

Minimum Energy Paths for Reliable Communication in Multi-hop Wireless Networks

Suman Banerjee
Department of Computer Science
University of Maryland at College Park
College Park, MD 20742, USA

Archan Misra
IBM T J Watson Research Center
19 Skyline Drive
Hawthorne, NY 10532, USA

CS-TR 4315
December 2001

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 misleading — 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. Finally, through detailed simulations, we show that our schemes can lead to upto 30-70% energy savings over best known current schemes, under realistic environments.

1 Introduction

Multi-hop wireless networks typically possess two important characteristics:

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

Energy-aware routing protocols (e.g., [1, 2, 3]) for such networks typically select routes that minimize the total transmission power over all the nodes in the selected path.

In *constant-power* scenarios, where the transmission power of a node is chosen independent of the distance of the link, conventional minimum-hop routing [4, 5] will be most energy efficient when the links are error free. In alternative *variable-power* scenarios, where the nodes can dynamically vary their transmitter power levels, the transmission power is typically a function of the distance between the transmitter and receiver nodes. Mathematically speaking, the transmission power P is proportional to some higher order of the distance D , i.e.,

$$P \propto D^K \quad K \geq 2 \quad (1)$$

where K is a constant that depends on the propagation medium and antenna characteristics¹. In these

¹ K is typically around 2 for short distances and omnidirectional antennae, and around 4 for longer distances.

scenarios, proposals for energy-efficient routing protocols (e.g., [1, 6]) typically aim to choose a path with a *large number of small-range hops*, since they consume less power than an alternative route that has a smaller number of hops, but a larger distance for individual hops. In general, most formulations for computing energy efficient paths employ algorithms for computing minimum-cost paths, with the link metric determined by the energy required to transmit a single packet over that link. Setting this link cost to 1 (and thus computing minimum hop paths) suffices in constant-power scenarios, since the transmission energy is the same for all links.

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 consider 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 provide a reliable transport layer. Therefore, the energy cost associated with a candidate path should thus reflect not merely the energy spent in just transmitting a *single* packet across the link, but rather the “total effective energy” spent in packet delivery, which includes the energy spent in potential retransmissions as well².

We first consider how the error rate of individual links affects the overall number of transmissions needed to ensure reliable packet delivery. Such an analysis helps to clearly delineate how the energy associated with the reliable delivery of a packet differs from the energy associated with simply transmitting the packet. As part of this analysis, we consider two different operating models:

- a) **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.

- b) **Hop-by-Hop Retransmissions (HHR)**: where each individual link provides reliable forwarding to the next hop using localized packet retransmissions.

We shall see that, in both cases, it is important to consider the link’s error rate as part of the route selection algorithm, since the choice of links with relatively high error rates can significantly increase the effective energy spent in reliably transmitting a single packet. This is true in both the constant-power and variable-power scenarios — in either scenario, ignoring the error rate of the link leads 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 shall see that the choice between a path with many short-range hops and another with fewer long-range hops is non-trivial, but involves a *tradeoff 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.

We then study how routing algorithms can be used to minimize our new objective function: the energy required to *reliably* transmit a packet to the destination, *the effective transmission energy*. Since most decentralized ad-hoc routing protocols (e.g., AODV [7], DSR [8]) attempt, at least approximately, to select a minimum-cost path (where the path cost is a sum of the individual link costs), we define a new link cost as a function of both the link distance and the link error rate. We shall 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.

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

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

the link error rates in the link cost, energy savings of 30% to 70% can often be achieved under realistic operating conditions.

The rest of the paper is organized as follows. Section 2 provides an overview of previous related work. Section 3 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 4 shows how to form link costs that lead to the selection of minimum effective energy paths. In Section 5 we present the results of our simulation studies on certain on ad-hoc topologies, for both the fixed-power and the variable-power scenarios. Finally, Section 6 concludes the paper.

2 Related Work

Metrics used by conventional routing protocols for the wired Internet typically do not need to consider any energy-related parameters. Thus, RIP [4] uses hop count as the sole route quality metric, thereby selecting minimum-hop paths between the source and destinations. OSPF [5], 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 [1], is one such minimum total transmission energy protocol, where the link cost was set to the transmission power and Dijkstra’s shortest path algorithm was 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 [2]. 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.

Most ad-hoc routing protocols essentially aim to compute minimum-cost paths; in contrast to generic (non ad-hoc) routing protocols, they contain special features to reduce the signaling overheads and convergence problems caused by node mobility and link failures. So, ad-hoc protocols, such as AODV or DSR, can in principle be adapted to yield minimum-energy 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 [6, 9] is designed for scenarios where the nodes can dynamically adjust their transmission powers—PARO attempts to generate a path with a large 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.

Researchers in energy-aware routing have also considered other objective functions, besides the one of minimum total energy. One alternative approach considers the battery capacity of individual nodes; such *battery-aware routing* algorithms typically aim to extend the lifetime of all the ad-hoc nodes by distributing the transmission paths among nodes that currently possess greater battery resources. Such algorithms are based on the observation that minimum-energy routes can often unfairly penalize a subset of the nodes; for example, if several minimum energy routes have a common host, the battery of that host will be exhausted quickly. Among such battery-aware algorithms, [10] 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. Since this mechanism could still lead to the choice of a path having a node that was nearing exhaustion (especially if the other nodes on the path had high residual capacity), the basic MMBCR algorithm and its variant (CMMBCR) [11] formulates path selection as a min-max problem.

In this approach, the capacity of a route is defined as the battery level of the critical (most drained) node; the algorithm then selects the path with the highest capacity.

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 reliable packet delivery. Our problem formulation and routing solution implicitly assumes 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 receiver ³.

3 Transmission Energy for Reliable Packet Delivery under Link Errors

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 (i, j) 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 on the channel. 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 (i, j) and is essentially 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 essentially 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 . According, the optimal transmission power associated with a link of distance D in the variable-power scenario is given by:

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

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:

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

It is now easy to understand, at least qualitatively, the impact of neglecting the link error rates while determining a specific path between the source and destination nodes. For the fixed-power case, the minimum-hop path may not be the most “effective” energy-efficient, since an alternative path with more hops may prove to be better if its overall error rate is sufficiently low. For the variable-power case, a path with a greater number of smaller hops may not always be better; the savings achieved in the individual transmission energies (given by Equation 3) may be nullified by a larger increase in link errors and consequently retransmissions.

We now analyze the interesting consequences of this behavior for the variable-power scenario (for both the EER and HHR cases); we omit the analysis for the fixed-power scenario which is simpler, and a special case of our ensuing analysis.

3.1 Effect of Link Errors on Optimal Route in EER Case

In the EER case, a transmission error on any link leads to a complete retransmission over the entire 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

³Similar ideas were proposed for link sensing in the Internet MANET Encapsulation Protocol [12].

adjacent nodes (breaking up a link of distance D into two shorter links of distance D_1 and D_2 such that $D_1 + D_2 = D$) always reduces the total E_{opt} . In fact, PARO works using precisely such an estimation. From a reliable transmission energy perspective, such a comparison is, however, 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) located at 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 inserting by a power-aware routing protocols. For notational ease, let the 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 optimality arguments, it is easy to see that the minimum transmission energy case occurs when each of the hops are of equal length $\frac{D}{N}$. 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 $\frac{1}{1-p}$. 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:

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

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 variable and taking derivatives, it is easy to see 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 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




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

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 analyses, a path with multiple shorter hops is thus not always beneficial than one with a smaller number of long-distance hops.* Accordingly, protocols such as PARO must be modified to consider the impact of introducing an additional hop on the overall probability of reliable delivery.

3.1.1 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




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

[13, 14]), 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* [15] simulator, 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 $p_{link} = \{0.01, 0.05\}$. The remarkably close agreement between our analytical predictions and TCP-driven simulation results verifies the practical utility of our analytical model.

3.2 Effect of Optimal Route 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, which typically always employ link-layer retransmissions. In this case, the link layer retransmissions on a specific link essentially 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




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

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:

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

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*. Of course, the logarithmic scale for the energy cost compresses the differences in the value of $E_{\text{total rel}}^{\text{HHR}}$ for different p_{link} . By itself, this result is not very interesting: if all links have the same error rate, it is beneficial to substitute a single hop with multiple shorter hops.

A more interesting study is to observe the total energy consumption, for a fixed N , for different values of p_{link} . Clearly, for moderately large values of p_{link} , the number of total transmissions (and hence, the energy consumption) increases super-linearly with an increase in the link error rate. The graph thus shows the importance of choosing links with appropriate

link error rates, even in the HHR case. (In the EER case, Figure 1 clearly demonstrates that the effect of larger link error rates is much more drastic — when $N = 10$, for example, increasing the loss probability from 0.1 to 0.2 can increase the energy consumption ten-fold.) An energy-aware algorithm that does not consider the error rates of associated links would not distinguish between two paths, each of 10 nodes having the same D values but *different* packet error rates. However, our analysis clearly shows that the effective energy consumed by a path consisting of links with higher packet error rates would be much larger than a path with smaller error rates.

We obtain another meaningful observation by comparing the values for $E_{\text{total rel}}$ for the EER and HHR cases (Figures 1 and 3), for identical values of N and K . It is easy to see that, for moderate to high values of p_{link} , the EER framework results in at least an order of magnitude higher energy consumption than the HHR case. 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.

4 Link Costs for Energy-Aware Routing for Reliable Communication

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 also 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.

Before presenting the appropriate link costs, it is necessary to define graph used for computing the shortest cost paths. Consider a graph, with the set of vertices representing the communication nodes and a link $l_{i,j}$ representing the direct hop between nodes i and j . For generality, assume an asymmetric case where $l_{i,j}$ is not the same as $l_{j,i}$; moreover, $l_{i,j}$ 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 transmission energy costs over each constituent link.

4.1 The Hop-by-Hop Retransmission (HHR) Scenario

Consider a path P from a source node S (indexed as node 1) to node D 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:

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

Various ad-hoc routing protocol, such as TORA or AODV, can then use this link cost to compute the appropriate minimum-energy routes. Newer ad-hoc routing protocols, such as PARO, can also be

easily adapted to use this new link cost formulation. Thus, in the 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 $C_{A,B}$.

4.2 End-to-End Retransmission (EER) Scenario

In the absence of hop-by-hop retransmissions, the expression for 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 costs⁴, 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; making the link cost for error-prone links prohibitively high, we can ensure that such links are usually excluded by shortest-cost path computations.

In particular, for a path consisting of k identical links (i.e. have the same link error rate and link trans-

⁴We do not consider solutions that require each node or link to separately advertise two different metrics. If such advertisements were allowed, we can indeed compute the optimal path accurately. For example, if we considered two separate metrics— a) $E_{i,j}$ and b) $\log(1 - p_{i,j})$, then a node can accurately compute the next hop neighbor (using a distance-vector approach) to a destination D by using the cumulative values $\sum E_{i,j}$ and $\sum \log(1 - p_{i,j})$ advertised by its neighbor set.

mission 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 links⁵. 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, we require that a link should advertise a single cost for that link for distributed route computation, in accordance with current routing schemes. 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, regular ad-hoc routing protocols, or newer ones such as PARO, can use this new link cost function C^{approx} to evaluate their routing decisions.

As with our theoretical studies in Section 3, 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.

⁵There should be an L factor in the numerator too (as in Equation 14, but since this is identical for all links, it can effectively be ignored.

5 Simulation Studies and Performance Evaluation

The analysis of the previous section provides a foundation for devising energy-conscious protocols for reliable data transfer. In this section, we report on extensive simulation-based studies on the performance impacts of our proposed modifications in the ns-2 [15] simulator. The traffic for our simulation studies consists of two types:

- i) For studies using the EER framework, we used TCP flows implementing the NewReno version of congestion control.
- ii) 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:

- a) The minimum-hop routing algorithm, where the cost of all links was identical and independent of both the transmission energy and the error rate.
- b) 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.
- c) Our Retransmission-Energy Aware (**RA**) algorithm, where the link cost includes the packet error rates, and thus considers the (theoretical) 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 different values of L).

In the fixed-power scenario, the minimum-hop and EA algorithms exhibit identical behavior; 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 over on a square region, to study the effects of various schemes on energy requirements




Figure 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.

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 algorithms (two for the fixed-power scenario, three for the variable-power scenario) 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.

5.1 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, most modulation schemes exhibit the following

generic relationship between p_b and P_r :

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

where N is the noise spectral density (noise power per Hz) and $\text{erfc}(x)$ is defined as the complementary function of $\text{erf}(x)$ and is given by

$$\text{erfc}(x) = 1 - \int_0^x e^{-t^2} dt$$

As specific examples, the bit error rate for is given by $p_b = \text{erfc}\left(\sqrt{\frac{P_r}{N}}\right)$ for coherent OOK (on-off keying), by $p_b = (M - 1) \times \text{erfc}\left(\sqrt{\frac{P_r \log_2(M)}{N}}\right)$ for M-ary FSK (frequency shift keying) and by $p_b \approx 2 \sin\left(\frac{\pi}{M\sqrt{2}}\right) \times \text{erfc}\left(\sqrt{\frac{2P_r}{N}}\right)$ for QPSK (quadrature phase shift keying). Thus, the bit error rate for binary PSK (BPSK) is given by:

$$p_b = 2 \sin\left(\frac{\pi}{2\sqrt{2}}\right) \times \text{erfc}\left(\sqrt{\frac{2P_r}{N}}\right) \quad (17)$$

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 $L.p_b$, where p_b is the bit error rate and L 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 T/D^K where T is the transmitter power. We choose BPSK as our representative candidate and hence, use Equation 17 to derive the bit-error-rate.

We study both the 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 the same for all links (by using the flexibility of adjusting the transmission power based on link distances). Therefore, for this scenario, we only need to model the additional link error rate due to ambient noise at the receiver. We chose the maximum error rate for a link due to ambient noise ($p_{ambient}$), for the different experiments in this case, and chose the actual error rate for a link uniformly at random from the interval $(0, p_{ambient})$.

5.2 Metrics

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

- 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.
- 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.

5.3 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.

5.3.1 HHR Model

We first present the results for the case where each hop (link) implemented 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 retransmissions (HHR case). We experimented




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

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. The normalized energy requirements of the minimum-hop and the EA schemes is unity in this case. With increasing link error rates, the benefits of using our re-transmission aware scheme becomes 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 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; typical 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. As can be seen, the trends for both UDP and TCP flows, in terms of energy requirements are similar, when link-layer retransmissions are present. 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




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




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

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 at many times 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 [16]:

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

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

5.3.2 EER Model

We now provide the results of our experiments under the EER scheme.

EER with TCP: We look 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. However, in this fixed transmission power case, the received signal strength over larger distance links is lower, and consequently, by Equation 17, has a higher bit error rate. Since there are no link layer retransmissions, 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.

5.4 Variable Transmission Power Scenario

In this case, the nodes are capable to adapting the 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.

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.

5.4.1 HHR Model

We first present the results for the case where each hop (link) implemented 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




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

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 useful to note, that this error rate is only the maximum error rate for the link. The actual link error rate is 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.

5.4.2 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




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




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




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




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

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). For lower error rates (p_{max} between 0.1 to

0.14) the RA scheme transmits nearly an order of magnitude more TCP sequence numbers than the EA scheme. 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”,




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

corresponding to $L = 5$, in this example⁶. Varying the L -value from this optimal value leads to poorer energy-efficiency (higher energy/packet). There is thus clearly a 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 retransmissions). 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.*

6 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,

⁶Finer measurements with many more L -values would yield the exact L that minimizes this curve.

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. Our analysis shows that if the number of hops exceeds an optimal value (which can be as small as 5 – 10 in realistic scenarios), the rise in the overall error probability negates any apparent reduction in the transmission power.
2. Any routing algorithm must evaluate a candidate link (and the path) on the basis of both its power requirements and its error rate. Even in the HHR framework, where retransmissions are typically localized to a specific hop, the choice of an error-prone link can lead to significantly higher effective energy expended per packet.
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 advantages of using our proposed retransmission aware routing scheme is significant irrespective of whether fixed or variable transmission power is used by the nodes to transmit across links.

We also studied modifications to the link cost that would enable conventional minimum-cost path algorithms to select optimal “effective energy” routes. While the appropriate cost for link (i, j) turned out to be $\frac{E_{i,j}}{1-p_{i,j}}$ for the HHR framework, it was not possible to define an exact link cost for the EER case. For the EER scenario, we studied the performance of approximate link costs of the form $\frac{E_{i,j}}{(1-p_{i,j})^L}$ for various values of L . Our simulation studies show that the incorporation of the error rate in the link cost leads to significant energy savings (potentially as high as 70%) compared to existing minimum-energy algorithms. It also turns out that, in the HHR model, the L parameter in the link cost provides a knob to trade off energy efficiency with network throughput (capacity). While larger values of L always lead to the selection of shorter-hop routes and larger session

throughput, the energy-efficiency typically increases and then decreases with increasing L .

As part of future research, we intend to extend our analyses (which assumed each link to be operating independently of other links) to scenarios, such as IEEE 802.11-based networks, where the logical links share the same physical channel and hence, interfere with one another. Indeed, since an energy-aware routing protocol defines the next-hop node (and hence, implicitly defines the associated transmission power), the choice of the routing algorithm is expected to affect both the overall network capacity and individual session throughputs in such scenarios.

References

- [1] S. Singh and C.S. Raghavendra, "PAMAS-PowerAware Multi-Access Protocol with Signaling for Ad Hoc Networks", *ACM Communication Review*, July 1998.
- [2] K. Scott and N. Bamboos, "Routing and Channel Assignment for Low Power Transmission in PCS", *Proceedings of ICUPC'96*, October 1996.
- [3] J.-H. Chang and L. Tassiulas, "Energy Conserving Routing in Wireless Ad-hoc Networks", *Proceedings of INFOCOM '00*, March 2000.
- [4] G. Malkin, "RIP Version 2", RFC 2453, *IETF*, November 1998.
- [5] J. Moy, "OSPF Version 2", RFC 2178, *IETF*, April 1998.
- [6] J. Gomez-Castellanos, A. Campbell, M. Naghshineh and C. Bisdikian, "PARO: A Power-Aware Routing Optimization Scheme for Mobile Ad Hoc Networks", draft-gomez-paro-manet-00.txt, Work in Progress, *IETF*, March 2001.
- [7] C. Perkins, E. Royer and S. Das, "Ad Hoc On Demand Distance Vector (AODV) Routing", draft-manet-ietf-aodv-8.txt, *IETF*, Work in Progress, March 2001.
- [8] D. Johnson and D. Maltz, "Dynamic Source Routing in Ad Hoc Wireless Networks", In *Mobile Computing*, chapter 5 pp. 153-181, Kluwer Academic Publishers, 1996.
- [9] J. Gomez and A. Campbell, "Power-Aware Routing Optimization for Wireless Ad Hoc Networks", High Speed Networks Workshop (HSN 2001), June 2001.
- [10] S. Singh, M. Woo and C.S. Raghavendra, "Power-Aware Routing in Mobile Ad Hoc Networks", *Proceedings of Mobicom'98*, October 1998.
- [11] C.K. Toh, H. Cobb and D. Scott, "Performance Evaluation of battery-Life-Aware Routing Schemes for Wireless Ad Hoc Networks", *Proceedings of IEEE ICC'01*, June 2001.
- [12] 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* August 1999, Work in progress.
- [13] T. V. Lakshman, U. Madhow and B. Suter, "Window-based Error Recovery and Flow Control with a Slow Acknowledgment Channel: a Study of TCP/IP Performance", *Proceedings of INFOCOM'97*, April 1997.
- [14] S. Floyd, "Connections with Multiple Congested Gateways in Packet-Switched Networks Part 1: One-way Traffic", *Computer Communication Review* October 1991.
- [15] "The ns-2 network simulator", <http://www-mash.cs.berkeley.edu/ns>.
- [16] M. Matthys, J. Semke, J. Mahdavi and T. Ott, "The Macroscopic Behavior of the TCP Congestion Avoidance Algorithm", *Computer Communications Review*, July 1997.