

The Energy-Robustness Tradeoff for Routing in Wireless Sensor Networks

Bhaskar Krishnamachari[†], Yasser Mourtada^{*}, and Stephen Wicker^{*}

[†]Department of Electrical Engineering, University of Southern California, Los Angeles, CA 90089, USA
bhaskark@ieee.org

^{*}School of Electrical and Computer Engineering, Cornell University, Ithaca, NY 14853, USA
{yasser,wicker}@ece.cornell.edu

Abstract—Wireless sensor networks consisting of large numbers of inexpensive energy-constrained nodes are an area of emerging networking research. Routing algorithms in these networks are required to provide tolerance to temporary or lasting faults in individual devices. The conventional methodology is to set radio transmit powers to the minimum levels required for connectivity and use multipath routing to provide robustness. We show in this paper through an analytical example and detailed simulation results that using a single path routing scheme with higher transmit power can also be an energy-efficient solution for robustness to node failures.

I. INTRODUCTION

Wireless sensor networks consisting of large numbers of inexpensive energy-constrained devices are expected to find a wide range of applications from vehicle tracking to habitat monitoring [1], [2], [3], [4]. Although sensor networks are primarily static in nature, the use of inexpensive devices is likely to result in higher rates of failures for individual nodes. It is therefore important for routing algorithms in this space to provide tolerance to such failures in an energy-efficient manner. This is the subject of our work.

One basic solution for robustness that has been proposed with several variations is multipath routing: the use of multiple disjoint or partially disjoint routes to convey information from source to destination.

There is considerable prior literature on multipath routing techniques which date back to work pertaining to the telephone system where they are used to minimize call blocking. The solutions for multipath routing in all kinds of networks have primarily aimed at providing a set of low-cost disjoint paths between the source and destination [12], [13], [14]. In recent years there has been a focus on multipath routing in mobile ad-hoc networks (MANETs), where the primary concern is path failure due to mobility. For example, alternate path routing is investigated in [7] as a mechanism for load balancing and protection against route failure in MANETs. An on-demand multipath routing scheme is proposed in [8] as a means to reduce

This work was supported by the Defense Advanced Research Projects Agency (DARPA), and administered by the Army Research Office under ESP MURI Award No. DAAD19-01-1-0504. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the Defense Advanced Research Projects Agency (DARPA), and Army Research Office.

query floods. An efficient heuristic scheme for selecting multiple reliable paths in MANETs is presented in [10]. Multipath algorithms for wireless networks are also proposed and studied in [6]. In the specific context of wireless sensor networks, the Directed Diffusion algorithm [5] is a routing protocol that allows for multiple alternate paths to be maintained by setting appropriate gradient levels. Partially disjoint multipath routing schemes described as “braided multipath” schemes for wireless sensor networks are studied in [9]. The energy-robustness tradeoff is also studied in [9], but with a focus on distinguishing between complete disjoint multipath routing and the braided multipath scheme.

The basic idea behind multipath routing is fault-tolerance through redundancy. An alternative philosophy for fault-tolerance is to minimize the number of failure modes by reducing the number of intermediate nodes prone to failure. In sensor networks this can be done even with single-path routing algorithms utilizing higher transmission ranges. We will show in this paper that there are in fact situations in which this alternative strategy is superior to multipath routing in terms of energy-efficiency.

The rest of this paper is organized as follows. In section II, we consider a simple scenario involving a network of five sensor nodes and perform some analytical calculations on the energy and robustness metrics associated with various routing configurations. The example yields insight into why single path routing with higher transmit powers can potentially be a more energy-efficient mechanism for providing robustness. We then turn more detailed simulations involving 50 sensors. The experimental setup for these simulations is described in section III. The results from these simulations are then presented and discussed in section IV. We conclude with a brief discussion in section V.

II. ILLUSTRATION

We begin our exploration of the energy-robustness tradeoffs by considering a simple, small configuration of five sensor nodes. The nodes are placed as seen in figure 1.

Each of the configurations shown in figure 1 represents a possible way to route information from the source to the receiver. If we assume that nodes can only communicate with other nodes

ing to the scheme H_8 is clearly the best strategy as far as robustness is concerned. However this does address the issue of energy-efficiency.

Since we are concerned with two different objectives: minimizing energy while maximizing robustness, it is helpful to make use of the notion of *Pareto optimality*. A routing scheme H_i is said to *dominate* a routing scheme H_j if it results in an equal or greater robustness level with strictly less energy cost or if it results in an equal or lesser energy cost with strictly higher robustness level, i.e. if $\Pi_{H_i} \geq \Pi_{H_j}, E_{H_i} < E_{H_j}$, or if $E_{H_i} \leq E_{H_j}, \Pi_{H_i} > \Pi_{H_j}$. Routing schemes which are not dominated by others in the set of considered schemes are said to be *Pareto optimal* and constitute the Pareto set.

From Table I, we see that for $\alpha = 2$, the Pareto set is $\{H_1, H_3, H_8\}$. What is remarkable is that in this particular case *all* the Pareto optimal routing strategies are single path routes. The multipath routing scheme H_2 is dominated by H_3 which provides greater robustness for the same energy level, and the multipath routing schemes H_4, H_5, H_6, H_7 are dominated by H_8 which provides greater robustness for less energy cost. Since the energy costs depend on the path loss exponent α , the Pareto set is also dependent on this parameter. For $\alpha = 4$, as seen in Table I, the Pareto optimal routing schemes are $\{H_1, H_2, H_3, H_8\}$. Here again, the multipath routing strategies H_4 and H_5 are dominated by H_3 because node B acts as a bottleneck; the additional energy expenditure for multipath does not yield an increase in robustness in these cases. It is also remarkable that even in this higher path loss situation the multipath routing strategies H_6 and H_7 , which result in 2 and 3 node-disjoint paths respectively, result in higher energy consumption than the high transmit power direct transmission in scheme H_8 .

Although this is a simple analytical example with a small number of nodes and arbitrary placements, it provides an insight into why multipath routing is not always the best solution when the primary concerns are energy efficiency and robustness to intermediate node failures. The Pareto optimal sets we examined in the two cases $\alpha = 2$ and $\alpha = 3$ contain the three possible single path routing schemes.

We now turn to simulation results involving greater number of nodes with random placement of nodes.

III. EXPERIMENTAL SETUP

For the experiments, 50 nodes are placed in a square area with unit sides. We consider the flow of information from a single source to a single destination. The source is placed at $(0, 0)$, and the destination sink node is placed at $(1, 1)$. The simulation is repeated 100 times with random placements for the remaining 48 intermediate nodes. These nodes are placed at random in the square, independently with a 2D uniform distribution. For each simulation the transmission radius R within which each pair of nodes can communicate is increased in increments of 0.05 from 0.05 to 1.5.

To test out some different single and multipath routing strategies we chose to simulate the family of forward- k routing algorithms that work as follows. The sink first floods a query to all nodes in the network, and the source node responds by routing information in the reverse direction along the paths followed by the initial query. In the forward- k protocol, the source sends its information to the first k neighboring nodes that had sent the sink-initiated query to it. Each intermediate node also forwards this data to its “best” k neighbors. The forward-1 protocol is essentially a single shortest path routing mechanism. The forward-2 and forward-3 protocols that we simulate are examples of braided multipath routing protocols. For comparison we also simulate basic flooding initiated by the source. This is a useful comparison point because of the following result:

Proposition: *If the transmission radius R for a static wireless network is fixed, the routing scheme H_{flood} consisting of flooding information outward from the source results in the optimum robustness value. In other words, $\Pi_{H_{flood}} = \max_H \Pi_H$.*

This result holds because the directed graph corresponding to the flooding scheme H_{flood} is maximally dense in that each directed edge and each node of the underlying topology is utilized. This maximizes the robustness to node failures.

IV. SIMULATION RESULTS

As we mentioned before, much of the prior investigation into multipath routing in wireless networks has focused on providing multiple node-disjoint paths for routing between a source and a destination node. It is intuitive that a forward- k strategy results in greater number of node-disjoint paths as k increases. Figure 2 shows how the number of node disjoint paths varies for the various schemes. It is noteworthy that the flooding is particularly effective as far as this metric is concerned.

We now turn to our robustness metric Π_H - the probability that the source is able to send information to the sink in the presence of uniform random node failures. Figures 4 and 5 show how this metric varies with the transmission radius for failure rates of 5% and 20% respectively. We make two observations from these figures. The first is that for a given transmission radius, the single path routing mechanism does indeed provide much lower robustness than the multipath routing schemes. The second is that for low failure rates, the three multipath routing mechanisms all provide nearly the same level of robustness. In essence the additional redundancy provided by having more than 2 node-disjoint paths results in negligible gains in robustness for low levels of node failures. At the failure rate of 20% there is slightly greater differentiation between the different multipath routing schemes but one can again see the law of diminishing returns at play - flooding provides only negligibly greater robustness than the forward-3 routing protocol.

Thus far we have ignored one critical aspect: the energy expenditure. While the multipath routing schemes provide greater robustness for a fixed value of the transmission radius, they do,




Fig. 2. Number of Node Disjoint Source Sink Routes with respect to the Transmission Radius




Fig. 5. Probability that a route exists with respect to Transmission Radius (20 % failure rate)




Fig. 3. Number of Nodes Transmitting with respect to Transmission Radius




Fig. 6. Probability that a route exists with respect to normalized energy cost (5 % failure rate)




Fig. 4. Probability that a route exists with respect to Transmission Radius (5 % failure rate)




Fig. 7. Probability that a route exists with respect to normalized energy cost (20 % failure rate)

What about
transmission load?

of course, do so at the cost of a greater number of transmissions. This can be seen in figure 3. Flooding requires an order of magnitude higher number of transmissions than even forward-3, showing it is clearly not an energy-efficient mechanism for providing robustness to node failures. This is still far from a clear picture of the energy-robustness tradeoffs. We have two parameters that we can tune to increase the energy and robustness metrics: one is the value of k , which in effect changes the routing structure without affecting the underlying topology. By increasing k , we apply energy in the form of greater number of transmissions in order to realize robustness gains through multiple paths. The second tunable parameter is the transmission radius: even if we stick to single path routing, increasing this parameter increases the robustness to node failures because it decreases the number of hops, leaving fewer possible failure modes.

Hence we plot the robustness metric with respect to the energy metric $E_H = m_H R^\alpha$ which incorporates both the transmission radius R as well as the number of transmission m_H . This is shown as scatter plots in figures 6 and 7 for failure rates of 5 and 20 % respectively, for $\alpha = 2$.

Now we have a dramatically different view. The Pareto optimal points are towards the top left hand corner of the scatter plot. Towards the left hand side of the plot in figure 6, where the energy costs are kept low, it is clear that the single path routing mechanism forward-1 provides the best robustness to node failure. While there is a region where the multipath scheme forward-2 dominates, the remaining schemes are all dominated. It is clear from the plot that from both energy and robustness perspectives it is better to transmit directly from the source to sink than to use either forward-3 or flooding, or even forward-2 with a higher transmission range setting. Figure 7 shows the same behavior for higher failure rates as well. If there are severe energy constraints, this figure suggests that it is better to allocate the energy to increasing transmission range, than to transmit along multiple paths. This validates the insight gleaned from the simple example we explored in section II.

V. CONCLUSIONS

Wireless sensor networks with large numbers of inexpensive individual devices are particularly prone to node failures. In several prior studies multipath routing schemes have been proposed in order to provide tolerance to such failures. We studied the issue of robustness to node failures in the particular context of energy-starved sensor networks, and showed that the robustness obtained from multipath routing can sometimes come at too high a cost.

Multipath routing is but one mechanism for trading off energy in order to increase robustness. An alternative to routing through many paths is the use of higher transmit powers with fewer paths, even a single path. We showed through the simple analytical example and a larger set of simulations that, depending on the constraints, this may be a more energy-efficient

mechanism for robustness to node failures. In any case, wireless sensor network designers would be well advised to consider using higher transmission powers in order to boost robustness, in addition to multipath routing.

REFERENCES

- [1] J. Warrior, "Smart Sensor Networks of the Future," *Sensors Magazine*, March 1997.
- [2] G.J. Pottie, W.J. Kaiser, "Wireless Integrated Network Sensors," *Communications of the ACM*, vol. 43, no. 5, pp. 551-8, May 2000.
- [3] D. Estrin *et al.* *Embedded, Everywhere: A Research Agenda for Networked Systems of Embedded Computers*, National Research Council Report, 2001.
- [4] A. Cerpa *et al.*, "Habitat monitoring: Application driver for wireless communications technology," *2001 ACM SIGCOMM Workshop on Data Communications in Latin America and the Caribbean*, Costa Rica, April 2001.
- [5] C. Intanagonwiwat, R. Govindan and D. Estrin, "Directed Diffusion: A Scalable and Robust Communication Paradigm for Sensor Networks," *ACM/IEEE International Conference on Mobile Computing and Networks (MobiCom 2000)*, August 2000, Boston, Massachusetts
- [6] X. Lin and I. Stojmenovic, "Location based localized alternate, disjoint and multi-path routing algorithms for wireless networks," *Journal of Parallel and Distributed Computing*, to appear.
- [7] M.R. Pearlman, Z.J. Haas, P. Scholander, and S.S. Tabrizi, "On the Impact of Alternate Path Routing for Load Balancing in Mobile Ad Hoc Networks," *ACM MobiHOC'2000*, Boston, MA, August 11, 2000
- [8] A. Nasipuri, R. Castaneda and S. R. Das, "Performance of Multipath Routing for On-Demand Protocols in Ad Hoc Networks," *ACM/Kluwer Mobile Networks and Applications (MONET) Journal*, Vol. 6, No. 4, 2001, pp. 339-349.
- [9] D. Ganesan, R. Govindan, S. Shenker and D. Estrin, "Highly Resilient, Energy Efficient Multipath Routing in Wireless Sensor Networks," *Mobile Computing and Communications Review (MC2R)*, Vol 1., No. 2. 2002.
- [10] P. Papadimitratos, Z. Haas, and E. G. Sirer, "Path-Set Selection in Mobile Ad Hoc Networks," to appear in *Mobihoc 2002*.
- [11] R.J. Gibbens, F.P. Kelly, and P.B. Key, "Dynamic Alternative Routing," *Routing in Communications Networks*, ed. M. Steenstrup, Prentice Hall, 1995.
- [12] S.-W. Lee and C.-S. Wu, "A k-best paths algorithm for highly reliable communication networks," *IEICE Trans. Comm.*, vol. E82-B, No. , p. 586-590, April 1999.
- [13] R. Ogier and N. Shacham, "A distributed algorithm for finding shortest pairs of disjoint paths," *IEEE INFOCOM'89*.
- [14] D. Sidhu, R. Nair, and S. Abdallah, "Finding Disjoint Paths in Networks," *ACM SIGCOMM 91*.
- [15] M. R. Garey and D. S. Johnson, *Computers and Intractability: a guide to the theory of NP-completeness*, W. H. Freeman & Company, 1979.

fdeus@mail.sis.pitt.edu

