Power Adaptation based Optimization for Energy Efficient Reliable Wireless Paths

Suman Banerjee[1] and Archan Misra[2]

[1] Computer Sciences Dept., University of Wisconsin, Madison, WI 53706, USA. Email: suman@cs.wisc.edu
[2] IBM T.J. Watson Research Center, Hawthorne, NY 10532, USA. Email: archan@us.ibm.com

Abstract. We define a transmission power adaptation-based routing technique that finds optimal paths for minimum energy reliable data transfer in multi-hop wireless networks. This optimal choice of the transmission power depends on the link distance between the two nodes and the channel characteristics. Typical energy efficient routing techniques use a transmission power such that the received signal power at the destination minimally exceeds a desired threshold signal strength level. In this paper we argue that such a choice of the transmission power does not always lead to optimal energy routes, since it does not consider differences in the receiver noise levels.

We first analyze the optimal transmission power choices for both the ideal case from an information-theoretic perspective, and for realistic modulation schemes. Subsequently we define our technique for transmission power adaptation that can be used in existing routing protocols for multi-hop wireless networks. Our simulations show that current best-known schemes incur upto 10% more energy costs in low noise environments, and upto 6.67 times the energy costs in high noise environments compared to our proposed scheme.

1 Introduction

We define a distributed route computation technique that computes minimum-energy paths for reliable multi-hop wireless communication. This technique consists of two parts — (1) adaptation of the transmission power level for each wireless link, and (2) assignment of link costs to individual links which will be used by standard routing protocols to compute end-to-end paths. Both these mechanisms take into account the specific channel characteristics of the links and are necessary and sufficient to compute the optimal energy-efficient path for reliable communication. In this work we show that our technique is optimal, i.e. it computes the minimum-energy path for *reliable packet delivery* in a multi-hop wireless network.

Typical minimum-energy routing protocols for multi-hop wireless networks assign the transmission power required to sustain communication over a link as the link cost. Subsequently they use standard route computation techniques to obtain appropriate end-to-end paths. Due to the properties of wireless signal attenuation these algorithms observe that the total energy requirements for packet transfer over the entire path can be minimized by choosing a route consisting of a large number of small-distance hops over an alternative one with a small number of large-distance hops [6, 12]. However,

these algorithms do not necessarily yield minimum-energy paths for reliable packet delivery. This is because the the link metrics in these algorithms depend solely on the energy spent in a single transmission attempt and do not capture the additional energy expended on retransmissions in the presence of link errors. In [1] the authors had shown how link costs should be assigned to account for the total energy spent in reliable packet delivery, which also includes energy consumed in packet retransmissions. However the authors in [1] continue to make the assumption that the transmission power level on an individual link is determined solely by the corresponding link distance. In this work, we argue that such a choice of transmission power level is not optimal in finding minimum energy costs — appropriate adaptation of the transmission power based on the channel noise characteristics is a crucial component in reaching this objective.

2 Optimal Transmission Power for Individual Links

In our analytic study we first use an information-theoretic approach to show how the optimal transmission power and the associated minimum reliable transmission energy depends on both the link distance and the channel characteristics. Since the resulting bounds are essentially theoretical and not practically realizable due to severe buffering and delay constraints, we then apply our framework to practical channel models. Due to space constraints in this extended abstract, we only summarize the main results. Specific details of the analysis can be found in [2].

2.1 Information-Theoretic Bounds on Optimal Transmission Power

We utilize the information-theoretic bound on the maximum capacity of the well-known band-limited Gaussian channel in order to ascertain the existence of an optimal transmission power for energy efficient reliable data forwarding across a link. The Gaussian channel models an environment where the noise component (both thermal and due to interfering transmissions) at the receiver is assumed to have a Gaussian spectral distribution and is additive in nature. In our analysis we use Shannon's channel capacity theorem to show that if a node transmits at a power level P_t, it follows that the *normalized reliable transmission energy* (i.e. the energy needed per bit of reliable transfer) is related to its transmit power level, P_t, as :

$$E(P_t) = \frac{P_t}{W \times \log_2(1 + \frac{P_t}{D^\alpha \times \eta \times W})} \qquad (1)$$

where W is the spectral width of the channel, η is the spectral noise density, and D is the distance between the transmitter and the receiver. This an increasing function of P_t.

Thus in theory, the smaller we make the transmission power, the more energy-efficient the communication process. The greatest energy efficiency (lowest cost per reliably transferred bit) is achieved as $P_t \downarrow 0$.

This minimum energy value can be obtained by applying the L'Hospital's rule to Equation 1, as is given by:

$$E_{opt} = \ln 2 \times \eta \times D^\alpha \qquad (2)$$

Therefore, *every channel is associated with a fundamental theoretical (non-zero) lower bound on the minimum energy needed to reliably transfer a single bit.*

The above results show that maximum energy efficiency is achieved by transmitting at as low a power level as possible, and that a non-zero communication rate can be sustained even if the received power is much smaller than the channel noise. This is clearly not possible in any *practical* communication system with realistic bounds on the transfer latency. Indeed, Shannon's result is based on the use of asymptotically long coding sequences, resulting in unbounded transmission delays. Now we will consider a *practical* communication sub-system. We shall then see that there exists a non-zero optimal transmission power-level P_t^*: while smaller values of the transmission power result in a sharp increase in the total number of retransmissions needed, values larger than the optimum end up wasting unnecessarily large amounts of energy in a single transmission.

2.2 Optimal Transmission Power for Practical Modulation Schemes

Consider a wireless channel with a simple retransmission based error correction mechanism; if a data packet is lost in transmission across any link, it is retransmitted by the upstream node of that link. Assuming independent packet losses, the expected number of transmissions needed for the reliable transfer of one data packet across any link is $1/(1-p)$, where p is the packet error rate of the link. Let us assume that data packets are L bits long. Then for such a wireless channel with data transmission rate, f, the normalized reliable transmission energy (energy required per bit of reliable transfer) is given by:

$$E(P_t) = \frac{L}{f} \times P_t \times \frac{1}{1 - p(P_t)} \times \frac{1}{L} \tag{3}$$

The packet error rate depends on the chosen transmission power, P_t, and hence so does the normalized energy requirements. The packet error rate can be computed as:

$$p(P_t) = 1 - (1 - p_b(P_t))^L \tag{4}$$

where p_b is the bit error rate of the wireless channel. Consider Binary Phase Shift Keying (BPSK) as an example of a practical channel modulation scheme, where the bit error rate is given by:

$$p_b = 0.5 \times erfc(\sqrt{\frac{P_t \times W}{D^\alpha \times N_r \times f}}) \tag{5}$$

where D is the link distance, W is the channel bandwidth, N_r is the noise signal power, and and f is the data transmission rate. BPSK modulation is used in wireless environments, for example in the 1 Mbps version of the IEEE 802.11 wireless LAN standard.

In Figure 1 we plot the variation in $E(P_t)$ for a channel employing BPSK modulation as a function of the transmission power P_t and a packet size of 1000 bytes. We set the channel parameters to be representative of the 802.11b standard, with a bit rate of 1 Mbps and a noise bandwidth (post de-spreading) of 2 MHz. The link distance D
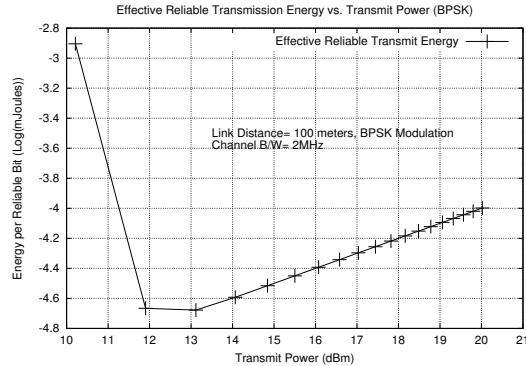
Fig. 1. Reliable Transfer Energy Behavior (BPSK)

is 100 meters and the spectral noise N_r is 4.0×10^{-11} W. We can see that the optimal transmission power for this channel exists, and is ≈ 20 mW.

In our analysis we have also explored how the optimal transmission power for a link varies with changes in link distance. Our results show that this optimal transmission power *increases at a slightly slower rate than D^α* where α is the constant is the signal attenuation model. Accordingly, conventional protocols, such as [1, 6, 12], which assign link costs to be proportional to D^α, penalize longer links more than needed. From Equations 3, 4, and 5 it also follows that the optimal transmission power depends on the spectral noise in the channel. Our study also showed that a policy of maintaining a "constant target Signal to Noise Ratio" at the receiver (by appropriately adjusting the transmission power) is a good informal rule that achieves "close to optimal" energy efficiency. In the full version [2] of this paper, we also show how the optimal energy efficiency of a practical modulation scheme compares with the theoretical lower bound provided by Equation 2.

3 Minimum Energy Routing — Assigning Link Costs

Consider a node that transmits a packet with transmit power P_t across a link. Let the corresponding energy required for this single transmission be E_t. (Assuming packets are of constant size, E_t differs from P_t by a proportionality constant.) Let E_{max} be the packet transmission energy corresponding to the maximum transmission power level at which the node can transmit. Similarly let E_{min} be the packet transmission energy corresponding to the minimum transmission power necessary to ensure that the signal strength at the receiver exceeds a desired threshold.

In this extended abstract we only describe the scenario where the wireless links implement link-layer retransmissions to recover from packet losses. However our proposed technique is equally applicable for alternative link layer reliability mechanisms like forward error correction. In another alternative scenario, if the wireless links do

not implement any recovery mechanisms, data needs to be recovered using end-to-end retransmissions. Description of these alternative cases can be found in the full version of the paper [2].

Let $p(E_t)$ denote the packet error probability corresponding to the packet transmission energy, E_t. The optimal value of the energy required for reliable packet delivery across a single link is given by the solution to:

$$\frac{d}{dE_t} E_t(reliable, HHR) = 0 \tag{6}$$

$$\frac{d^2}{dE_t^2} E_t(reliable, HHR) \geq 0 \tag{7}$$

where $E_t(reliable, HHR) = E_t/(1 - p(E_t))$. It follows that this optimal value, E_t^*, that minimizes the energy cost for a link satisfies:

$$E_t^*.p'(E_t^*) - p(E_t^*) = 1 \tag{8}$$

where $p'(.)$ denotes the first derivate of $p(.)$ with respect to E_t. E_t^* can be computed using efficient numerical techniques. However, this optimal solution may exceed E_{max}. In such a case our choice of optimal transmission energy, $E_t^* = E_{max}$. Similarly if the computed E_t^* falls below E_{min}, we assign $E_t^* = E_{min}$. In the routing protocol, we assign each wireless link a cost which is given by energy required for reliable packet transmission across that link, given by $E_t^*/(1-p(E_t^*))$. The end-to-end route can therefore be computed in a distributed manner by any standard routing protocol capable of computing minimum cost paths. It follows that shortest cost path found by the routing algorithm will be the *optimal energy-efficient route* for that end-to-end path.

4 Simulation Studies and Performance Evaluation

In this section, we report on extensive simulation-based studies on the performance impacts of our proposed modifications in the *ns-2* simulator. In these studies, we only consider the hop-by-hop retransmission scenario. We performed studies using both TCP and UDP traffic sources to study the effect of our routing schemes on these transport layer mechanisms. However, due to space constraints we report results for TCP experiments, performed using its NewReno variant. The details of the UDP experiments can be found in [2]. For all these simulation studies, we use link-layer retransmissions to recover from packet losses.

To study the performance of our suggested schemes, we implemented and observed three separate routing algorithms: (a) Energy-Aware (EA) scheme in which link costs are assigned based only the energy requirements for a single transmission attempt across the link. This scheme does not account for energy requirements for retransmissions and is equivalent to [6, 12], (b) Re-transmission Aware (RA) scheme [1] in which link costs also account for energy required for retransmissions, and (c) our proposed Optimal Re-transmission Aware (RA-Opt) scheme which additionally performs optimal adaptation of transmission power.

4.1 Network Topology and Link Error Modeling

For our studies, we used different topologies having upto 100 nodes distributed over on a square region, to study the effects of various schemes on energy requirements and throughputs achieved.

1. *Grid topologies:* For comparison purposes, we first present results on the performance of the schemes on a 100 node grid topology similar to one used in [1]. The nodes on this topology were separated 100 units apart along each axis, and the maximum transmission radius of the node was limited to 150 units. Thus, each node has between 3 and 8 neighboring nodes on this topology [3].

2. *Random topologies:* We also present simulation results for randomly generated topologies. In the random topologies, the nodes were distributed uniformly at random in a 1000×1000 square grid. We experimented with different transmission radii for the nodes. In our random topology generator, we specified the desired number of links [4]. To avoid uni-directional links, we assigned the same transmission radii to all nodes. Note that a hop-by-hop retransmission scheme works only for bi-directional links. In the results presented in this section for the random topologies, we specified the number of wireless links to be one-eighth of a complete graph on these set of nodes. The consequent transmission radii for each node was about 210 units.

Each of the routing algorithms was then run on these 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 topology had 3 active flows, providing a total of 12 flows. In the random topology, we chose 12 random source-destination pairs from the entire set of nodes.

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 2 Mbps. Each of the simulations was run for a fixed duration.

We choose BPSK as our representative modulation scheme and hence, use Equation 5 to derive the bit-error-rate. We varied the ambient noise to obtain different data points. For the non-adaptive transmission power algorithms (EA and RA) we chose a transmit power of 20 mW. The spectral noise for the different channels was chosen to

---

[3] Our-energy aware routing formulation does not directly define a transmission range. It is possible that a longer link with lower receiver noise may consume less effective energy that a shorter link with higher receiver noise. Real-life scenarios, however, impose both an upper bound on the maximum possible transmission power as well as a minimum energy threshold for successful packet reception — if the received power level is below this threshold, no reception is possible even in the absence of any receiver noise. Since any signal suffers channel attenuation $\propto D^4$, the transmission range is an alternative way of assuming that the received power level beyond a distance of 150 units is always lower than the minimum reception threshold, even if the transmitter operates at the maximum power level.

[4] We count each pair of nodes that are within the transmission range of each other as one wireless link.
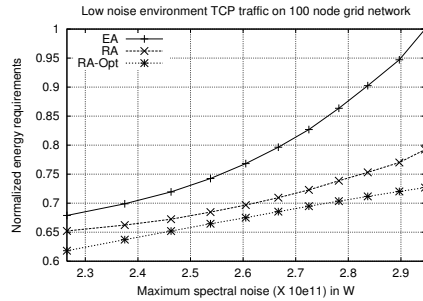
Fig. 2. Normalized energy costs for TCP flows (Low Noise Grid Topology).
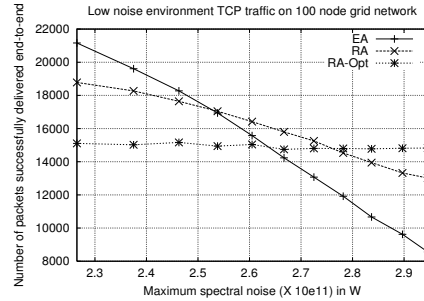

Fig. 3. Comparison of throughput for TCP flows (Low Noise Grid Topology).

vary between two configurable parameters, N_{min} and N_{max} corresponding to minimum and maximum noise respectively. less than a configurable parameter N_{max}.

We simulated two different environments:

1. *Low noise environment:* In this case, we chose N_{min} to be $1.8 \times 10^{-11}W$, while N_{max} was varied between $2.0 \times 10^{-11}W$ and $3.0 \times 10^{-11}W$. For the non-adaptive schemes (EA and RA) a maximum spectral noise of $2.0 \times 10^{-11}W$ leads to a corresponding channel packet error rate of 0.1 on a 100 unit link. Our adaptive transmission algorithm (RA-Opt) appropriately chose a transmission power for each link so that the energy consumption for reliable data transfer across that link is minimized.
2. *High noise Environment:* In this scenario, we chose N_{min} to be $2.8 \times 10^{-11}W$; we varied N_{max} between $3.0 \times 10^{-11}W$ and $4.0 \times 10^{-11}W$.

Our results show that the RA-Opt scheme outperforms the other schemes in environments (other than zero noise environments). Additionally, our scheme shows significant benefits as the noise in the environment increases, as a comparison between these two environments show.

4.2 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 maximum-energy experiment between all the schemes among all these experiments. This provides an easy representation for comparison of the different schemes with each other and with changing maximum spectral noise for different sets of studies.

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.

4.3   Low noise Environment for Grid Topologies

Figure 2 shows the normalized energy consumption for the different schemes for the TCP flows. For example, when N_{max} was set to $2.9 \times 10^{-11}W$, the relative energy requirements of RA-Opt, RA and EA were 0.72, 0.77 and 0.96 respectively. As expected, the energy requirements of all the schemes increase with increase in spectral noise. The EA scheme has the highest energy requirements among all the schemes when the maximum channel noise on links was high. Both the RA and the RA-Opt scheme performs significantly better than this scheme for the entire range of spectral noise. RA-Opt has the best performance among the three different schemes for the entire range of spectral noise. The EA scheme consumes about 10% to 33% more energy per packet, while the RA scheme consumes about 8% to 10% more energy per packet than the RA-Opt scheme,

It is interesting to note that in this low noise environment the energy costs of both the EA and RA schemes have a convexity property, while that of the RA-Opt scheme has a concavity property. This implies that the benefits of the RA-Opt scheme becomes more and more significant with increase in the spectral noise.

In Figure 3 we plot the effective reliable throughput metric for the different schemes in the same low noise experiments. The number of packets transmitted reliably over a fixed duration for the EA scheme falls rapidly with increase in spectral noise. This is expected because the EA scheme does not consider channel properties in choosing routes. In contrast, the number of packets reliably transferred by the RA scheme falls in a more gradual fashion. The decreasing trend in both these schemes is due to the increasing link error rates with the increase in spectral noise. As the link error rates increase, packets sees an increase in end-to-end delays, due to the delays spent in increased number of retransmissions necessary to ensure reliability.

However, the same metric stays relatively constant for the RA-Opt case. This is because, the RA-Opt scheme aggressively adapts the transmission power so as to minimize the energy costs for reliable packet delivery across a link. *The corresponding transmission power to achieve this optimal cost is such that the link error rate stays fairly stable across the entire range of spectral noise.*

4.4   High noise Environment for Grid Topologies

Now we present results for the higher noise environment. Note that in this environment, the value of N_{min} is significantly larger than its corresponding value in the low noise environment. In Figure 4, we plot the normalized energy required per packet in the high noise environment for UDP flows. For example, when N_{max} is set to $3.375 \times 10^{-11}W$ the relative energy requirements of RA-Opt, RA and EA schemes are 0.16, 0.30 (i.e.
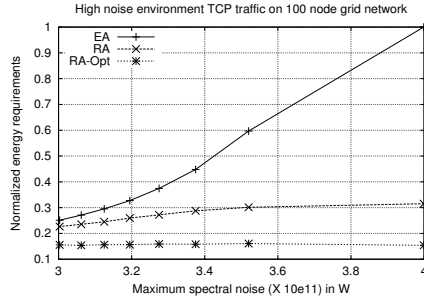
Fig. 4. Normalized energy costs for TCP flows (High Noise Grid Topology).
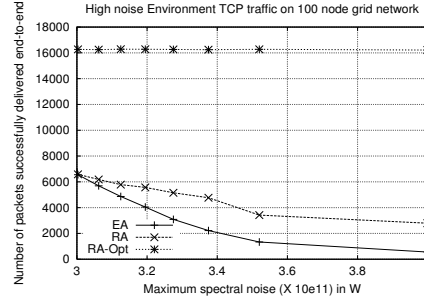
Fig. 5. Comparison of throughput for TCP flows (High Noise Grid Topology).

1.9 times that of RA-Opt) and 0.60 (i.e. 2.75 times that of RA-Opt) respectively. The benefits of the RA-Opt scheme is significantly higher than in the low noise environment (note that the scale of the Y-axis is much larger than the corresponding plots for low noise environments). The RA scheme consumes between 1.5 times to 2 times more energy per packet in this environment than the RA-Opt scheme, while the EA scheme consumes 6.67 times more energy in the worst case in these experiments.

It is interesting to observe the behavior of the throughput achieved by TCP flows in this high noise environment. In trying to optimize the energy consumption, the RA-Opt scheme adapts the transmission power which suitably drives down the channel error rates in all environments. Therefore, the throughput achieved by the RA-Opt scheme is largely unaffected by the noise characteristics. In Figure 5 we plot the throughput for TCP flows in the high noise environment. We can observe that the RA-Opt scheme achieves the same throughput both in the low and high noise environments. Both the EA and RA schemes suffer in the high noise environment, as can be seen in the significant drop in their throughputs achieved.

Both the EA and RA schemes achieve similar throughputs when the maximum spectral noise is $3 \times 10^{-11}W$ with the minimum spectral noise being $2.8 \times 10^{-11}W$. This is because the range of error rates between different links are similar in this scenario, and so RA is unable to choose significantly better paths than EA. RA-Opt is able to make such a choice by increasing the transmission power at nodes to drive down the error rates significantly.

4.5 Random Topologies

We now present results of our studies with randomly generated topologies. Observing that the relative energy requirements for both TCP and UDP flows are similar in nature, we only present the results for the TCP flows in a high noise environment.

In Figure 6, we plot the energy requirements for the different schemes in the high noise environment. The benefits of the RA-Opt scheme is apparent in the plot. The relative energy costs of the RA-Opt, RA and EA schemes are 0.22, 0.49 (i.e. 123% more than RA-Opt) and 0.53 (i.e. 141% more than RA-Opt) respectively. At low noise levels
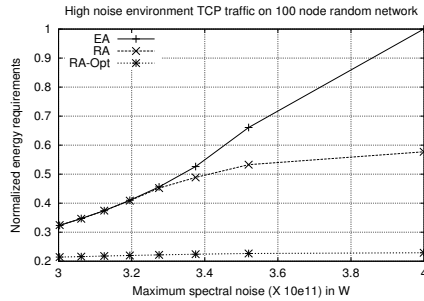
Fig. 6. Normalized energy costs for TCP flows (High Noise Random Topology).
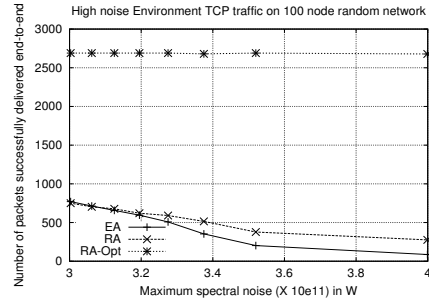
Fig. 7. Comparison of throughput for TCP flows (High Noise Random Topology).

(less than $3.28 \times 10^{-11}W$) the RA scheme and the EA scheme performs equally. However, by adapting the transmission power at the nodes, the RA-Opt scheme performs significantly better than these schemes.

In Figure 7 we show the throughput of the different schemes. Like before, the RA-Opt scheme maintains a stable throughput, while the other schemes see a significant degradation in performance.

5 Related Work

Typical energy-aware routing protocols (e.g. PAMAS [12] and PARO [6]) aim to minimize the total power consumed over the entire transmission path. However, these protocols do not address the problem of reliable data delivery, and hence, ignores the additional energy expenditure due to re-transmissions, if they become necessary. Such a formulation often leads to the formation of a path with a large number of hops. In [11], the authors extend this basic approach by including energy expenditure for packet reception in the link cost.

The basic idea of reducing communication energy costs for packet transmission over individual wireless links has also been explored in the literature. However, the focus has been on the use of intelligent *link scheduling* algorithms, rather than on transmission power control. For example, Zorzi and Rao [15] proposed the use of short, periodic probe packets to detect the condition of the wireless channel. Actual data packet transmissions were deferred when the channel is in a 'bad' state. A similar idea for energy-efficient scheduling of packets from a base station to a set of downstream wireless hosts has been explored in [3, 10, 4].

Ad-hoc routing protocols 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 [9] or DSR [7], can (in principle) be adapted, with suitable modifications, to yield minimum-energy paths by setting the link metric to be a function of the transmission energy.

Apart from minimum energy path problem, research in in energy-aware routing has also focused on other problems with related objectives. For example, *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 [13, 14, 8]. While minimum energy algorithms are most efficient, these network lifetime maximizing schemes are more "fair." A combination of both these approaches can therefore be useful as shown in [14, 8].

Link error probabilities have been considered for single hop spread spectrum links in [5]. In contrast, we focus on end-to-end energy costs for multi-hop wireless networks.

6 Conclusions

We have defined an optimal energy routing scheme for reliable data transfer on a multi-hop wireless network. This scheme outperforms well-known existing routing schemes for a wide range of channel characteristics. This improvement in energy efficiency is achieved by explicitly considering the impact of receiver noise on packet errors, and by adjusting the transmission power to minimize the total energy spent in reliably forwarding a single bit. The scheme is general across different modulation techniques that can be employed for data delivery on wireless networks.

We investigated the issue existence of an optimal transmission energy for a given link. By employing information-theoretic bounds, we have shown how any link is associated with a fundamental lower bound on energy efficient reliable communication. Moreover, this fundamental bound is directly proportional to the channel attenuation rate ($\propto D^{\alpha}$), and is achieved by choosing arbitrary low transmission power. In contrast, practical communication systems are associated with a well-defined optimal transmission power, such that any decrease or increase from this optimal value results in a sharp increase in the total transmission energy spent in reliable data transfer.

We have studied the applicability of this technique in choosing optimal energy paths. By appropriately choosing the cost metric, it is possible to optimize other objective functions, e.g. end-to-end latencies, data delivery throughput, etc. We, therefore, believe that our scheme has a wider applicability to a range of operating modes depending on the the optimization objectives. Simulation studies indicate that performing adaptive power control based on the individual link conditions (error rates) can provide energy savings of $\approx 10\%$ in low-noise environments, and as much as $\approx 40\%$ in high-noise environments.

The analysis in this paper assumed that the use of retransmissions as the sole means of providing a reliable link layer. As discussed earlier, the fundamental technique can, however, also be applied to alternate reliability schemes such as forward error correcting codes through appropriate changes to the relationships in Equations 4 and 3. Since our power adaptation mechanism implicitly relies on relative stable variations in the packet error-rate, this technique is especially useful in static, or low-mobility, multi-hop networks, where link parameters such as distance or attenuation coefficients do not exhibit very rapid changes. (Of course, our formulation is applicable in the presence of typical wireless environment effects such as fading; the average bit error rate at the link-layer is typically a more stable statistical metric obtained by averaging over such physical layer

variations). We have also assumed the existence of appropriate MAC-layer contention resolution mechanisms for common-channel networks, which present an abstraction of zero-interference to the higher layers. Such an abstraction is also provided by the use of distinct physical channels based on TDMA/FDMA/CDMA techniques. Finally, our approach to energy optimization is useful not just for forming energy-efficient routing paths, but also for independently optimizing the transmission energy on each individual link. The routing algorithm and individual link-layer power control techniques can operate on different time scales; while routes can be re-computed over longer time periods, individual transmitters can adjust their link transmission power over shorter time-scales.

References

1. S. Banerjee and A. Misra. Minimum energy paths for reliable communication in multi-hop wireless networks. In *Proceedings of ACM Mobihoc*, June 2002.
2. S. Banerjee and A. Misra. Power adaptation based optimization for energy efficient reliable wireless paths. Technical Report, Department of Computer Sciences, University of Wisconsin-Madison, WI 53706, USA, Available at http://www.cs.wisc.edu/~suman/pubs/optpower-tr.pdf, 2004.
3. P. Bhagwat, P. Bhattacharya, A. Krishna, and S.K. Tripathi. Using channel state dependent packet scheduling to improve tcp throughput over wireless lans. *Wireless Networks*, 3(1), 1997.
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. J.H. Gass Jr., M.B. Pursley, H.B. Russell, and J.S. Wysocarski. An adaptive-transmission protocol for frequency-hop wireless communication networks. *Wireless Networks*, 7(5):487–495, September 2001.
6. J. Gomez and A. Campbell. Power-aware routing optimization for wireless ad hoc networks. In *High Speed Networks Workshop (HSN)*, June 2001.
7. D. Johnson and D. Maltz. Dynamic source routing in ad hoc wireless networks. In *Mobile Computing*, pages 153–181, 1996.
8. A. Misra and S. Banerjee. Mrpc: Maximizing network lifetime for reliable routing in wireless environments. In *Proceedings of IEEE Wireless Communications and Networking Conference*, March 2002.
9. C.E. Perkins and E.M. Royer. Ad-hoc on-demand distance vector routing. In *Proceedings of the 2^{nd} IEEE Workshop on Mobile Computing Systems and Applications*, February 1999.
10. 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*, April 2001.
11. K. Scott and N. Bamboos. Routing and channel assignment for low power transmission in PCS. In *Proceedings of ICUPC*, October 1996.
12. S. Singh and C.S. Raghavendra. Pamas-power aware multi-access protocol with signaling for ad hoc networks. In *ACM Communications Review*, July 1998.
13. S. Singh, M. Woo, and C.S. Raghavendra. Power-aware routing in mobile ad-hoc networks. In *Proceedings of Mobicom*, October 1998.
14. C.K. 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.
15. M. Zorzi and R. Rao. Error control and energy consumption in communications for nomadic computing. *IEEE Transactions on Computers*, 46(3), March 1997.