper metric should include the total energy (including that expended for any retransmissions necessary) spent in reliably delivering the packet to its final destination.

We first study how link error rates affect this retransmission-aware metric, and how it leads to an efficient choice between a path with a large number of short-distance hops and another with a smaller number of large-distance hops. Such studies motivate the definition of a link cost that is a function of both the energy required for a single transmission attempt across the link and the link error rate. This cost function captures the cumulative energy expended in reliable data transfer, for both reliable and unreliable link layers. Through detailed simulations, we show that our schemes can lead to upto 30-70% energy savings over best known current schemes, under realistic environments.

Use of minimum energy paths for packet transmission may not always maximize the operational lifetime of multi-hop wireless networks. In the latter part of this paper, we define a new power-aware routing algorithm, called MRPC, that can be used to increase the operational lifetime of the network. MRPC integrates a node-specific metric, (residual battery power) and our proposed link-specific metric to select a path with the largest packet capacity at the ‘critical’ node. Finally, we also present, CMPRC, a conditional variant of MRPC that switches from minimum energy routing to MRPC only when the packet forwarding capacity of nodes falls below a threshold and through simulations we quantify the benefits of these schemes.

Index Terms

Low-power protocols, Wireless ad-hoc networks, Energy-efficiency, Reliable

I. INTRODUCTION

Multi-hop wireless networks typically possess two important characteristics:

- 1) The battery power available on the lightweight mobile nodes (such as sensor nodes or smart-phones) is relatively limited.
- 2) Communication costs in terms of transmission energy required are often much higher than computing costs on individual devices.

Energy-aware routing protocols (e.g., [19], [18], [2]) for such networks typically select routes that minimize the total transmission power aggregated over all nodes in the selected path.

If all nodes use the same constant transmission power, irrespective of the link distance, and if the links are error-free, then conventional minimum-hop routing (e.g. RIP [11] and OSPF [14]) will be most energy efficient. However,

This is an extended version of two papers[1] and [13] that appeared in MobiHoc’02 and WCNC’02 respectively.

minimum-hop solutions are not applicable in *variable-power* scenarios, where the nodes can dynamically vary their transmitter power levels. In such cases, greater energy-efficiency can be obtained if the nodes choose the transmission power depending the distance between the transmitter and receiver nodes.

For wireless links, a signal transmitted with power P_t over a link with distance D gets attenuated and is received with power

$$P_r \propto \frac{P_t}{D^K} \quad K \geq 2, \quad (1)$$

where K is a constant that depends on the propagation medium and antenna characteristics¹. Therefore, the transmission power for these links are chosen proportional to D^K . To calculate energy efficient paths, each link is assigned cost proportional to the energy required for a single transmission across the link. Minimum energy paths in this case will be paths with minimum aggregate path cost. Therefore, in these environments a path with a large number of small hops are typically chosen over an alternate path with a small number of large hops. This is the strategy used by a number of energy-efficient routing techniques, e.g., PAMAS [19], PARO [8].

In this paper, we discuss why such a formulation of the link cost fails to capture the actual energy spent in *reliable* packet delivery — a more accurate formulation needs to incorporate the link error rates to account for the potential cost of retransmissions needed for reliable packet delivery. Wireless links typically employ link-layer frame recovery mechanisms (e.g. link-layer retransmissions, or forward error correcting codes) to recover from packet losses. Additionally, protocols such as TCP or SCTP employ additional source-initiated retransmission mechanisms to ensure end-to-end reliability. Therefore, the energy cost associated with each candidate path should not merely reflect the energy spent in a single transmission attempt of a packet across the path, but rather the “total energy” spent in packet delivery. This energy cost should include the energy spent in potential retransmissions as well².

A. Minimum Energy Paths

We first consider how the error rate of individual links affects the energy required to ensure reliable packet delivery. As part of this analysis, we consider two different operating models:

- 1) **End-to-End Retransmissions (EER)**: where the individual links do not provide link-layer retransmissions and error recovery— reliable packet transfer is achieved only via retransmissions initiated by the source node.
- 2) **Hop-by-Hop Retransmissions (HHR)**: where each individual link provides reliable forwarding to the next hop using localized packet retransmissions.

In both cases, the choice of links with relatively high error rates leads to packet re-transmissions to guarantee reliable delivery. This will significantly increase the effective energy spent in reliable packet transmission. This is true in both the constant-power and variable-power scenarios — in either scenario, ignoring the error rate of the link can lead to the selection of paths with high error rates and consequently, high retransmission overhead. The analysis of the effects of link error rates on the effective energy consumption is more interesting for the variable-power case: we show that the choice between a path with many short-range hops and another with fewer long-range hops involves a *trade-off between the reduction in the transmission energy for a single packet and the potential increase in the frequency of retransmissions*. Our analysis of the variable-power scenarios shows that schemes which consider the link-error rates would perform significantly better than currently proposed minimum-energy routing protocols, which do not.

Based on the above analytic observations, we define a new link cost metric that captures the effects of both the link distance and the link error rate. Minimum cost route computations based on this metric leads to energy efficient paths for reliable communication for both EER and HHR scenarios. We show that such a link cost can be exactly defined to obtain optimal solutions only for the HHR scenario; for the EER framework, we can only devise an approximate cost function. By using simulation studies, we also demonstrate how the choice of parameters in the approximate EER cost formulation represents a tradeoff between energy efficiency and the achieved throughput. Since most decentralized ad-hoc routing protocols (e.g., AODV [16], DSR [9], TORA [15]) attempt, at least approximately, to select a minimum-cost path (where the path cost is a sum of the individual link costs), our formulation can be easily applied to these protocols.

¹ K is typically around 2 for short distances and omni-directional antennae, and around 4 for longer distances.

² This is especially relevant in multi-hop wireless networks, where variable channel conditions often cause packet error rates as high as 15–25%.

While the link quality has been previously suggested as a routing metric to reduce queuing delays and loss rates, its implicit effect on the energy efficiency has not been studied before. By incorporating the link error rates in the link cost, we show that 30% to 70% energy savings can often be achieved under realistic operating conditions.

B. Maximizing Network Lifetime

While minimum energy paths provide the most efficient way to transfer *individual* packets between the source and destination, in many ad-hoc wireless scenarios, the metric of actual interest is not this energy requirements to transfer individual packets but the total operational lifetime of the network. To avoid the extinction of nodes due to exhaustion of their battery power, some power-aware routing algorithms [20], [22] try to balance the depletion of battery-power at the different nodes. It is easy to see that the two routing objectives — minimizing energy requirements to transfer individual packets and maximizing the lifetime of a network, can be mutually contradictory. Strategies to maximize network lifetime typically take into account the variable traffic volume passing through different nodes and avoid rapid battery depletion at a few ‘unlucky’ node.

In the latter part of this paper we leverage our minimum energy path selection technique to define a new power-aware route selection protocol called Maximum Residual Packet Capacity (MRPC) that can be used to increase the operational lifetime of the network. In this protocol the choice of the routes are based both on node-specific parameters (e.g. the residual battery power at a node) and link specific metrics (as defined to compute minimum energy paths). MRPC is conceptually similar to the MMBCR algorithm [22] in that, at any point in time, it tries to select the route that maximizes the residual capacity currently available at the most critical node (the one with the least residual packet forwarding capacity). MRPC accommodates scenarios where the nodes can adjust their transmission power dynamically (based on the distance between the nodes), and also incorporates the effect of link layer error rates and consequent packet re-transmissions.

In MRPC, the cost of choosing a particular link at any instant is defined as the *idealized maximum number of packets* that can be transmitted by the transmitting node over the specific link, assuming the complete absence of any other cross traffic at that node. We use simulation studies to show how MRPC leads to superior performance (longer network lifetimes or larger number of successfully transmitted packets) than alternative suggested algorithms.

Since minimum-energy routes are more energy efficient, a conditional variant of the MMBCR algorithm was also proposed in [22]. In this scheme (called Conditional MMBCR, or CMMBCR) minimum energy routes were chosen till the residual battery power of constituent nodes of the routes fell below a specified threshold level. Once this threshold level is crossed, routes are chosen using the MMBCR algorithm, which equitably distributes the battery consumption among the different nodes thus protecting against the early exhaustion of a few nodes. We also present the conditional analogue of MRPC, the Conditional MRPC (CMRPC) algorithm and then evaluate its performance, vis-a-vis the CMMBCR algorithm. CMRPC performs minimum energy routing (using our link cost formulation for reliable transmissions) as long as the remaining battery power at the constituent nodes lie above a specified threshold. Beyond this point, CMRPC switches to the MRPC-based max-min path selection algorithm.

C. Roadmap

The rest of the paper is organized as follows. Section II provides an overview of previous related work. Section III formulates the effective transmission energy problem as a function of the number of hops, and the error rates of each hop, for both the EER and HHR case and analyses its effect on the optimum number of hops in the variable-power scenario. It also demonstrates the agreement between our idealized energy computation and real TCP behavior. Section IV shows how to form link costs that lead to the selection of minimum energy paths. In Section V we present the results of simulation studies to study the performance of our minimum-cost algorithms for both the fixed-power and the variable-power scenarios. In Section VI we describe our MRPC algorithm for maximizing the operational lifetime of a network, and subsequently present simulation results for this algorithm (and its conditional variant) in Section VII. Finally, we conclude in Section VIII.

II. RELATED WORK

Metrics used by conventional routing protocols for the wired Internet typically do not need to consider any energy-related parameters. Thus, RIP [11] uses hop count as the sole route quality metric, thereby selecting minimum-hop paths between the source and destinations. OSPF [14], on the other hand, can support additional link metrics such as available bandwidth, link propagation delay etc. There is, however, no well-defined support for using link-error rates as a metric in computing the shortest cost path. Clearly, in fixed-power scenarios, the minimum-hop path would also correspond to the path that uses the minimum total energy for a single transmission of a packet.

In contrast, energy-aware routing protocols for variable-power scenarios aim to directly minimize the total power consumed over the entire transmission path. PAMAS [19] is one such minimum total transmission energy protocol, where the link cost is set to the transmission power and Dijkstra's shortest path algorithm is used to compute the path that uses the smallest cumulative energy. In the case where nodes can dynamically adjust their power based on the link distance, such a formulation often leads to the formation of a path with a large number of hops. A link cost that includes the receiver power as well is presented in [18]. By using a modified form of the Bellman-Ford algorithm, this approach resulted in the selection of paths with smaller number of hops than PAMAS.

Routing protocols for wireless ad-hoc environments (e.g. AODV [16], DSR [9]) contain special features to reduce the signaling overheads and convergence problems caused by node mobility and link failures. While these protocols do not necessarily compute the absolute minimum-cost path, they do aim to select paths that have lower cost (in terms of metrics such as hop count or delay). Such protocols, can in principle be adapted to yield energy-efficient paths simply by setting the link metric to be a function of the transmission energy. In contrast, other ad-hoc routing protocols have been designed specifically to minimize transmission energy cost. For example, the Power-Aware Route Optimization (PARO) algorithm [8], [7] is designed for scenarios where the nodes can dynamically adjust their transmission power — PARO attempts to generate a path with a larger number of short-distance hops. According to the PARO protocol, a candidate intermediary node monitors an ongoing direct communication between two nodes and evaluates the potential for power savings by inserting itself in the forwarding path— in effect, replacing the direct hop between the two nodes by two smaller hops through itself.

Minimum-energy path computation algorithms can sometimes unfairly penalize a subset of the nodes, e.g., if several minimum energy routes have a common node in the path, the battery of that node will be exhausted quickly. Therefore researchers in energy-aware routing have also considered an alternate objective function — maximizing the lifetime of the network. In some of the proposed mechanisms, the lifetime of a network can be extended by equitably distribute the energy expenditure among the different nodes of the network. The resulting end-to-end paths consist of nodes which currently possess greater battery resources. Among such battery-aware algorithms, [20] formulated a node metric, where the capacity of each node was a decreasing function of the residual battery capacity. A minimum cost path selection algorithm then helps to steer routes away from paths where many of the intermediate nodes are facing battery exhaustion.

The MMBCR algorithm [22] uses a min-max route selection technique. The algorithm chooses that route that has the largest value for its most critical (“bottleneck”) node—i.e., the node with the least residual battery capacity. The MMBCR algorithm never tries to minimize the total transmission energy along a path. It is *always* concerned with spreading the transmission cost evenly among available nodes. Therefore, it can lead to overall higher energy consumption and consequently, a reduction of the average node lifetime. To counteract this, a conditional variant, called CMMBCR [22], uses a low energy path for packet transmission, as long as the battery power level on all the nodes in the selected path lie above a certain threshold γ . However, this low energy path computation does not take into account energy costs due to packet retransmissions. Once one or more of nodes on all possible paths falls below the battery protection threshold, the algorithm switches to the MMBCR mode.

The effect of link quality on packet transmissions across individual links have been studied in the literature. The focus of some of this study has been on the use of intelligent *link scheduling* algorithms, rather than on transmission power control [23], [17], [4]. In another approach, Gass et. al. [6] have proposed a transmission power adaptation scheme to control the link quality of individual frequency-hopping wireless links.

All these protocols and algorithms, do not, however, consider the effect of the link error rates on the overall number of retransmissions, and thus the energy needed for end-to-end reliability in packet delivery. Our problem formulation and route computation explicitly accounts for potential retransmissions of packets along the entire path in calculating the energy costs. To the best of our knowledge, our work defines the first known technique that addresses the energy efficient routing in multi-hop wireless networks from a reliable packet delivery standpoint. In this work, we assume that each node in the ad-hoc network is aware of the packet error link on its outgoing links. Sensing the channel noise conditions can be done either at the link layer, a capability that is built into most commercial wireless 802.11 interfaces available today, or through higher layer mechanisms such as periodic packet probes or aggregated packet reception reports from the receiver³.

III. ENERGY COST ANALYSIS AND MINIMUM ENERGY PATHS

In this section, we demonstrate how the error rate associated with a link affects a) the overall probability of reliable delivery, and consequently, b) the energy associated with the reliable transmission of a single packet. For any particular link $\langle i, j \rangle$ between a transmitting node i and a receiving node j , let $T_{i,j}$ denote the transmission power and $p_{i,j}$ represent the packet error probability. Assuming that all packets are of a constant size, the energy involved in a packet transmission, $E_{i,j}$, is simply a fixed multiple of $T_{i,j}$.

Any signal transmitted over a wireless medium experiences two different effects: attenuation due to the medium, and interference with ambient noise at the receiver. Due to the characteristics of the wireless medium, the transmitted signal suffers an attenuation proportional to D^K , where D is the distance between the receiver and the transmitter. The ambient noise at the receiver is independent of the distance between the source and distance, and depends purely on the operating conditions at the receiver. The bit error rate associated with a particular link is essentially a function of the ratio of this received signal power to the ambient noise. In the constant-power scenario, $T_{i,j}$ is independent of the characteristics of the link $\langle i, j \rangle$ and is a constant. In this case, a receiver located farther away from a transmitter will suffer greater signal attenuation (proportional to D^K) and will, accordingly, be subject to a larger bit-error rate. In the variable-power scenario, a transmitter node adjusts $T_{i,j}$ to ensure that the strength of the (attenuated) signal received by the receiver is *independent of D* and is above a certain threshold level Th . Accordingly, the optimal transmission power associated with a link of distance D in the variable-power scenario is given by:

$$T_{opt} = Th \times \gamma \times D^K, \quad (2)$$

where γ is a proportionality constant and K is the coefficient of attenuation ($K \geq 2$). Since Th is typically a technology-specific constant, we can see that the optimal transmission energy over such a link varies as:

$$E_{opt}(D) \propto D^K. \quad (3)$$

If links are considered error-free, then minimum hop paths are the most energy-efficient for the fixed-power case. Similarly, in the absence of transmission errors, paths with a large number of small hops are typically more energy efficient in the variable power case. However in the presence of link errors none of the above choices may give optimal energy efficient paths in the two cases. We now analyze the interesting consequences of this behavior for the variable-power scenario (for both the EER and HHR cases); The analysis for the fixed-power scenario is simpler, and is a special case of the variable-power scenario.

A. Optimal Minimum Energy Paths in EER Case

In the EER case, a transmission error on any link leads to an end-to-end retransmission over the path. Given the variable-power formulation of E_{opt} in Equation 3, it is easy to see why placing an intermediate node along the straight line between two adjacent nodes (breaking up a link of distance D into two shorter links of distance D_1 and D_2 such

³Similar ideas were proposed for link sensing in the Internet MANET Encapsulation Protocol [3].

that $D_1 + D_2 = D$) always reduces the total E_{opt} . In fact, PARO [7] works using precisely such an estimation. From a reliable transmission energy perspective, such a comparison is inadequate since it does not include the effect on the overall probability of error-free reception.

To understand the energy-tradeoff involved in choosing a path with multiple short hops over one with a single long hop, consider communication between a sender (S) and a receiver (R) separated by a distance D . Let N represent the total number of hops between S and R , so that $N - 1$ represents the number of forwarding nodes between the endpoints. For notational ease, let these nodes be indexed as $i : i = \{2, \dots, N\}$, with node i referring to the $(i - 1)^{th}$ intermediate hop in the forwarding path; also, node 1 refers to S and node $N + 1$ refers to R . In this case, the total optimal energy spent in simply transmitting a packet once (without considering whether or not the packet was reliably received) from the sender to the receiver over the $N - 1$ forwarding nodes is:

$$E_{total} = \sum_{i=1}^N E_{opt}^{i,i+1}, \quad (4)$$

or, on using Equation 3,

$$E_{total} = \sum_{i=1}^N \alpha D_{i,i+1}^K, \quad (5)$$

where $D_{i,j}$ refers to the distance between nodes i and j and α is a proportionality constant. To understand the transmission energy characteristics associated with the choice of $N - 1$ intermediate nodes, we compute the lowest possible value of E_{total} for any given layout of $N - 1$. Using very simple symmetry arguments, it is easy to see that the minimum transmission energy case occurs when each of the hops are of equal length D/N . In that case, E_{total} is given by:

$$E_{total} = \sum_{i=1}^N \alpha \frac{D^K}{N^K} = \frac{\alpha D^K}{N^{K-1}} \quad (6)$$

For computing the energy spent in *reliable delivery*, we now consider how the choice of N affects the the probability of transmission errors and the consequent need for retransmissions. Clearly, increasing the number of intermediate hops the likelihood of transmission errors over the entire path.

Assuming that each of the N links has an independent packet error rate of p_{link} , the probability of a transmission error over the entire path, denoted by p , is given by

$$p = 1 - (1 - p_{link})^N \quad (7)$$

The number of transmissions (including retransmissions) necessary to ensure the successful transfer of a packet between S and D is then a geometrically distributed random variable X , such that

$$\text{Prob}\{X = k\} = p^{k-1} \times (1 - p), \quad \forall k$$

The *mean* number of individual packet transmissions for the successful transfer of a single packet is thus $1/(1 - p)$. Since each such transmission uses total energy E_{total} given by Equation 6, the total expected energy required in the reliable transmission of a single packet is given by:

$$\begin{aligned} E_{total\ rel}^{EER} &= \alpha \frac{D^K}{N^{K-1}} \cdot \frac{1}{1 - p} \\ &= \frac{\alpha D^K}{N^{K-1}(1 - p_{link})^N} \end{aligned} \quad (8)$$

The equation clearly demonstrates the effect of increasing N on the total energy necessary; while the term N^{K-1} in the denominator increases with N , the error-related term $(1 - p_{link})^N$ decreases with N . By treating N as a continuous




Fig. 1. Total Energy Costs vs. Number of Forwarding Nodes (EER)

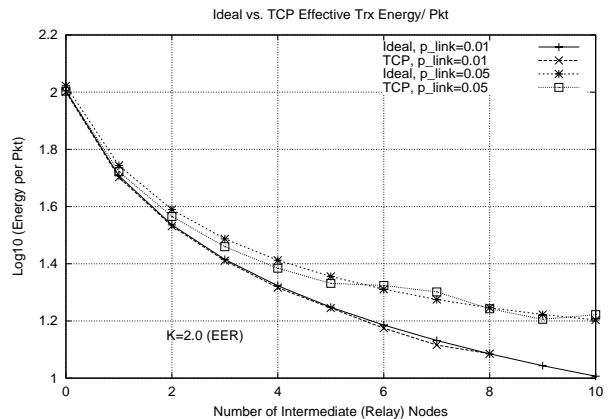


Fig. 2. Idealized / TCP Energy Costs vs. Number of Forwarding Nodes (EER)

variable and differentiating, it follows that the optimal value of the number of hops, N_{opt} is given by:

$$N_{opt} = \frac{(K - 1)}{-\log(1 - p_{link})}$$

Thus a larger value of p_{link} corresponds to a smaller value for the optimal number of intermediate forwarding nodes. Also, the optimal value for N increases linearly with the attenuation coefficient K . There is thus clearly an optimal value of N ; while lower values of N do not exploit the potential reduction in the transmission energy, higher values of N cause the overhead of retransmissions to dominate the total energy budget.

To study these tradeoffs graphically, we plot $E_{total\ rel}^{EER}$ against varying N (for different values of p_{link}) in Figure 1. For this graph, α and D (which are really arbitrary scaling constants) in the analysis are kept at 1 and 10 respectively and $K = 2$. The graph shows that for low values of the link error rates, the probability of transmission errors is relatively insignificant; accordingly, the presence of multiple short-range hops nodes leads to a significant reduction in the total energy consumption. However, when the error rates are higher than around 10%, the optimal value of N is fairly small; in such scenarios, any potential power savings due to the introduction of an intermediate node are negated by a sharp increase in the number of transmissions necessary due to a larger effective path error rate. *In contrast to earlier work, our analysis shows that a path with multiple shorter hops is thus not always beneficial than one with a smaller number of long-distance hops.*

1) *Energy Costs for TCP Flows:* Our formulation (Equation 8) provides the total energy consumed per packet using an ideal retransmission mechanism. TCP's flow control and error recovery algorithms could potentially lead to different values for the energy consumption, since TCP behavior during loss-related transients can lead to unnecessary retransmissions. While the effective TCP throughput (or goodput) as a function of the end-to-end loss probability has been derived in several analyses (see [10], [5]), there exists no model to predict the total number of packet transmissions (including retransmissions) for a TCP flow subject to a variable packet loss rate. We thus use simulation studies using the *ns-2* simulator⁴, to measure the energy requirements for reliable TCP transmissions. Figure 2 plots the energy consumed by a persistent TCP flow, as well as the ideal values computed using Equation 8, for varying N and for $p_{link} = \{0.01, 0.05\}$. We observe good agreement between our analytical predictions and TCP-driven simulation results. This verifies the practical utility of our analytical model.

B. Optimal Minimum Energy Paths in HHR Case

In the case of the HHR model, a transmission error on a specific link implies the need for retransmissions on that link alone. This is a better model for multi-hop wireless networking environments, since wireless link layers typically

⁴ Available at <http://www.isi.edu/nsnam/ns>

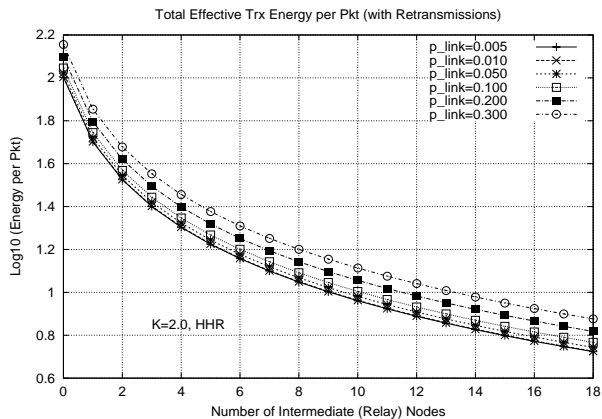


Fig. 3. Total Energy Costs vs. Number of Forwarding Nodes (HHR)

employ link-layer retransmissions. In this case, the link layer retransmissions on a specific link ensure that the transmission energy spent on the other links in the path is *independent* of the error rate of that link. For our analysis, we do not bound the maximum number of permitted retransmissions: a transmitter continues to retransmit a packet until the receiving node acknowledges error-free reception. (Clearly, practical systems would typically employ a maximum number of retransmission attempts to bound the forwarding latency.) Since our primary focus is on energy-efficient routing, we also do not explicitly consider the effect of such retransmissions on the overall forwarding latency of the path in this paper.

Since the number of transmissions on each link is now *independent of the other links* and is geometrically distributed, the total energy cost for the HHR case is

$$E_{\text{total rel}}^{\text{HHR}} = \sum_{i=1}^N \alpha \frac{D_{i,i+1}^k}{1 - p_{i,i+1}} \quad (9)$$

In the case of N intermediate nodes, with each hop being of distance $\frac{D}{N}$ and having a link packet error rate of p_{link} , this reduces to:

$$E_{\text{total rel}}^{\text{HHR}} = \alpha \frac{D^K}{N^{K-1} * (1 - p_{\text{link}})} \quad (10)$$

Figure 3 plots the total energy for the HHR case, for $K = 2$ and different values of N and p_{link} . In this case, it is easy to see that the *total energy required always decreases with increasing N , following the $\frac{1}{N^{K-1}}$ asymptote*. The logarithmic scale for the energy cost compresses the differences in the value of $E_{\text{total rel}}^{\text{HHR}}$ for different p_{link} . If all links have the same error rate, it would therefore, be beneficial to substitute a single hop with multiple shorter hops.

In Figure 3 we can also observe the effect of increasing value of p_{link} for a fixed N . As expected, a higher link error rate leads to larger number of re-transmissions and a higher energy consumption. It is important to note that the effect of increasing link error rates is more significant in the EER case — in Figure 1, for $N = 10$, increasing the loss probability from 0.1 to 0.2 can increase the energy consumption ten-fold.

A comparison of the energy consumption in the EER and HHR cases for identical values of N and K shows that the energy consumption in the EER case is at least an order of magnitude larger, for even moderate values of the link error rate. By avoiding the end-to-end retransmissions, the HHR approach can significantly lower the total energy consumption. These analyses reinforce the requirements of link-layer retransmissions in any radio technology used in multi-hop, ad-hoc wireless networks.

IV. ASSIGNING LINK COSTS FOR MINIMUM ENERGY RELIABLE PATHS

In contrast to traditional Internet routing protocols, energy-aware routing protocols typically compute the shortest-cost path, where the cost associated with each link is some function of the transmission (and/or reception) energy associated with the corresponding nodes. To adapt such minimum cost route determination algorithms (such as Dijkstra's or the Bellman-Ford algorithm) for energy-efficient reliable routing, the link cost must now be a function of not just the associated transmission energy, but the link error rates as well. Using such a metric would allow the routing algorithm to select links that present the optimal tradeoff between low transmission energies and low link error rates. As we shall shortly see, defining such a link cost is possible only in the HHR case; approximations are needed to define suitable cost metrics in the EER scenario.

Consider a graph, with the set of vertices representing the communication nodes and a link $\langle i, j \rangle$ representing the direct hop between nodes i and j . For generality, assume an asymmetric case where $\langle i, j \rangle$ is not the same as $\langle j, i \rangle$; moreover, $\langle i, j \rangle$ refers to the link used by node i to transmit to node j . A link is assumed to exist between node pair $\{i, j\}$ as long as node j lies within the transmission range of node i . This transmission range is uniquely defined for the constant-power case; for the variable-power case, this range is really the *maximum permissible range* corresponding to the maximum transmission power of a sender. Let $E_{i,j}$ be the energy associated with the transmission of a packet over link $l_{i,j}$, and $p_{i,j}$ be the link packet error probability associated with that link. (In the fixed-power scenario, $E_{i,j}$ is independent of the link characteristics; in the variable-power scenario, $E_{i,j}$ is a function of the distance between nodes i and j .) Now, the routing algorithm's job is to compute the shortest path from a source to the destination that minimizes the sum of the energy costs over each constituent link.

A. Hop-by-Hop Retransmissions (HHR)

Consider a path P from a source node S (indexed as node 1) to node D (indexed as node $N + 1$) that consists of $N - 1$ intermediate nodes indexed as $2, \dots, N$. Then, choosing path P for communication between S and D implies that the total energy cost is given by:

$$E_P = \sum_{i=1}^N \frac{E_{i,i+1}}{1 - p_{i,i+1}} \quad (11)$$

Choosing a minimum-cost path from node 1 to node $N + 1$ is thus equivalent to choosing the path P that minimizes Equation 11. It is thus easy to see that the corresponding link cost for link $l_{i,j}$, denoted by $C_{i,j}$, is given by:

$$C_{i,j} = \frac{E_{i,j}}{1 - p_{i,j}} \quad (12)$$

Well known ad-hoc routing protocols, such as AODV, DSR and TORA, can use this link cost to compute the appropriate energy-efficient routes. Some of the existing energy-efficient routing techniques, e.g. PARO, can also be easily adapted to use this new link cost formulation to compute minimum-energy routes. Thus, in such a modified version of the PARO algorithm, an intermediate node C would offer to interject itself between two nodes A and B if the sum of the link costs $C_{A,C} + C_{C,B}$ was less than the 'direct' link cost $C_{A,B}$.

B. End-to-End Retransmissions (EER)

In the absence of hop-by-hop retransmissions, the total energy cost along a path contains a multiplicative term involving the packet error probabilities of the individual constituent links. In fact, assuming that transmission errors on a link do not stop downstream nodes from relaying the packet, the total transmission energy can be expressed as :

$$E_P = \frac{\sum_{i=1}^N E_{i,i+1}}{\prod_{i=1}^N (1 - p_{i,i+1})} \quad (13)$$

Given this form, the total cost of the path cannot be expressed as a linear sum of individual link costs⁵, thereby making the exact formulation inappropriate for traditional minimum-cost path computation algorithms. We therefore concentrate on alternative formulations of the link cost, which allow us to use conventional distributed shortest-cost algorithms to compute “approximate” minimum energy routes.

A study of Equation 13 shows that using a link with a high p can be very detrimental in the EER case: an error-prone link effectively drives up the energy cost for all the nodes in the path. Therefore, a useful heuristic function for link cost should have a super-linear increase with increase in link error rate; by making the link cost for error-prone links prohibitively high, we can ensure that such links are usually excluded during shortest-cost path computations.

In particular, for a path consisting of k identical links (i.e. have the same link error rate and link transmission cost), Equation 13 will reduce to

$$E_P = \frac{kE}{(1-p)^k} \quad (14)$$

where, p is the link error rate and E is the transmission cost across each of these links. This leads us to propose a heuristic cost function for a link, as follows:

$$C_{i,j}^{approx} = \frac{E_{i,j}}{(1-p_{i,j})^L} \quad (15)$$

where $L = 2, 3, \dots$, and is chosen to be identical for all links⁶. Clearly, if the exact path length is known and all nodes on the path have the identical link error rates and transmission costs, L should be chosen equal to that path length. However, in accordance with current routing schemes, we require that a link should associate only a single link cost with itself, irrespective of the lengths of specific routing paths that pass through it. Therefore, we need to fix the value of L , independent of the different paths that cross a given link. If better knowledge of the network paths are available, *then L should be chosen to be the average path length of this network*. Higher values of L impose progressively stiffer penalties on links with non-zero error probabilities.

Given this formulation of the link cost, the minimum-cost path computation effectively computes the path with the minimum “approximate” energy cost given by:

$$E_P \sim \sum_{i=1}^N \frac{E_{i,i+1}}{(1-p_{i,i+1})^L} \quad (16)$$

As before, protocols like AODV, DSR, TORA and PARO can use this new link cost function $C_{i,j}^{approx}$ (Equation 15) to make their routing decisions.

As with our theoretical studies in Section III, the analysis here does not directly apply to TCP-based reliable transport, since TCP’s loss recovery mechanism can lead to additional transients. In the next section, we shall use simulation-based studies to study the performance of our suggested modifications to the link cost metric in typical ad-hoc topologies.

V. PERFORMANCE EVALUATION: MINIMUM ENERGY PATHS

The analysis of the previous section provides a foundation for devising energy-conscious protocols for reliable data transfer. In this section, we report on simulation-based studies that examine the performance of our suggested techniques for computing energy efficient reliable communication paths. We performed our simulations in the ns-2 simulator. We experimented with different types of traffic sources:

⁵We do not consider solutions that require each node or link to separately advertise two different metrics. It is possible to define optimal energy efficient paths if nodes distributed two separate metrics — a) $E_{i,j}$ and b) $\log(1-p_{i,j})$. The cumulative values $\sum E_{i,j}$ and $\sum \log(1-p_{i,j})$ can be used by nodes to compute such optimal paths.

⁶There should be an L factor in the numerator too (as in Equation 14, but since this is identical for all links, it can be ignored).

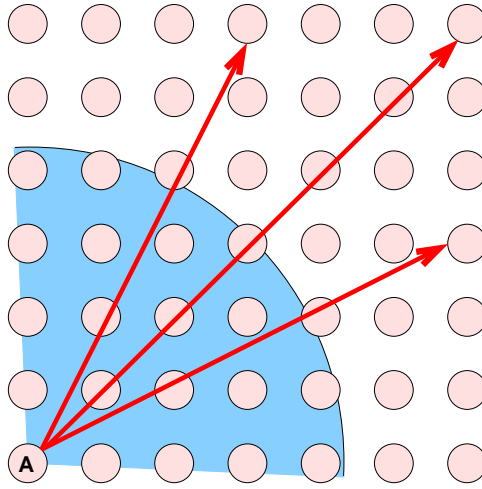


Fig. 4. The 49-node topology. The shaded region marks the maximum transmission range for the corner node, A . There are three flows from each of the 4 corner nodes, for a total of 12 flows.

- 1) For studies using the EER framework, we used TCP flows implementing the NewReno version of congestion control.
- 2) For studies using the HHR framework, we used both UDP and TCP flows. In UDP flows, packets are inserted by the source at regular intervals.

To study the performance of our suggested schemes, we implemented and observed three separate routing algorithms:

- 1) The minimum-hop routing algorithm, where the cost of all links is identical and independent of both the transmission energy and the error rate.
- 2) The Energy-Aware (**EA**) routing algorithm, where the cost associated with each link is the energy required to transmit a single packet (without retransmission considerations) across that link.
- 3) Our Retransmission-Energy Aware (**RA**) algorithm, where the link cost includes the packet error rates, and thus considers the impact of retransmissions necessary for reliable packet transfer. For the HHR scenario, we use the link cost of Equation 12; for the EER model, we use the ‘approximate’ link cost of Equation 15 with $L = 2$. In Section V-D.2, we also study the effect of varying the L -parameter.

In the fixed-power scenario, minimum-hop routing creates the most energy efficient paths if link error rates are ignored. Therefore, minimum-hop and EA algorithms exhibit identical behavior in this case. Accordingly, it suffices to compare our RA algorithm with minimum-hop routing alone. For our experiments, we used different topologies having upto 100 nodes randomly distributed on a square region, to study the effects of various schemes on energy requirements and throughputs achieved. In this section, we discuss in detail results from one representative topology, where 49 nodes were distributed over a 70X70 unit grid, equi-spaced 10 units apart (Figure 4). The maximum transmission radius of a node is 45 units, which implies that each node has between 14 and 48 neighbors on this topology,

Each of the routing techniques were then run on these static topologies to derive the least-cost paths to each destination node. To simulate the offered traffic load typically of such ad-hoc wireless topologies, each of the corner node on the grid had 3 active flows, providing a total of 12 flows. Since our objective was to study the transmission energies alone, we did not consider other factors such as link congestion, buffer overflow etc. Thus, each link had an infinitely larger transmit buffer; the link bandwidths for all links (point to point) was set to 11 Mbps. Each of the simulations was run for a fixed duration.

A. Modeling Link Errors

The relation between the bit-error-rate (p_b) over a wireless channel and the received power level P_r is a function of the modulation scheme. However, in general, several modulation schemes exhibit the following generic relationship between p_b and P_r :

$$p_b \propto \operatorname{erfc}\left(\sqrt{\frac{\text{constant} \times P_r}{N \times f}}\right)$$

where N is the noise spectral density (noise power per Hz), f is the raw channel bit-rate and $\operatorname{erfc}(x)$ is defined as the complementary function of $\operatorname{erf}(x)$ and is given by

$$\operatorname{erfc}(x) = 1 - \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$$

As specific examples, the bit error rate is given by $p_b = \operatorname{erfc}\left(\sqrt{\frac{P_r}{2 N f}}\right)$ for coherent OOK (on-off keying), by $p_b = (M - 1) \times \operatorname{erfc}\left(\sqrt{\frac{P_r \times \log_2(M)}{2 N f}}\right)$ for M-ary FSK (frequency shift keying) and by

$$p_b = 0.5 \times \operatorname{erfc}\left(\sqrt{\frac{P_r}{N f}}\right), \quad (17)$$

for BPSK (binary phase-shift keying).

Since we are not interested in the details of a specific modulation scheme but merely want to study the general dependence of the error rate on the received power, we make the following assumptions:

- i) The packet error rate p , equals $S.p_b$, where p_b is the bit error rate and S is the packet size. This is an accurate approximation for small error rates p_b ; thus, we assume that the packet error rate increases/decreases in direct proportion to p_b .
- ii) The received signal power is inversely proportional to D^K , where D is the link distance, and K is the same constant as used in Equation 2. Thus P_r can be replaced by P_t/D^K where P_t is the transmitter power. We choose BPSK as our representative candidate and hence, use Equation 17 to derive the bit-error-rate.

We report results for both fixed and variable power scenarios in our simulations.

- **Fixed transmission power:** In this case, all the nodes in the network use a fixed power for all transmissions, which is independent of the link distance. While such an approach is clearly inefficient for wireless environments, it is representative of several commercial radio interfaces that do not provide the capability for dynamic power adjustment. From Equation 17, it is clear that links with larger distances have higher packet error rates.

For our experiments in this case, we first chose a maximum error rate (p_{max}) for a unit hop along the axes for the grid topology given in Figure 4. Using Equation 2 and 17, it is then possible to calculate the corresponding maximum error rates on the other links.

To add the effect of random ambient noise in the channel, we chose the actual packet error rate on each link uniformly at random from the interval $(0, p_{max})$, where p_{max} is the maximum packet error rate computed for that link. For different experiments, we varied the p_{max} for the unit hop links (and correspondingly the maximum error rates for the other links).

- **Variable transmission power:** In this case, we assume that all the nodes in the network are dynamically able to adjust transmission power across the links. Each node chooses the transmission power level for a link so that the signal reaches the destination node with the *same constant* received power. Since we assume that the attenuation of signal strength is given by Equation 1, the energy requirements for transmitting across links of different lengths is given by Equation 3.

Since all nodes now receive signals with the same power, the bit error rate, given by Equation 17, is now dependent only on the distance-independent receiver noise component. Accordingly, if we assume that the noise level

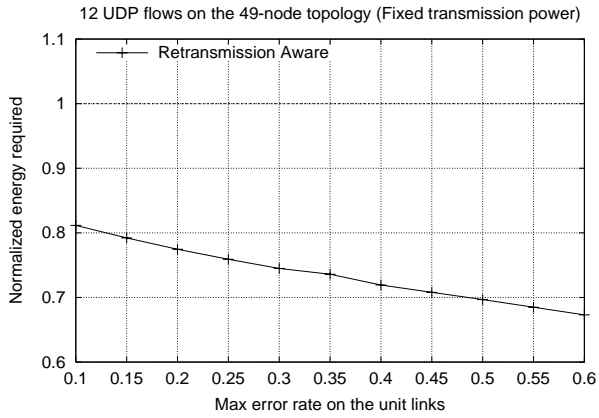


Fig. 5. UDP flows with link layer re-transmissions (HHR) for fixed transmission power scenario.

at different receivers are independent of one another, it follows that the bit-error rates of different links are essentially random and do not depend on the link distance. We simply need to model the random ambient noise at each receiver. We chose the maximum error rate for a link due to ambient noise, (p_{ambient}), for the different experiments in this case. We then chose the actual error rate for any particular link uniformly at random from the interval $(0, p_{\text{ambient}})$.

B. Metrics

To study the energy efficiency of the routing protocols, we observed two different metrics:

- 1) **Normalized energy:** We first compute the average energy per data packet by dividing the total energy expenditure (over all the nodes in the network) by the total number of unique packets received at any destination (sequence number for TCP and packets for UDP). We defined the normalized energy of a scheme, as the *ratio of the average energy per data packet for that scheme to the average energy per data packet required by the minimum-hop routing scheme*. Since, the minimum-hop routing scheme clearly consumes the maximal energy, the normalized energy parameter provides an easy representation of the percentage energy savings achieved by the other (EA and RA) routing algorithms.
- 2) **Effective Reliable Throughput:** This metric counts the number of packets that was reliably transmitted from the source to the destination, over the simulated duration. Since all the plots show results of runs of different schemes over the same time duration, we do not actually divide this packet count by the simulation duration. Different routing schemes will differ in the total number of packets that the underlying flows are able to transfer over an identical time interval.

C. Fixed Transmission Power Scenario

We first present results for the case where each node uses a fixed and constant transmission power for all links. In this case, it is obvious that the EA routing scheme degenerates to the minimum-hop routing scheme.

1) *HHR Model:* We first present the results for the case where each link implements its own localized retransmission algorithm to ensure reliable delivery to the next node on the path.

HHR with UDP: Figure 5 shows the the total energy consumption for the routing schemes under link-layer re-transmissions (HHR case). We experimented with a range of link error rates to obtain these results. As can be seen, the RA scheme shows a significant improvement over the minimum-hop (identical in this environment to the EA) scheme, as expected. With increasing link error rates, the benefits of using our re-transmission aware scheme become more significant. For example, at a maximum link error rate for the unit hop links (p_{max}) of 0.25, the RA scheme consumes about

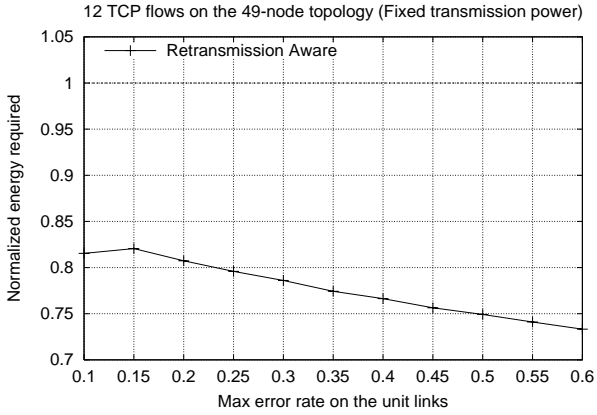


Fig. 6. Energy required for TCP flows with link layer re-transmissions (HHR) for fixed transmission power scenario.

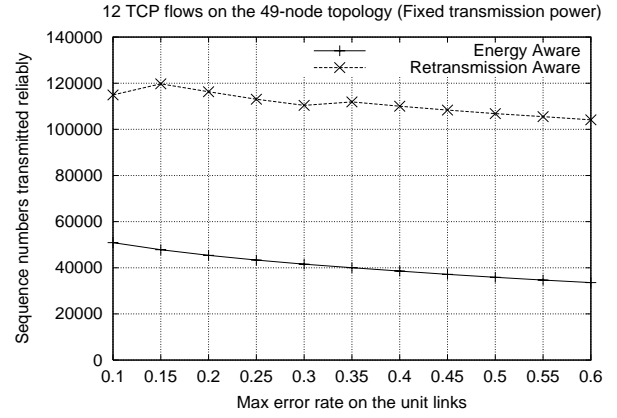


Fig. 7. Reliable packet transmissions for TCP flows with link layer re-transmissions (HHR) for fixed transmission power scenario.

24% lower energy than the other two schemes. Note, that in this case, 0.25 is only the maximum link error rate for the unit links; different unit links will have actual error rates varying between 0.0 and 0.25.

It is perhaps important to emphasize that it is only the *normalized energy* for the RA scheme which decreases with increasing link error rate. The *absolute energy* expenditure will obviously increase with an increasing value of p_{max} for all routing algorithms.

HHR with TCP: In Figure 6, we observe the same metric for TCP flows. When link-layer retransmissions are available, the trends for both UDP and TCP flows, in terms of energy requirements are similar. However, it is more interesting to observe the consequences of using these different schemes on the number of data packets transmitted reliably to the destinations of the flows. This is shown in Figure 7. The RA scheme consistently delivers a larger volume of data packets to the destination within the same simulated duration, even while it is consuming less energy per sequence number transmitted. This is because of two reasons. First, the RA scheme chooses path with lower error rates. Thus the number of link-layer retransmissions seen for TCP flows using the RA scheme is lower, and hence the round trip time delays are lower. The throughput, T , of a TCP flow, with round trip delay, τ and loss rate, p , varies as [12]:

$$T(\tau, p) \sim \frac{1}{\tau} \times \frac{1}{\sqrt{p}} \quad (18)$$

The RA scheme has smaller values of both p and τ and so has a higher throughput.

2) *EER Model:* We now provide the results of our experiments under the EER scheme.

EER with TCP: We looked at the energy requirements when end-to-end TCP re-transmissions are the sole means of ensuring reliable data transfer. The minimum-hop algorithm always chooses a small number of larger distance links. In this fixed transmission power case, the received signal strength is lower for larger distance links. Consequently (by Equation 17), the minimum-hop algorithm chooses links with higher bit error rates. Since there are no link layer re-transmissions, the loss probability for each data segment is fairly high. Therefore this scheme achieves a very low TCP throughput (less than 1% of that achieved by the RA scheme) and still used 10-20% more energy. Hence it was difficult to do meaningful simulation comparisons of the RA scheme with this minimum-hop algorithm.

D. Variable Transmission Power Scenario

In this case, the nodes are able to adapt their transmission power, so that the received signal strength is identical across all links. To achieve this, clearly, links with larger distances require a higher transmission power than links with smaller distances. In this situation, we varied the link error rate due to ambient noise at the receiver of the links to compare the different schemes.




Fig. 8. UDP flows with link layer re-transmissions (HHR) for variable transmission power scenario.

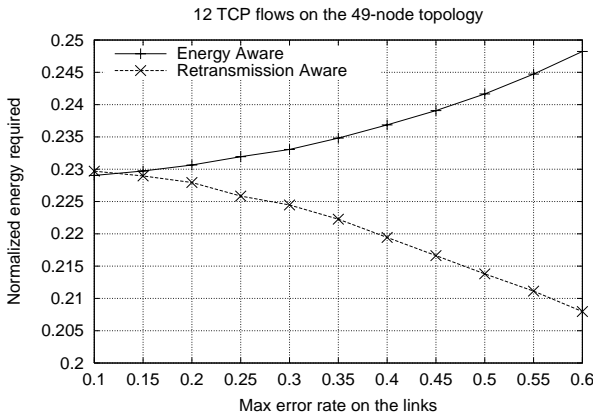


Fig. 9. Energy required for TCP flows with link layer re-transmissions (HHR) for variable transmission power scenario.




Fig. 10. Reliable packet transmissions for TCP flows with link layer re-transmissions (HHR) for variable transmission power scenario.

Unlike the fixed transmission power case, the EA routing algorithm in this case chooses paths with a large number of small hops, and has lower energy consumption than the minimum hop routing algorithm. Therefore, in these results, we compare our RA scheme with both EA and minimum-hop routing.

1) *HHR Model*: We first present the results for the case where each link implements its own localized retransmission algorithm to ensure reliable delivery to the next node on the path.

HHR with UDP: Figure 8 shows the the total energy consumption for the routing schemes under link-layer retransmissions (HHR case). We experimented with a range of channel error rates to obtain these results. Both EA and RA schemes are a significant improvement over the minimum-hop routing scheme, as expected. However, with increasing channel error rates, the difference between the normalized energy required per reliable packet transmission for the RA and the EA schemes diverges. At some of the high channel error rates ($p_{ambient} = 0.5$), the energy requirements of the RA scheme is about 25% lower than the EA scheme. It is again important to note that this error rate is only the maximum error rate for the link. The link error rates of individual links are typically much smaller.

Once again, it is only the normalized energy for the RA scheme which decreases. The absolute energy required obviously increases with an increasing value of p_{max} .

HHR with TCP: In Figure 9, we observe the same metric for TCP flows. As before, the energy requirements of for the RA scheme is much lower than the EA scheme. Additionally, we can again observe (Figure 10) that the number of data packets transmitted reliably for the RA scheme is much higher than that of the EA scheme.

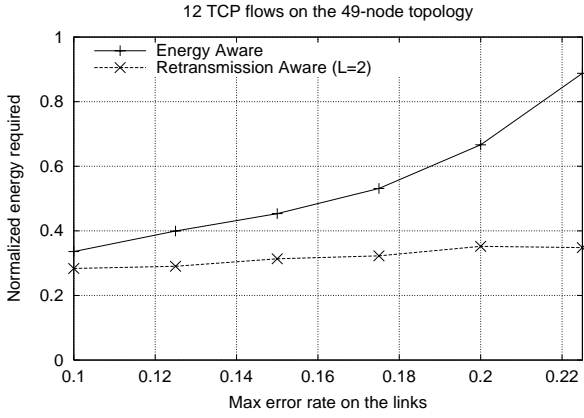


Fig. 11. TCP flows with no link layer re-transmissions (EER) for variable transmission power scenario.

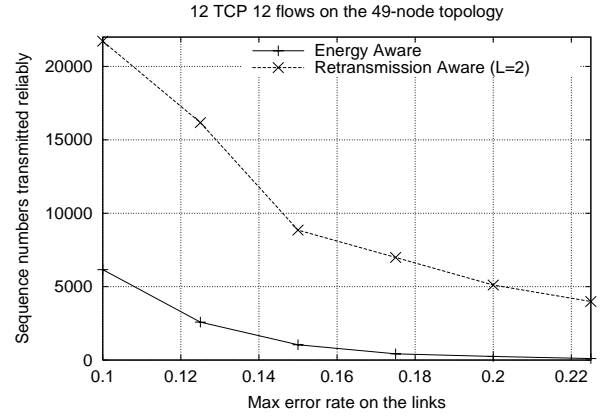


Fig. 12. TCP flows with no link layer re-transmissions (EER) for variable transmission power scenario.




Fig. 13. Varying the L parameter to tradeoff normalized energy and number of reliably transmitted sequence numbers.

2) *EER Model*: Finally, we provide the results of our experiments under the EER framework.

EER with TCP: For the EER case, like before, it was often difficult to simulate links with high error rates— even with a small number of hops, each TCP packet is lost with a high probability and no data ever gets to the destinations.

The energy savings achieved by the RA algorithm is more pronounced when no link-layer retransmission mechanisms are present. For some of the higher link error rates simulated in this environment (e.g., $p_{max} = 0.22$), the energy savings of the RA scheme was nearly 65% of the EA scheme, as can be seen in Figure 11. Again, it is interesting to observe the data packets transmitted reliably by the EA and the RA schemes, simulated over the same duration (Figure 12). The RA scheme transmits nearly an order of magnitude more TCP sequence numbers than the EA scheme, even for relatively small maximum error rates (p_{max} between 0.1 to 0.14). While the total TCP goodput approaches zero for both schemes, as the link error rates increase, the rate of decrease in the TCP goodput is much higher for the EA scheme than the RA scheme.

Varying L : In Figure 13, we varied the L -parameter of Equation 15 for a specific error rate on the links (i.e., $p_{max} = 0.175$). The number of reliably transmitted packets increased monotonically with the value of L . However, the curve in the figure has a minimum “energy per reliably transmitted packet”, corresponding to $L = 5$, in this example⁷. Varying the L -value from this optimal value leads to poorer energy-efficiency (higher energy/packet). There is thus clearly a

⁷Finer measurements with many more L -values would yield the exact L that minimizes this curve.

trade-off between the achieved throughput, and the effective energy expended. To achieve a higher throughput, it is necessary to prefer fewer hops, as well as links with low error rates (higher error rate links will cause higher delays due to re-transmissions). This plot illustrates the following important point: *it is possible to tune the L -parameter to choose an appropriate operating point that captures the tradeoff between a) the achieved TCP throughput, and b) the effective energy expended per sequence number received reliably.* Of course, the right choice of L is expected to be topology dependent. We leave the problem of developing an adaptive algorithm for optimizing L in a specific network as an open problem for future research.

VI. MAXIMIZING NETWORK LIFETIME: THE MRPC ALGORITHM

Selecting the path with the least transmission energy for reliable communication may not always maximize the lifetime of the ad-hoc network. We now discuss how we can include our retransmission-aware formulation of the link cost in an algorithm, called MRPC, that attempts to increase the operational lifetime of multi-hop wireless networks. Unlike previous protocols, MRPC considers both the node characteristics (residual battery energy at the transmitting node) and the link characteristics (link distance and link error rates) link, while evaluating the suitability of alternative paths. Given the current battery power levels at the different nodes, MRPC selects a route that has the *maximum reliable packet carrying capacity* among all possible paths, assuming no other cross traffic passes through the nodes on that path.

To formalize the algorithm, let us assume that the residual battery power at a certain instance of time at node i is B_i . As before, let the transmission energy required by node i to transmit a packet over link $\langle i, j \rangle$ to node j be $E_{i,j}$. Let the source and destination nodes for a specific session (route) be S and D respectively. If the route-selection algorithm then selects a path P from S to D that includes the link $\langle i, j \rangle$, then the maximum number of packets that node i can forward over this link is clearly $B_i/E_{i,j}$. Accordingly, we can define a **node-link metric**, $M_{i,j}$ for the link $\langle i, j \rangle$ as :

$$M_{i,j} = \frac{B_i}{E_{i,j}} \quad (19)$$

The key point in this formulation is that the cost metric includes both a node-specific parameter (the battery power) and a link-specific parameter (the packet transmission energy for reliable communication across the link).

Clearly, the “lifetime” of the chosen path P , defined by the maximum number of packets that may be potentially forwarded between S and D using path P , is determined by the *weakest intermediate node*— one with the smallest value of $M_{i,j}$. Accordingly, the “lifetime” associated with route P is:

$$Life_P = \min_{(i,j) \in P} \{M_{i,j}\} \quad (20)$$

The MRPC algorithm then selects the route $P_{candidate}$ that maximizes the “lifetime” of communication between S and D . Formally, the chosen route is such that:

$$P_{candidate} = \arg \max \{Life_P | P \in \text{all possible routes}\} \quad (21)$$

While the computed route may be optimal at the time of computation, the random traffic patterns will potentially make the currently selected paths sub-optimal at some point in the future. Thus, MRPC is really a route selection algorithm; a routing protocol that uses MRPC for multi-hop wireless networks will include mechanisms for periodic and distributed route computation.

A. Distributed Implementation

Given the cost and lifetime formulations for MRPC (Equations 19 and 20), it is then easy to use a modified version of Dijkstra’s minimum cost algorithm for decentralized route computation. Note that $Life_P$ is not an additive function of the individual node-link costs. It can, however, be computed over a path by applying the \min operator in an iterative fashion.

To apply Dijkstra's algorithm for determining the minimum-cost path, we define the distance metric from any node to the given destination as the value of $Life_P$ over the optimal path from that node to D . Now consider a node A that sees advertisements from its neighbors, $\{X, Y, Z, \dots\}$ with corresponding distance metrics $Life_X, Life_Y, Life_Z, \dots$ for a given destination D . Node A can then compute the best path to D (using its optimal neighbor) by using the following simple algorithm:

- 1) For each of the neighboring nodes ($j \in \{X, Y, Z, \dots\}$), compute the link cost $M_{A,j}$ using Equation 19.
- 2) For each of the neighboring nodes ($j \in \{X, Y, Z, \dots\}$), compute the potential new value of $Life_{pot}$ using

$$Life_{pot}(A, j) = \min \{M_{A,j}, Life_j\}$$

- 3) Select as the next-hop neighbor towards D as that node which results in the maximum value of $Life_{pot}$, i.e., chose node k such that

$$k = \arg \max_{j \in \{X, Y, Z\}} \{Life_{pot}(A, j)\}$$

and assigning $Life_A = Life_{pot}(A, k)$.

It is easy to see that using this recursive formulation allows all nodes in the ad-hoc network to iteratively build their optimal route towards a specific destination D . The distance-vector formulation presented here can easily be incorporated in protocols, such as AODV, DSR and TORA, that are specifically designed for ad-hoc mobile environments. The intent of this paper is not to indicate the choice of a specific routing protocol, but to define a set of power-aware metrics for use by a protocol during route selection.

B. Applying MRPC to Energy-Aware Cost Formulations

The basic MRPC formulation for power-aware routing does not need to specify the value of the transmission energy cost associated with a specific link. Note that Equation 19 is expressed as a function of a generic link cost $M_{i,j}$. Accordingly, by specifying different forms of $M_{i,j}$ it is possible to tailor the MRPC mechanism for specific technologies and/or scenarios.

For the fixed-power scenario, the energy involved in a single packet transmission attempt, $E_{i,j}$, is a constant for all $\langle i, j \rangle$ and is independent of the distance between neighboring nodes i and j . For the variable-power scenario, $E_{i,j}$ will typically be $\propto D_{i,j}^K$, where $D_{i,j}$ is the distance between nodes i and j .

In Section III, we showed why a routing algorithm for reliable packet transfer should include the link's packet error probability in formulating the transmission energy cost. By ignoring the packet error probability, the link cost concentrates (wrongly) only on the energy spent in transmitting a single packet; the correct metric is the *effective* packet transmission energy for reliable transmission, which includes the energy spent in one or more re-transmissions that might be necessary in the face of link errors. In Section IV, we had suggested a transmission energy metric of the form $C_{i,j} = E_{i,j}/(1 - p_{i,j})^L$ where $p_{i,j}$ is the link's packet error probability, $L \geq 1$. For hop-by-hop re-transmissions (a reliable link layer) L should be chosen to be 1; in the absence of hop-by-hop re-transmissions (i.e. re-transmissions are only performed end-to-end), the transmission cost is well approximated by $L \in [3, 5]$.

It should thus be clear that MRPC degenerates to MMBCR [22] *only if* all nodes are incapable of dynamically adapting their power based on the transmission range, and *only if* all links have the same intrinsic error rates. In all other cases, MRPC makes a more intelligent choice, since it takes into consider the potential variability in the energy needed for reliable packet transfer.

For the simulation results reported later, we assume that the nodes are capable of performing link-layer retransmissions for lost or corrupted frames (i.e. HHR case). We also assume that nodes choose the frame transmission power depending on the link attenuation effects (i.e. variable power scenario).




Fig. 14. The 49-node topology. The shaded region marks the maximum transmission range for the corner node, A when $R=1.5$.

C. CMRPC

The CMRPC algorithm is the MRPC equivalent of the CMMBCR algorithm presented in [22]. The CMMBCR algorithm is based on the observation that using residual battery energy as the sole metric throughout the lifetime of the ad-hoc network can actually lower the overall lifetime, since it never attempts to minimize the total energy consumption. Accordingly, the CMMBCR algorithm uses regular minimum-energy routing as long as there is even one candidate path, where the remaining battery power level in all the constituent nodes lies above a specified threshold γ . When no such path exists, CMMBCR switches to MMBCR, i.e., it picks the path with the maximum residual capacity on the “critical node”.

Our CMRPC algorithm differs from CMMBCR in that the cost-functions at all times include the link-specific parameters (e.g. error rates) as defined in Section IV. The algorithm can thus be specified as follows. Let Ψ be the set of all possible paths between the source S and destination D . At any point of time, let Ω represent the set of paths such that:

$$\text{for any route } Q \in \Omega \text{ } Life_Q \geq \gamma$$

i.e., Ω represents the set of paths whose most critical nodes have a lifetime greater than a specified threshold. The routing scheme thus consists of the following actions:

- 1) If $\Omega \neq \emptyset$ (there are one or more paths with $Life$ greater than the threshold, the algorithm selects a path $\bar{Q} \in \Omega$ that minimizes the total transmission energy for reliable transfer, i.e.,

$$\bar{Q} = \arg \min_{Q \in \Omega} \left\{ \sum_{(i,j) \in Q} M_{i,j} \right\} \quad (22)$$

- 2) Otherwise, switch to the MRPC algorithm, i.e., select \bar{Q} such that

$$\bar{Q} = \arg \max \{ Life_Q | Q \in \Psi \}$$

The threshold γ is a parameter of the CMRPC algorithm— a lower value of γ implies a smaller protection margin for nodes nearing battery power exhaustion. Accordingly, the performance of the CMRPC algorithm will be a function of γ .

VII. PERFORMANCE EVALUATION: MAXIMIZING NETWORK LIFETIME

In this section, we use simulation studies to understand the performance benefits and tradeoffs associated with the MRPC and the CMRPC routing algorithms. We compare the performance of 6 different routing schemes:

- 1) *Min-Hop Routing*: This is the conventional “energy-unaware” Internet routing algorithm, where each link is assigned an identical cost.
- 2) *Min-Energy Routing*: This algorithm, referred to RA in Section V, simply selects the path corresponding to the minimum packet transmission energy for reliable communication, without considering the battery power of individual nodes.
- 3) *MMBCR*: This power-aware routing algorithm, described earlier, selects the path whose critical node has the highest residual battery energy.
- 4) *CMMBCR*: This algorithm, also described earlier, switches from minimum-energy paths to MMBCR when the residual battery energy fall below a specified threshold.
- 5) *MRPC*: Our algorithm always uses the path with the highest value of $Life_P$.
- 6) *CMRPC*: This conditional version of MRPC uses the min-energy algorithm as long as $Life_P$ associated with the chosen route P lies above a specified threshold; once the critical link-node cost function falls below this value, the algorithm switches to MRPC-based routes.

A. Simulation Model

For our experiments, we used the same set of topologies as in Section V to study the effects of various schemes on energy requirements and throughputs achieved. In this section, we discuss in detail results from the same 49 node topology as presented in Section V. The corner nodes and the mid-points of each side of the rectangular grid were chosen as traffic sources and destinations—the bold lines in Figure 14 show the session end-points. Each (source, destination) pair had 2 simultaneous sessions activated in the opposite direction, giving rise to a total of 16 different sessions. For the results reported here, each session consisted of a UDP traffic generated by a CBR source whose inter-packet gap was distributed uniformly between 0.1 – 0.2 secs. The error rate on each link was independently distributed uniformly between $(0.05, p_{max})$; we experimented with varying values of p_{max} . Routes were recomputed at 2 second intervals in these simulations to capture the effect of changes in the residual packet capacity on the link metrics.

Whenever nodes died (when its battery power gets completely drained) during the course of a simulation, our simulation code would check whether the graph became partitioned. The simulations were run until each of the 16 sessions failed to find any route from their source to the corresponding destination. To avoid the termination of a simulation due to battery power exhaustion at source or destination nodes, all source and sink nodes were configured to have ‘infinite’ power resources. All the other ‘intermediate’ nodes were configured with identical initial battery power levels.

To study the performance of the various algorithms, we performed experiments where the maximum transmission radius, R , of each node was varied. Figure 14 shows the set of neighboring nodes for a corner node when the transmission radius is set to 1.5. We note the expiration sequence, as well as the node expiry times, for each simulation. The *expiration sequence* (sorted in ascending order of the expiration times) provides a useful indicator of how each algorithm affects the lifetime of the individual nodes, and the entire network. In addition to the expiration sequence, we also calculate the *total packet throughput* by counting the total number of packets successfully received at the destination nodes, and the *energy costs per packet* by dividing the total energy expenditure by the total packet throughput. Except for the expiration sequences, all other metrics were obtained by averaging over multiple runs.

B. Expiration Sequence and Energy Costs

Figures 15, 16 and 17 plot the node expiration sequence, the total packet throughput and the effective energy per packet respectively for all the routing algorithms when the transmission radius, R , was set equal to 1.5 units and $p_{max} = 0.3$. The protection threshold was set to 75% of the total battery capacity for CMMBCR and 75% of the average initial packet capacity for CMRPC. We can see that, as expected, the minimum-hop algorithm performs the worst, since it not only fails to balance the workload among the intermediate nodes, but also uses large-distance (i.e. the diagonal) hops and consequently larger transmission energy. In contrast, while the minimum-energy algorithm does use smaller individual hops, it is also susceptible to high variability in the expiration sequence. The plot effectively demonstrates the improved performance of MRPC over all the other schemes in presence of link errors. In the MMBCR scheme,

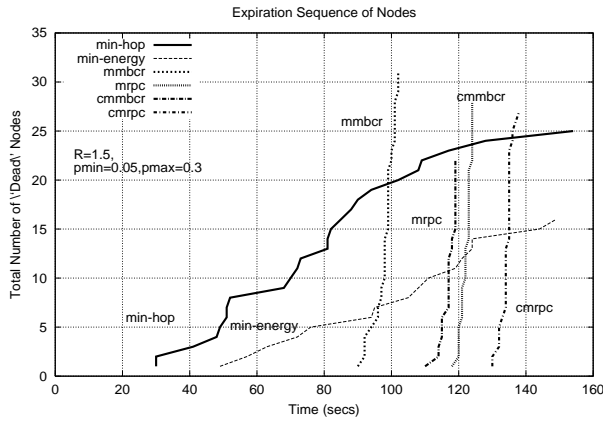


Fig. 15. Expiration Sequence for Different Algorithms, R=1.5

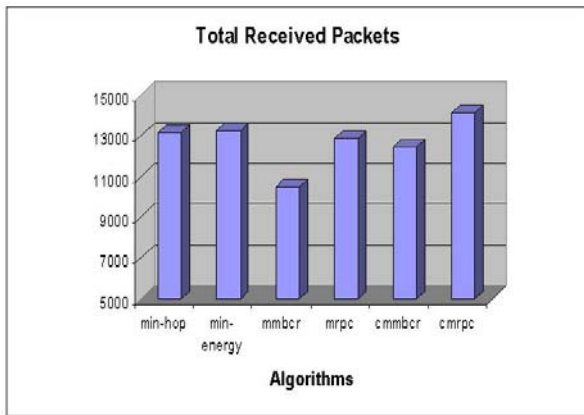


Fig. 16. Total Packet Throughput (UDP Sources), R=1.5

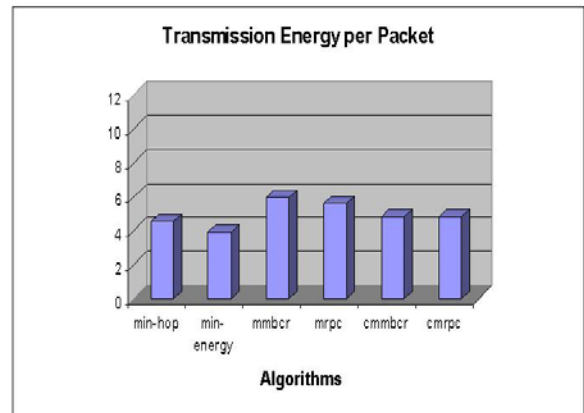


Fig. 17. Avg. Transmission Energy per Received Packet (UDP Sources), R=1.5

the network partitions the sources and the destinations by ~ 102 secs. The MRPC algorithm, instead, is able to ensure packet transfer till ~ 120 secs. The figure also demonstrates the relative performance benefits of the conditional variants of MMBCR and MRPC. Both CMRPC and CMMBCR perform much better than MRPC and MMBCR respectively, since during the initial “minimum-energy routing” phase, it considers the differential transmission energy consumed by different links.

The performance variation among the algorithms can be observed more clearly in Figures 16 and 17. While the min-energy (RA) algorithm obviously results in the lowest effective energy per packet, it also results in a much smaller total packet throughput. In contrast to MMBCR, MRPC is not only able to transmit a much larger number of packets but also at a lower per-packet energy consumption. Similarly, CMRPC outperforms CMMBCR in both the total packet throughput as well as the energy efficiency.

The relative performance of MRPC and CMRPC depend on the choice of the threshold γ . CMRPC does not always outperform MRPC and the benefits depend on the choice of the threshold value. If γ is close to 100%, CMRPC degenerates to MRPC; on the other hand, if γ is close to 0%, CMRPC degenerates to min-energy routing. Figures 18 and 19 show the total network throughput and energy per transferred packet as the CMRPC protection threshold γ is varied. Clearly, the average energy per packet increases with increasing γ , as CMRPC performs minimum-energy routing for a smaller duration. On the other hand, the total network throughput is maximized at an intermediate value for γ (around




Fig. 18. CMRPC: Total Packet Throughput vs. γ

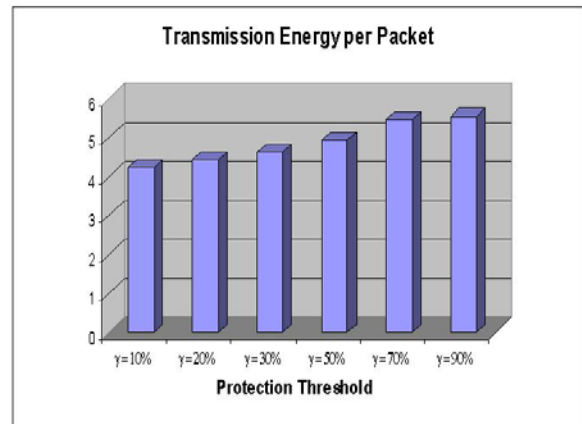


Fig. 19. CMRPC: Avg. Transmission Energy per Received Packet vs. γ

20% in Figure 18). While smaller values (longer min-energy routing) lead to higher variability in the expiration times, larger values fail to exploit minimum-energy paths even if the residual battery capacities are sufficiently large.

VIII. CONCLUSION

In this paper, we have shown why the *effective total transmission energy*, which includes the energy spent in potential retransmissions, is the proper metric for reliable, energy-efficient communications. The energy-efficiency of a candidate route is thus critically dependent on the packet error rate of the underlying links, since they directly affect the energy wasted in retransmissions. Our analysis of the interplay between error rates, number of hops and transmission power levels reveals several key results:

- 1) Even if all links have identical error rates, it is not always true that splitting a large-distance (high-power) hop into multiple small-distance (low-power) hops results in overall energy savings.
- 2) Any routing algorithm must evaluate a candidate link (and the path) on the basis of both its power requirements and its error rate.
- 3) Link-layer retransmission support (HHR) is almost mandatory for a wireless, ad-hoc network, since it can reduce the effective energy consumption by at least an order of magnitude.
- 4) The advantage of using our proposed re-transmission aware routing scheme is significant irrespective of whether fixed or variable transmission power is used by the nodes to transmit across links.

In the model used in this paper, we have only considered the energy consumption for packet transmission, $E_{i,j}$, in formulating our link cost $C_{i,j}$ in Equations 12 and 15. Different studies (e.g., [21]) have shown that the energy expended for packet reception is sometimes comparable to the energy consumed for packet transmission for some wireless technologies. This packet reception cost can be easily accommodated in our energy cost formulation. For example in the HHR framework, we can simply modify Equation 12 as $C_{i,j} = (E_{i,j} + R_{i,j})/(1 - p_{i,j})$, where $R_{i,j}$ is the energy consumed for packet reception on link $\langle i, j \rangle$.

We have also presented a power-aware algorithm, MRPC and its conditional variant, CMRPC, for energy-efficient routing in ad-hoc wireless networks. Our algorithm uses a combined node-link metric to choose routes. Through detailed simulations we show that a protocol based on this algorithm is able to significantly reduce energy requirements at the nodes, as well as extend the operational lifetime of the network.

REFERENCES

- [1] S. Banerjee and A. Misra. Minimum energy paths for reliable communication in multi-hop wireless networks. In *Proceedings of Mobihoc*, June 2002.

- [2] J.-H. Chang and L. Tassiulas. Energy conserving routing in wireless ad-hoc networks. In *Proceedings of Infocom*, Mar. 2000.
- [3] S. Corson, S. Papademetriou, P. Papadopoulos, V. Park, and A. Qayyum. An Internet MANET encapsulation protocol (IMEP) specification, draft-ietf-manet-imep-spec02. txt. *IETF*, Aug. 1999.
- [4] A. El Gamal, C. Nair, B. Prabhakar, E. Uysal-Biyikoglu, and S. Zahedi. Energy-efficient Scheduling of Packet Transmissions over Wireless Networks. In *Proceedings of IEEE Infocom*, June 2002.
- [5] S. Floyd. Connections with multiple congested gateways in packet-switched networks part 1: One-way traffic. *Computer Communication Review*, 21(5), Oct. 1991.
- [6] J. Gass Jr., M. Pursley, H. Russell, and J. Wysocarski. An adaptive-transmission protocol for frequency-hop wireless communication networks. *Wireless Networks*, 7(5):487–495, Sept. 2001.
- [7] J. Gomez and A. Campbell. Power-aware routing optimization for wireless ad hoc networks. In *High Speed Networks Workshop (HSN)*, June 2001.
- [8] J. Gomez, A. Campbell, M. Naghshineh, and C. Bisdikian. Conserving transmission power in wireless ad hoc networks. In *Proceedings of International Conference on Networking Protocols*, Nov. 2001.
- [9] D. Johnson and D. Maltz. Dynamic source routing in ad hoc wireless networks. In *Mobile Computing*, pages 153–181, 1996.
- [10] T. Lakshman, U. Madhow, and B. Suter. Window-based error recovery and flow control with a slow acknowledgment channel: a study of TCP/IP performance. In *Proceedings of Infocom*, Apr. 1997.
- [11] G. Malkin. RIP version 2, RFC 2453. *IETF*, Nov. 1998.
- [12] M. Matthiis, J. Semke, J. Madhavi, and T. Ott. The macroscopic behavior of the TCP congestion avoidance algorithm. *Computer Communication Review*, 27(3), July 1997.
- [13] A. Misra and S. Banerjee. MRPC: Maximizing network lifetime for reliable routing in wireless environments. In *IEEE Wireless Communications and Networking Conference (WCNC)*, Mar. 2002.
- [14] J. Moy. OSPF version 2, RFC 2328. *IETF*, Apr. 1998.
- [15] V. Park and M. Corson. A highly adaptive distributed routing algorithm for mobile wireless networks. In *Proceedings of Infocom*, Apr. 1997.
- [16] C. Perkins and E. Royer. Ad-hoc on-demand distance vector routing. In *Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications*, Feb. 1999.
- [17] B. Prabhakar, E. Uysal-Biyikoglu, and A. El Gamal. Energy-efficient Transmission over a Wireless Link via Lazy Packet Scheduling. In *Proceedings of IEEE Infocom*, Apr. 2001.
- [18] K. Scott and N. Bamboos. Routing and channel assignment for low power transmission in PCS. In *Proceedings of ICUPC*, Oct. 1996.
- [19] S. Singh and C. Raghavendra. Pamas-power aware multi-access protocol with signaling for ad hoc networks. In *ACM Communications Review*, July 1998.
- [20] S. Singh, M. Woo, and C. Raghavendra. Power-aware routing in mobile ad-hoc networks. In *Proceedings of Mobicom*, Oct. 1998.
- [21] M. Stemm and R. Katz. Measuring and reducing energy consumption of network interfaces in hand-held devices. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Science, Special Issue on Mobile Computing*, 80(8), Aug. 1997.
- [22] C. Toh, H. Cobb, and D. Scott. Performance evaluation of battery-life-aware routing schemes for wireless ad hoc networks. In *Proceedings of ICC*, June 2001.
- [23] M. Zorzi and R. Rao. Error control and energy consumption in communications for nomadic computing. *IEEE Transactions on Computers*, 46(3), Mar. 1997.