
A Cross-Layer Approach to Energy Balancing in Wireless Sensor Networks

Daniele Puccinelli, Emmanuel Sifakis, and Martin Haenggi

Department of Electrical Engineering
University of Notre Dame {dpuccine,esifakis,mhaenggi}@nd.edu

Summary. We consider a *many-to-one* real-time sensor network where sensing nodes are to deliver their measurements to a base station under a time constraint and with the overall target of minimizing the energy consumption at the sensing nodes. In wireless sensor networks, the unreliability of the links and the limitations of all resources bring considerable complications to routing. Even in the presence of static nodes, the channel conditions vary because of multipath fading effects due to the motion of people or objects in the environment, which modify the patterns of radio wave reflections. Also, sensing nodes are typically battery-powered, and ongoing maintenance may not be possible: the progressive reduction of the available energy needs to be factored in.

The quality of the links and the remaining energy in the nodes are the primary factors that shape the network graph; link quality may be measured directly by most radios, whereas residual energy is related to the node battery voltage, which may be measured and fed into the microcontroller. These quantities may be used to form a cost function for the selection of the most efficient route. Moreover, the presence of a time constraint requires the network to favor routes over a short number of hops (*a.k.a.* the long-hop approach, in the sense that a small number of long hops is used) in order to minimize delay. Hop number information may be incorporated into the cost function to bias route selection toward minimum-delay routes. Thus, a cross-layer cost function is obtained, which includes raw hardware information (remaining energy), physical layer data (channel quality), and a routing layer metric (number of hops).

A route selection scheme based on these principles intrinsically performs node energy control for the extension of the lifetime of the individual nodes and for the achievement of energy balancing in the network; intuitively, the long-hop approach permits the time-sharing of the critical area among more nodes. A novel, practical algorithm based on these principles is proposed with the constraints of the currently available hardware platforms in mind. Its benefits are investigated with the help of computer simulation and are illustrated with an actual hardware implementation using Berkeley motes.

1 Introduction

A large portion of wireless sensor network research over the past few years has concentrated on routing [1]. In this extensive body of work, the most common metric used for the assessment of routing quality is the number of hops in a route [7]. However, nodes in wireless sensor networks are typically static, and link quality has to be taken into account to avoid routing over lossy links. In the literature on ad hoc networks, various metrics are suggested in order to take packet loss into account. An interesting example is ETX (Expected Transmission Count) [2], which uses per-link measurements of packet loss ratios in both directions of each wireless link to find the high throughput paths which require the smallest number of transmissions to get a packet to destination. Since sensor nodes have heavy energy constraints, the battery capacity of individual nodes should also be taken into account. Battery-aware routing algorithms [10] typically try to extend the lifetime of the network by distributing the transmission paths among nodes with greater remaining battery resources. The key observation is that minimum-energy routes can often unfairly penalize a subset of the nodes. Most efforts in this direction are targeted toward ad hoc networks, and are often not portable to wireless sensor networks, which are the focus of the present work.

The distinguishing features of wireless sensor networks make routing much more challenging. Wireless sensor networks often need to self-organize and operate in an untethered fashion. Sensor nodes are usually static, but location information is often unavailable, and geographic routing is thus not an option. Typically, constraints in terms of energy, processing, and storage capacities are extremely tight: resources need to be carefully managed, and lifetime extension of the sensing nodes is a major concern. The adoption of a *many-to-one* traffic scheme is the primary cause of energy unbalance leading to premature discontinuation of node activity. If the base station is not located within the reach of a source, a multihop scheme needs to be adopted, and other nodes are used as relays to guarantee a connection from that source to the base station. However, if many sources are in that situation, the nodes directly within the reach of the base station are located on most forwarding paths, have an increased workload, and their lifetime is likely to be shortened. In the special case of *one-to-one* traffic, schemes that find and rely on optimal routing paths also cause energy unbalancing, as they unevenly distribute the workload across the network and shorten the lifetime of the nodes along the optimal path. In [9], the occasional use of suboptimal routes is suggested as a countermeasure; this solution is shown to yield a significant lifetime improvement with respect to methods based on optimal routes alone.

An efficient attempt to balance the energy distribution in the network is necessarily reliant upon some form of monitoring of the quality of the links which allows an assessment of connectivity. The remaining battery power is also of interest, since we aim at having a reasonably uniform energy balancing. Our main goal in this paper is to investigate the benefits of a joint metric

which considers hop count, link quality, and remaining battery power. We look at the routing problem in wireless sensor networks from a cross-layer perspective, by introducing an algorithm to build a routing tree by means of a heuristic metric encompassing physical layer information such as link strength and battery power as well as routing layer information (number of hops). The scenario is as follows: in a wireless sensor network comprising of N severely energy-constrained (battery-operated) lower-end sensing nodes, we intend to use fairly efficient routes for node lifetime extension; hence, we propose a strategy with energy-balancing guarantees. We assess its validity and benefits by means of simulation, and describe a hardware implementation with the most recent generation of Berkeley nodes, MICAz.

The remainder of the paper is organized as follows. Section 2 describes our algorithm and illustrates it with examples obtained with a custom network simulator. Section 3 presents an evaluation of the performance of our scheme obtained by means of extensive simulation, and Section 4 details the hardware implementation. Section 5 includes some closing remarks.

2 A Routing Scheme with Energy-Balancing Guarantees

From a dynamic programming perspective, the problem of routing from a source s to a destination node d corresponds to defining a policy at a generic node r_i for the choice of a relay r_{i+1} in the direction of d given the list of nodes $s, r_0, r_1, \dots, r_{i-1}$. Routing over many short hops (*e.g.*, nearest-neighbor routing) enjoys a lot of support. From an energy-consumption viewpoint, it is claimed that dividing a hop over a distance d into n short hops yields an energy benefit of $n^{\alpha-1}$, where α is the path loss exponent [8]. Unfortunately, this comes from the use of oversimplified link models such as the disk model [3]: a transmission can either fail or succeed depending on whether the distance is larger or smaller than the so-called transmission radius. This model totally ignores fading and is therefore inaccurate for a realistic modeling of the wireless medium [8]. It is shown in [4] that for a block fading channel multihop routing does not offer any energy benefits for $\alpha=2$. Routing over fewer (but longer) hops carries indeed many advantages [5]; from our point of view, the most interesting are the reduction of energy consumption at the sensing nodes, the achievement of a better energy balancing, the more aggressive exploitation of sleep modes, and the lack of route maintenance overhead. If the sensing nodes only occasionally need to act as relays, they can sleep longer and only consume energy to make their own data available. With these ideas in mind, we propose a lightweight multihop routing scheme with energy-balancing guarantees. Our scheme is an example of flat routing, in the sense that all nodes are assigned similar roles (hierarchies are not defined). It can be considered dynamically proactive, since all routes are computed prior to being used and periodically updated.

In a network of N nodes ($N - 1$ sensing nodes and a base station), the generation of the routing tree is performed as follows:

- Every sensing node sends a test packet to the base station; due to the unreliability of wireless links, the base station is only bale of receiving a fraction of such test packets.
- The base station uses the received test packets to measure the quality of the links to the sensing nodes, and feeds this information back to them.
- The sensing nodes that receive a reply from the base station broadcast a route setup packet to advertise their nearest-neighbor status as well as the quality of their link to the base and their battery voltage. Route setup packets may be seen as pointers toward the direction of the base station.
- Other sensing nodes that receive these packets generate and broadcast similar packets to advertise that they are two hops away from the base station.
- The information travels upstream from the base station into the network until all nodes know their depth and the tree is fully defined. The route setup is targeted at downstream communication, and its effectiveness is reliant upon link symmetry. In environments where asymmetric links are abundant, link quality estimation from reverse link quality information often does not work and handshakes between nodes are necessary.

Table 1. List of symbols.

Symbol	Meaning
A_i	Address of node i .
$l_{i \rightarrow k}$	Quality of the link between node i and node k .
$M_{i,k}$	Metric of the route found at node i going through k .
L_i	$\min(l_{i \rightarrow j}, L_j)$, $M_{i,j} = \max_{0 \leq k \leq N-1} M_{i,k}$, $i \neq k$. $L_{BS} \triangleq 1$.
D_i	Depth of node i .
v_i	Battery voltage of node i .
V_i	$\min(v_i, V_k)$, $M_{i,k} = \max_{0 \leq j \leq N-1} M_{i,j}$, $i \neq j$. $V_{BS} \triangleq 1$.
$l[i, k]$	$\min(l_{i \rightarrow k}, L_k)$ (k -th entry in the table internal to node i).

We integrate a route selection scheme into the tree formation procedure by adding the following operations:

- Each sensing node maintains a table indicating the nodes it can reach, the quality of the links to such nodes, their battery voltage, and their depth. All this information can be inferred by the route setup packets mentioned above. In particular, the depth can be inferred by using a counter initialized to 0 in the setup packets sent by the base station. The nearest neighbors of the base station will thus set it to 1, their nearest neighbors will increment it to 2, and so forth.

Table 2. Setup packet sent by the base station.

A_{bs}	L_{bs}	V_{bs}	D_{bs}
bs	1	1	0

The setup packet from the base station has the default structure shown in Table 2. All symbols are explained in Table 1.

Table 3. Structure of the internal table for node i , which is able to communicate with nodes k and j .

A	L	V	D	M
k	$\min(l_{i \rightarrow k}, L_k)$	V_k	$D_k + 1$	$M_{i,k}$
j	$\min(l_{i \rightarrow j}, L_j)$	V_j	$D_j + 1$	$M_{i,j}$

Table 4. Structure of a setup packet sent by node i after choosing to route packets over node k .

A_i	L_i	V_i	D_i
i	$\min(l_{i \rightarrow k}, L_k)$	$\min(v_i, V_k)$	$D_k + 1$

- Sensing node i at depth d_i receives a setup packet from node k at depth d_k . In the remainder of the paper, we assume the quantities L and V to be normalized as to be comprised within the interval $[0, 1]$. If $d_k \geq d_i$, the setup packet from node k is discarded (it does not point toward the base station). If $d_k < d_i$, on the other hand, the k -th entry in the table internal to node i is processed as displayed in Table 4. It should be observed that the base station does not need to maintain an internal table. The rationale behind these operations is to keep track of possible routes to the base station and be able to order them on the basis of a joint metric favoring good links, relays with abundant energy resources, and low number of hops. By keeping track of the minimum link quality and the lowest voltage in the route, bottlenecks may be identified.
- In node i , a metric is computed for each entry in the internal table. For instance, a possible metric for the k -th entry is given by

$$M_{i,k} = \frac{\min(l_{i \rightarrow k}, L_k) + V_k + \frac{1}{D_k + 1}}{3}, \tag{1}$$

where k is a downstream nearest neighbor of i . This metric favors energy balancing: the number of hops from node i to the base station through

node k is given by $D_k + 1$, and using its inverse in the metric results into a larger metric for routes over a small number of hops. Node i operates the decision to route over the node whose entry in the table has the largest metric. We will use this metric in the numeric examples.

This procedure is to be repeated periodically for dynamic route maintenance and update (in case of topology changes, *e.g.* due to mobility or to the death of one or more nodes).

The effectiveness of this approach is reliant upon the existence of a solid MAC scheme minimizing in-network interference.




Fig. 1. An example of a sensor network with 5 nodes. Node 5 represents the base station. The numbers near the links indicate their quality, whereas the numbers below the nodes represent the battery voltage.

We will now show a simple example of a network with $N=5$ nodes (shown in Figure 1) in order to clarify how the algorithm works.

- Nodes 1 through 4 send a test packet to node 5. Due to the features of the wireless channel, node 5 can only communicate with nodes 3 and 4: the test packets sent by the other nodes are not received by node 5.
- Node 5 sends a setup packet to 3 and 4; this packet is shown in Table 5.

Table 5. Setup packet sent by node 5.

A_5	L_5	V_5	D_5
5	1	1	0

- Node 4 fills the row pertaining to node 5 in its internal table as follows:

$$l[4, 5] = \min(l_{4 \rightarrow 5}, L_5) = 0.9 \tag{2}$$

$$V_5 = 1 \tag{3}$$

$$D_5 + 1 = 1 \tag{4}$$

$$M_{4,5} = \frac{0.9 + 1 + 1}{3} = 0.97. \tag{5}$$

Table 6. Setup packet sent by node 4.

A_4	L_4	V_4	D_4
4	$l[4, 5]=0.9$	$\min(v_4, V_5) = 0.6$	1

On the basis of the table, the packet shown in table 6 is sent. Since there is only one possible route (node 5 is the only entry in the internal table), $L_4 = l[4, 5]$.

- Node 3 fills the row pertaining to node 5 in its internal table as follows:

$$l[3, 5] = \min(l_{3 \rightarrow 5}, L_5) = 0.4 \tag{6}$$

$$V_5 = 1 \tag{7}$$

$$D_5 + 1 = 1 \tag{8}$$

$$M_{4,5} = \frac{0.4 + 1 + 1}{3} = 0.8. \tag{9}$$

On the basis of the table, the packet shown in table 7 is sent. Since there is only one possible route (node 5 is the only entry in the internal table), $L_3 = l[3, 5]$.

- Only nodes 2 and 5 receive the test packets from 3 and 4. Node 5 does not process them, as they come from higher-depth nodes.

Table 7. Setup packet sent by node 3.

A_3	L_3	V_3	D_3
3	$l[3, 5]=0.4$	$\min(v_3, V_5) = 0.2$	1

- Node 2 fills the rows pertaining to nodes 3 and 4 on the basis of the received setup packets. The row for node 3 is filled as follows:

$$l[2, 3] = \min(l_{2 \rightarrow 3}, L_3) = 0.4 \quad (10)$$

$$V_3 = 0.2 \quad (11)$$

$$D_3 + 1 = 2 \quad (12)$$

$$M_{2,3} = \frac{0.4 + 0.2 + 0.5}{3} = 0.37. \quad (13)$$

The row for node 4 is filled as follows:

$$l[2, 4] = \min(l_{2 \rightarrow 4}, L_4) = 0.1 \quad (14)$$

$$V_4 = 0.6 \quad (15)$$

$$D_4 + 1 = 2 \quad (16)$$

$$M_{2,4} = \frac{0.1 + 0.6 + 0.5}{3} = 0.4 \quad (17)$$

Table 8. Setup packet sent by node 2.

A_2	L_2	V_2	D_2
2	$l[2, 4]=0.1$	$\min(v_2, V_4) = 0.6$	2

The route with the highest metric, namely the one related to entry 4, is selected, $L_2 = l[2, 4]$, and the test packet shown in Table 8 is sent out.

- Only nodes 1, 3, and 4 receive the test packet from 2. Nodes 3 and 4 do not process them, as they come from a higher-depth node.

- Node 1 fills the row pertaining to node 2 in its internal table as follows:

$$l[1, 2] = \min(l_{1 \rightarrow 2}, L_2) = 0.1 \quad (18)$$

$$V_2 = 0.6 \quad (19)$$

since node 2 has chosen route (2, 4, 5), and node 4 ($v_4 = 0.6$) is the bottleneck of route (1, 2, 4, 5) in terms of battery voltage.

$$D_2 + 1 = 3 \quad (20)$$

$$M_{1,2} = \frac{0.1 + 0.6 + 0.3}{3} = 0.33. \quad (21)$$

- Node 1 is thus able to reconstruct the route to node 5: the first step is node 2, the only entry in the table internal to node 1. Node 2 knows that packets need to be routed over node 4, whose entry has the best metric. Node 4 is a nearest neighbor of the base station, so the chosen route is (1, 2, 4, 5).

3 Performance Analysis

We have used a custom network simulator to assess the validity of our routing scheme in a scenario with $N=100$ nodes uniformly randomly deployed in a $30\text{m} \times 30\text{m}$ environment. The source and destination are placed at the opposite ends of one of the diagonals of the square deployment area. We will show that the use of our proposed joint metric can lead to a significant gain in terms of energy balancing with respect to approaches that only consider link quality. A *one-to-one* scenario with one node trying to communicate to the base station is simulated. The simulator evaluates link quality on the basis of a multipath channel model with additive Gaussian noise.

Our simulator initially finds a total of 72 routes: 8 over 3 hops, and 64 over 4 hops. Routes with more hops are discarded. Figure 2 shows the performance of a routing scheme using bottleneck link quality as the only metric for route selection. The route with the best bottleneck link quality is chosen, which does not at all guarantee energy-balancing: each chosen route is perused until its relaying nodes are fully depleted. The use of a joint metric between node k and node 0 along the route $(k, k-1, \dots, 1, 0)$ obtained as

$$M_k = \frac{\min(L_{k \rightarrow k-1}, \dots, L_{1 \rightarrow 0}) + \min(V_{k-1}, \dots, V_0) + \frac{1}{d_k}}{3} \quad (22)$$

dramatically improves energy balancing: Figure 3 shows that the death of the first node is delayed from step 12 to step 118. The impact of this result can be better appreciated by considering that a wireless sensor network typically




Fig. 2. The number of 3-hop and 4-hop routes steadily decreases as the batteries of the nodes are depleted: a routing scheme based on bottleneck link quality does not ensure a proper energy balancing.

follows a *many-to-one* traffic pattern: keeping as many nodes alive for as long as possible is extremely important for the sensing capabilities of the network.

Other metrics are also possible. A stronger emphasis may be placed on either energy balancing or link quality by using a weighted average of the type

$$M_k = \alpha \min(L_{k \rightarrow k-1}, \dots, L_{1 \rightarrow 0}) + \beta \min(V_{k-1}, \dots, V_0) + (1 - \alpha - \beta) \frac{1}{d_k}. \quad (23)$$

Figure 4 shows the performance of our scheme with a weight assignment of the form $\alpha=1/6$ and $\beta=1/2$, which places an even stronger emphasis on energy balancing, and further delays the death of the first node (from step 118 to step 131). With both this weighted metric and the unweighted bottleneck metric, a gain of an order of magnitude can be achieved with respect to the scheme based on bottleneck link quality. These gains are of course upper bounds, as in our simulation we assume that the energy consumption is negligible if the nodes are not actively relaying packets. These upper bounds can be approached with an aggressive use of sleep modes and low-power listening techniques which reduce the receive energy, which is normally comparable to the transmit energy.

We have seen the advantages of a joint metric, but the analysis above is based on the idea of using bottleneck quantities. Figure 5 shows that in-




Fig. 3. Our joint metric ensures better energy balancing: routes start to become unusable after 118 steps.

corporating average route link quality and average route battery voltage (as opposed to minima) into the selection metric fails to provide energy balancing and leads to disadvantages such as the premature loss of shorter routes. Bottlenecks should not be ignored: if a K -hop route includes $K - 1$ good hops and a bottleneck, it will have a good average but will also prematurely cease to exist if selected.

4 Hardware Implementation

We implemented our energy-aware routing scheme in hardware using the low power MICAz platform. MICAz represents the latest generation of Berkeley motes, and is commercialized by Crossbow. It is built around an ATMega128L microprocessor, and features a CC2420 802.15.4-compliant ZigBee-ready radio. IEEE 802.15.4 is a standard for low-rate, wireless personal area networks which provides specification for the physical and the MAC layer. At the physical layer, it defines a low-power spread spectrum radio operating at 2.4GHz with a bit rate of 250kb per second. ZigBee is a collection of high level communication protocols built on top of the IEEE 802.15.4 MAC layer.

We implemented the algorithm with the operating system TinyOS [6]; our code has a reasonably small footprint (about 10KB of ROM occupancy, and less than 500B of RAM occupancy).

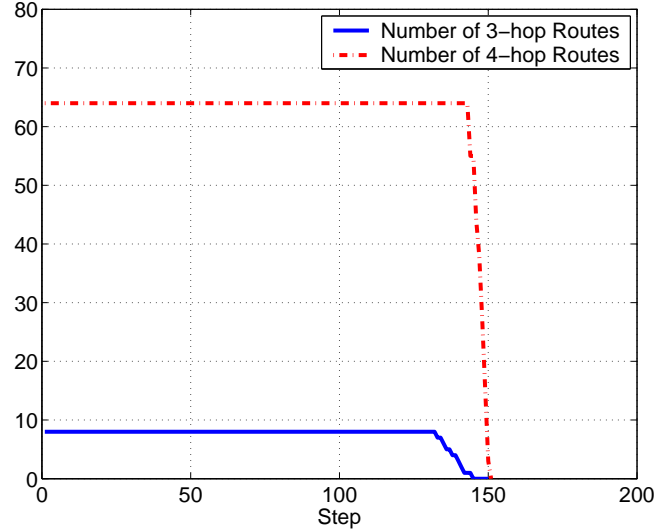


Fig. 4. The use of weights within the metric can lead to even better energy balancing results.

We placed 10 MICAz nodes in a lab environment in an arrangement shown in Figure 6. This network is to deliver a continuous stream of packets at a rate of 46 packets/min from node 1 to node 10 over a duration of 105 hours and 20 minutes (the time that was needed for the radio of the source node to become unusable due to the overly low voltage). Although the nodes themselves are static, people moving in the lab cause a certain amount of fading. Fresh batteries were used, the transmit power is -15dBm for both control and application traffic, the MAC scheme is standard CSMA/CA, and no acknowledgments are used to make links reliable. In this implementation, unused nodes sleep 66% of the time, and low power listening modes are not used. The lifetime gain in this experiment can therefore be expected to be fairly far from the upper bounds, but can still provide us with a valid proof of concept.

In the 105h experiment, 290720 packets were transmitted and 256313 were successfully received at the BS; this corresponds to a packet loss of about 12%. Given that there were no retransmissions at all and the transmit power was relatively small, this loss is quite acceptable. Among all the routes found, 9% were single-hop, 89% were two-hop, and 1.8% three-hop. The mean path loss over a distance of about 8m and an obstructed line-of-sight path prevents packets from being received at a transmit power of -15dBm. So the fact that single-hop routes exist indicates that the algorithm exploits positive fading states, i.e., is *opportunistic*, thereby allowing all relay nodes to sleep for some




Fig. 5. The use of average quantities (as opposed to bottleneck quantities) leads to the premature loss of 3-hop routes.

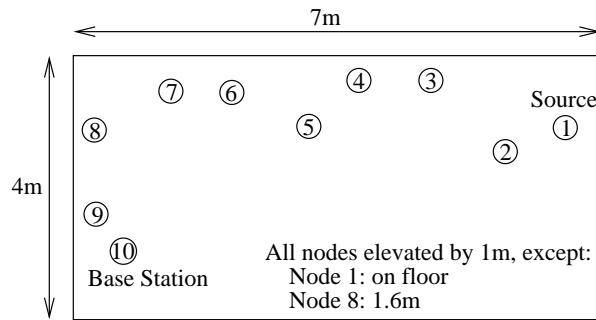


Fig. 6. Setup of the 10-node experiment in an indoor environment.

time. More precisely, one-hop routes dominate between hours 58 and 64; in this time lapse, the room was empty, and the particular arrangement of the furniture created a pattern of reflections that placed the destination in a good *fading spot* with respect to the source, thereby allowing one-hop communication. As soon as people entered the room again, the particular arrangement that had created this situation was modified, and the source was no longer able to exploit static fading to reach the destination directly.




Fig. 7. Battery discharge curves for the 10 nodes.

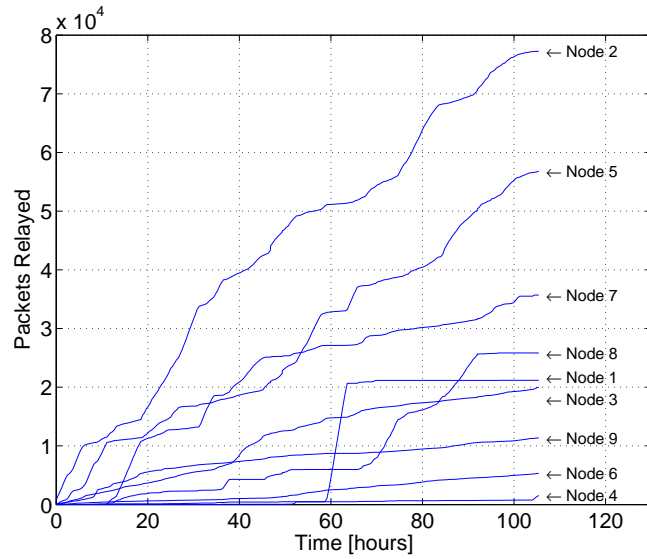


Fig. 8. Cumulated number of packets relayed by each node.

With short-hop routing, the discharge curves of the nodes would all lie below the one for the source node in Figure 7, since all the relays would

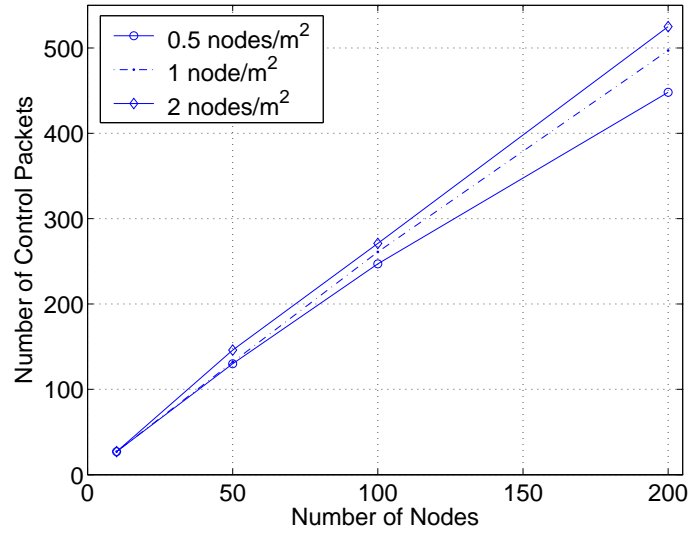


Fig. 9. Number of control packets as the network size increases.

not only transmit each packet but also receive each packet, and the receive energy is substantial. The use of sleep modes is certainly not aggressive in the algorithm used; by increasing the length of the sleeping periods, the high gains shown by our simulations can be approached. This is easier in networks with a high number of nodes, as the chances of losing connectivity because of the excessive use of sleep modes are reduced. This shows that our scheme has very interesting scalability properties; our long-hop approach enhances such properties, in the sense that with as few long hops as possible the volume of routing traffic is minimized. The number of control packets for route discovery is related to the number of nodes by the relation $N - 1 + 2n_1 + \sum_{i=2}^M n_i$, where N is the number of nodes, M is the maximum depth in the routing tree, and n_i is the number of nodes which can reach the base station in i hops. In fact: 1) $N - 1$ nodes send a test packet to the base station; 2) only n_1 nodes can actually reach it, and the base station thus sends n_1 feedback packets to such nodes; 3) these n_1 nodes broadcast one route setup packet each; 4) so do their n_2 nearest neighbors; 5) this goes on until the n_M nearest neighbors at M hops from the base station are reached. Figure 9 shows that the number of control packets as predicted by our simulator for three different node densities scales approximately linearly with the number of nodes, which is again promising for scalability.

5 Closing Remarks

The main contribution of the present work is the introduction of a new approach to route selection in wireless sensor networks. With the constraints of real hardware in mind, we suggest the use of a joint metric as part of a cross-layer approach to achieve energy-balancing. Simulation results show a gain of up to one order of magnitude in node lifetime extension with respect to routing schemes based on link quality. A successful hardware implementation with MICAz motes is indicative of the lightweight nature of our scheme. Control applications would greatly benefit from a scheme that relies on a reduced number of hops because of the inherent delay benefit, and the dramatic extension of network lifetime provided by our routing strategy is extremely appealing.

Acknowledgements

The support of NSF (grants ECS 03-29766 and CAREER CNS 04-47869) is gratefully appreciated.

References

1. J. N. Al-Karaki and A. E. Kamal. Routing Techniques in Wireless Sensor Networks: A Survey. *IEEE Wireless Communications*, 11:6–28, Dec 2004.
2. D. S. J. D. Couto, D. Aguayo, J. Bicket, and R. Morris. A High Throughput Path Metric for MultiHop Wireless Routing. In *Proceedings of the 9th Annual International Conference on Mobile Computing and Networking (MobiCom'03)*, San Diego, CA, Sept. 2003.
3. E. N. Gilbert. Random plane networks. *Journal of SIAM*, 9:533–543, Oct. 1961.
4. M. Haenggi. Twelve Reasons not to Route over Many Short Hops. In *IEEE Vehicular Technology Conference (VTC'04 Fall)*, Los Angeles, CA, Sept. 2004.
5. M. Haenggi and D. Puccinelli. Routing in Ad Hoc Networks: A Case for Long Hops. *IEEE Communications Magazine*. To appear in November 2005.
6. J. Hill, R. Szewczyk, A. Woo, S. Hollar, D. Culler, and K. Pister. System architecture directions for network sensors. In *Proceedings of the 9th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'00)*, Cambridge, MA, Nov. 2000.
7. D. B. Johnson and D. A. Maltz. Dynamic Source Routing in Ad-Hoc Wireless Networks. In T. Imielinski and H. Korth, editors, *Mobile Computing*. Kluwer Academic Publisher, 1996.
8. T. S. Rappaport. *Wireless Communications*. Prentice Hall, Upper Saddle River, NJ, 2001.
9. R. Shah and J. Rabaey. Energy aware routing for low energy ad hoc sensor networks. In *Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC'02)*, March 2002.

10. S. Singh, M. Woo, and C. Raghavendra. Power-aware routing in mobile ad hoc networks. In *Proceedings of the 4th Annual International Conference on Mobile Computing Networking (MobiCom'98)*, Dallas, TX, Oct. 1998.