

Highly-Resilient, Energy-Efficient Multipath Routing in Wireless Sensor Networks

Deepak Ganesan
UCLA
deepak@cs.ucla.edu

Ramesh Govindan
USC-ISI
govindan@isi.edu

Scott Shenker
ACIRI
shenker@aciri.org

Deborah Estrin
UCLA
destrin@cs.ucla.edu

ABSTRACT

Previously proposed sensor network data dissemination schemes require periodic low-rate flooding of data in order to allow recovery from failure. We consider constructing two kinds of multipaths to enable energy efficient recovery from failure of the shortest path between source and sink. Disjoint multipath has been studied in the literature. We propose a novel braided multipath scheme, which results in several partially disjoint multipath schemes. We find that braided multipaths are a viable alternative for energy-efficient recovery from isolated and patterned failures.

1. INTRODUCTION

Sensor networks [2] are envisaged as large-scale networks of small networked sensor nodes such as the Rene [4]. Such a node could have one or more sensors and would be densely deployed near the phenomena to be sensed, in a highly redundant manner to maximize lifetime, and deal with dynamics and failures.

Three criteria drive the design of large-scale sensor networks: *scalability* (these networks might involve thousands of nodes), *energy-efficiency* (in particular, wireless communication can incur significantly higher energy cost than computation [7]), and *robustness* (to environmental effects and node and link failures).

These networks may require novel routing techniques for scalable and robust data dissemination, such as Directed diffusion [5]. Of particular interest is the notion of path *reinforcement*; that a node in the network may make a local decision (based possibly on perceived traffic characteristics) to draw data from one or more neighbors in preference to other neighbors. We say that such path setup techniques use *localized* algorithms.

In this paper, we propose using multipath routing to increase resilience to node failure. We explore localized algorithms for two different approaches to constructing multipaths between two nodes. One is the classical node-disjoint multipath adopted by prior work, where the alternate paths do not

intersect the original path (or each other). The other approach abandons the requirement for disjoint paths and instead builds many *braided* paths. With braided paths, there are typically no completely disjoint paths but rather many partially disjoint alternate paths.

We use two important metrics in judging the performance of these competing approaches, *resilience* and *maintenance overhead*. There is an inherent tradeoff between these two quantities. Becoming more resilient typically consumes more energy. In this paper we investigate the tradeoffs that result from the two proposed routing algorithms.¹

The literature on multipath routing is vast and we do not attempt to be comprehensive in this summary of related work. To our knowledge, however, ours is the first attempt to evaluate the energy/resilience tradeoff for multipath routing in wireless sensors. Some of our design choices have been influenced by Dispersity Routing [1] and work on multipath in ad-hoc networks [6]

2. DISJOINT AND BRAIDED PATHS

Classical multipath routing has been explored for two reasons: load balancing and robustness. While load-balancing is essential to conserve energy in sensor networks, this is not the focus of our paper. Instead, we use multipath routing to rapidly find alternate paths between source and sink. Our rationale for this use of multipath is as follows. We assume that, from the application's perspective, a desirable goal is to deliver data along this primary (best available) path. However, to scalably (*i.e.* without flooding for rediscovery) recover from failure of this primary path, we construct and *maintain* a small number of alternative paths. Maintaining alternate paths, however, incur the overhead of sending low-rate data through alternate paths as keep-alives and does not preclude the pathological case of failure on all multipaths.

We consider two designs for multipath routing:

disjoint (Section 2.1) and *braided* (Section 2.2). The energy-resilience tradeoffs of these schemes are then explored via simulation (Section 4).

2.1 Disjoint Multipaths

The first multipath mechanism we consider constructs a small number of alternate paths that are *node-disjoint* with the primary path, and with each other. These alternate paths are thus

¹This paper provides a flavor of the multipath techniques suggested. Due to space constraints, some of the more detailed simulations and analysis has been omitted. Please refer to [3] for more details

Permission to make digital or hard copies of part or all of this work or personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

MobiHOC 2001, Long Beach, CA, USA
© ACM 2001 1-58113-390-1/01/10...\$5.00

unaffected by failures on the primary path, but can potentially be less desirable (e.g., have longer latency) than the primary path.

A constructive definition for an idealized k node-disjoint multipath (assuming global knowledge) is:

For the i_{th} node-disjoint path, choose the best path that is node-disjoint with the currently constructed multipath. The resulting is the *idealized k -disjoint multipath*.

Here's one possible localized algorithm for disjoint path construction. Assume for the moment that some low-rate samples have initially been flooded throughout the network. The sink then has some empirical information about which of its neighbors can provide it with the highest quality data (lowest loss or lowest delay). To this most preferred neighbor, it sends out a *primary-path* reinforcement. As with the basic directed diffusion scheme, that neighbor then locally determines its most preferred neighbor in the direction of the source, and so on.

Shortly, thereafter, the sink sends an *alternate path* reinforcement to its next most preferred neighbor. By constraining each node to accept only one reinforcement, the alternate paths setup are guaranteed to be mutually disjoint and disjoint with the primary path. A node that receives more than one reinforcement, *negatively reinforces* all reinforcements but the first. This mechanism can be trivially extended to construct k disjoint multipaths.

We call these *localized* disjoint multipaths. This search procedure may discover longer alternate paths than the idealized version, being restricted to local knowledge. This difference accounts for some performance differences between the two kinds of disjoint multipaths.

2.2 Braided Multipaths

While disjoint paths have some attractive resilience properties, they can be energy inefficient since node-disjoint paths could be potentially longer than the primary path. Our *braided multipath* relaxes the requirement for node disjointness. Alternate paths in a braid are partially disjoint from the primary path, not completely node-disjoint.

A constructive definition for our *braided* multipath is (Figure 1): For each node on the primary path, find the best path from source to sink that does *not* contain that node. This alternate best path need not necessarily be completely node-disjoint with the primary path. We call the resulting set of paths (including the primary path) the *idealized* braided multipath. As its name implies, the links constituting a braid either lie on the primary path, or can be expected to be geographically close to the primary path. In this sense, the alternate paths forming a braid would expend energy comparable to the primary path.

One localized technique for constructing braids is described below. As in Section 2.1, the sink sends out a primary path reinforcement to its most preferred neighbor. In addition, the sink sends an alternate path reinforcement to its next preferred neighbor. In addition, recursively each other node on the primary path *originates an alternate path reinforcement* to its next most preferred neighbor. By doing this, each node thus tries to route around its immediate neighbor on the primary path towards the source. When a node, not on the primary path receives an alternate path reinforcement, it propagates it towards its most preferred neighbor. When a node already on the primary path receives an alternate path reinforcement, it

does not propagate the received alternate path reinforcement any further.

Figure 2 illustrates a localized braid obtained by using the above mechanism. In this figure, n_{k+1} sends an alternate reinforcement to route around n_k that passes through a_k and a_{k-1} before rejoining the primary path at n_{k-2} . In practice, though, our local rules cannot always ensure this perfect detour around n_k . These effects vary with node density and other factors, and arise due to the absence of global knowledge.

3. EVALUATION METHODOLOGY

In this section, we precisely define our two metrics for multipath performance: maintenance overhead and resilience. We also describe the failure models for which we evaluated the resilience of our multipath mechanisms. Finally, we discuss our experimental methodology and list the parameters that affect the multipath schemes.

3.1 Maintenance Overhead

The *maintenance overhead* of a scheme is a measure of the energy required to maintain these alternate paths using periodic keep-alives. Assume that the source disseminates r events in some time interval T over the primary path. Then, we assume that ϵr events are sent on the alternate paths of the disjoint or the braided multipath, with each alternate path receiving equal proportions of this keep-alive traffic. Then, the energy required to maintain the alternate paths is proportional to the average length (in number of hops) of the alternate paths. To meaningfully calibrate the maintenance overhead, we normalize it with respect to the length of the shortest path. Thus, our maintenance overhead metric is:

$$(L_a - L_p)/L_p \quad (1)$$

where L_a is the average length of an alternate path, and L_p is the length of the primary path.

3.2 Failures

We study the resilience of our multipath routing schemes to two widely different failure models: independent node failures, and geographically correlated failures.

Isolated Failures: Our first failure model captures independent node failures and represent the effect of local environmental effects. More precisely, each node in the multipath has a probability of failure p_i during some small interval T . Then, for each of our multipath schemes, we define *resilience to isolated failure* to mean the probability of at least one alternate path being available within the interval T , given that *at least one node on the primary path has failed*. This latter constraint captures our use of multipath routing for recovery from shortest path failure.

Patterned Failures: Our second failure model captures geographically correlated failures. Specifically, a patterned failure results in the failure of all nodes a circle of radius R_p . The choice of a circle is somewhat arbitrary, but attempts to model the idealized wave propagation of most physical phenomena. The rough justification for this model is that sustained activity or environmental effects (such as rain fades) within a geographic region can cause such correlated failure, either due to loss of connectivity or due to energy dissipation.

We assume location of the centers of these circles is randomly distributed within the sensor field. Furthermore, lack-




Figure 1: Idealized Braid




Figure 2: Localized Braid




Figure 3: Perfect Braid

ing any other realistic model, we assume that the number of patterned failures within a given small time interval T is Poisson distributed, with some parameter λ_p .

Then, for each multipath scheme, its *resilience to patterned failure* is defined as the probability that, within a small interval T : at least one alternate path is available between source and sink, given that at least one node on the primary path falls within the circle defining a patterned failure.

3.3 Details of Methodology

In Section 4, we discuss our evaluation of disjoint and braided multipaths via simulation over the 802.11-like MAC in $ns-2$. For generating reinforcements, the most-preferred neighbor was the one from whom a given event was heard first. All our experiments were conducted by uniformly distributing a number of sensor nodes with transmission radius 40meters on a finite plane of dimension 400 meters square.

To compute a multipath’s resilience to isolated failures, we repeated the following set of steps a large number of times:

- Fail each node on the multipath with probability p_i .
- If a node on the primary path has failed, then, the assign a value of 1 to this set if at least one alternate path is available, 0 otherwise

The resilience of the multipath to isolated failures is the average value assigned to sets in which at least one node in the primary path fails. The number of runs of the experiments, and the number of sets in each run were adjusted to obtain acceptable 95% confidence intervals.

To compute a multipath’s resilience to patterned failures, we repeated the following set of steps a large number of times:

- Pick an integer n from a Poisson distribution with parameter λ_p .
- Randomly place n points on the plane.
- Fail all nodes within a radius R_p of the plane.
- If a node on the primary path has failed, then, the assign a value of 1 to this set if at least one alternate path is available, 0 otherwise.

The resilience of the multipath to patterned failures is the average value assigned to sets in which at least one node in the primary path fails. The number of runs of the experiments, and the number of sets in each run were adjusted to obtain acceptable 95% confidence intervals.

3.4 Qualitative Comparison

Before discussing our simulation results, we try to present some intuition for the energy/resilience tradeoffs of the two multipath schemes we have discussed so far. We use the corresponding idealized mechanisms to guide our intuition, since

their behavior easier to reason about than their localized counterparts.

The energy cost of alternate disjoint paths depends on the network density. At low network densities, alternate disjoint paths are significantly longer than, and have higher cost than, the primary path. At higher densities, the likelihood of finding node-disjoint alternate paths of shorter length increases, thereby reducing the energy cost of maintaining them. In contrast, the energy cost of an alternate path in the braid is comparable to that of the primary path, more or less independent of density. Thus the difference in *maintenance overhead* between disjoint and braided multipath is high at lower densities.

Disjoint paths give us independence, *i.e.*, any number of nodes can fail on the primary path without impacting the alternate path. However, the failure of a single node on each alternate path results in the failure of the multipath. By contrast, in braided multipaths, the various alternate paths are not independent, and a combination of failures on the primary path could sever all alternate paths. However, the number of *distinct* alternate paths through a braid is significantly higher than the number of nodes in its primary path. This contributes to the greater resilience of the braid.

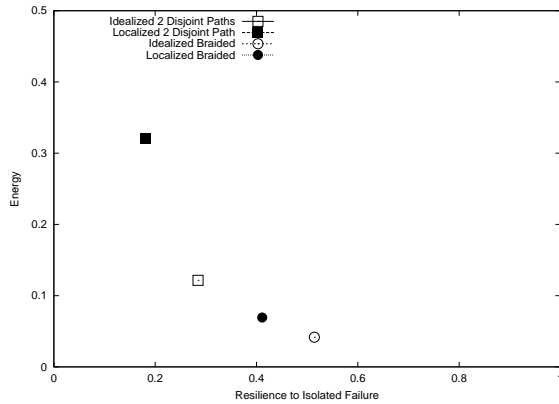
Patterned failures also affect disjoint and braided paths differently. A failure pattern that affects the primary path would be likely to affect alternate paths that are geographically near primary path, and affect less paths that are more distant. Since braiding encourages geographically closer alternate paths, disjoint multipaths are likely to be more resilient to pattern failures than braided multipaths.

4. SIMULATION RESULTS


In performing these simulation experiments, our goal was to understand the energy/resilience tradeoff between our various multipath schemes, their dependence on network density, and validation of the idealized models. One simple instance of the energy/resilience tradeoff is illustrated in Figure 4. We see that for isolated failures, 2-disjoint idealized multipaths are significantly less resilient, and have higher maintenance overhead than idealized braided multipaths. For patterned failures, the idealized schemes have comparable resilience, but 2-disjoint has higher maintenance overhead. Similar distinctions exist for the localized mechanisms.

Clearly, Figure 4 does not represent the whole picture. Our simulations [3] carefully study the impact on each metric of varying different parameters. We summarize the salient observations in this section.

Maintenance Overhead: Overall, braided idealized multipaths require lower maintenance overhead than 2-disjoint idealized multipaths, the difference being significant at lower densities. The localized braided heuristic does not exhibit the same properties as the idealized version at low densities, but tracks it closely at higher densities. Finally, the maintenance overhead of localized 2-disjoint is nearly an order of magni-



(a) Isolated: 400 nodes, 6 hop source-sink separation, $p_i = 0.2$



(b) Patterned: 400 nodes, 6 hop source-sink separation, $\lambda_p = 3, R_p = 20$

Figure 4: Illustrating the energy vs resilience tradeoff

tude higher than localized braid at high densities. In other words, we believe these results show that it might be easier to construct low-overhead braids than to construct low-overhead disjoint paths using localized algorithms.

Resilience to Isolated Failures: In general, the idealized braid is more resilient than the idealized disjoint multipath, the difference being significant at higher densities. Localized algorithms are slightly less resilient than their idealized counterparts. The reasons for these is that both the localized braid and the localized disjoint multipath can discover longer paths than their idealized counterparts.

Resilience to Patterned Failures: The resilience to patterned failure of the idealized braid compares well with the idealized 2-disjoint and 3-disjoint paths. The localized braid varies differently to density than the idealized scheme (explained in [3]). With increasing frequency of failure, or radius of failure, the resilience decreases, although the impact of radius is more dramatic. Increasing the level of disjointness (of disjoint paths) only gives us modest resilience gain, and incurs large cost difference.

5. CONCLUSIONS

We demonstrate that multipath routing can be used for energy-efficient recovery from failure in wireless sensor networks. We explore and evaluate a novel braided design which shows considerable promise.

For a disjoint multipath configuration whose patterned failure resilience is comparable to that of braided multipaths, the braided multipaths have about 50% higher resilience to isolated failures and a third of the overhead for alternate path maintenance.

We believe that it is harder to design localized energy-efficient mechanisms for constructing disjoint alternate paths, because the localized algorithms lack the information to find low latency disjoint paths.

Finally, increasing the number of disjoint paths does increase the resilience of disjoint multipaths but with a propor-

tionately higher energy cost. It is not the case that a small energy expenditure dramatically improves the resilience of disjoint paths.

6. REFERENCES

- [1] Anindo Banerjee. Simulation Study of the Capacity Effects of Dispersity Routing for Fault Tolerant Real-Time Channels. In *ACM Computer Communications Review*, volume 26, pages 194–205. ACM Press, October 1996.
- [2] Deborah Estrin, Ramesh Govindan, John Heidemann, and Satish Kumar. Scalable coordination in sensor networks. In *Proceedings of the Fifth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom 99)*, pages 263–270, N.Y., aug " 15–20" 1999. ACM Press.
- [3] Deepak Ganesan, Ramesh Govindan, Scott Shenker, and Deborah Estrin. Highly-resilient, energy-efficient multipath routing in wireless sensor networks. Technical report, University of California, Los Angeles, 2001.
- [4] Jason Hill, Robert Szewczyk, Alec Woo, Seth Hollar, David Culler, and Kristofer Pister. System architecture directions for network sensors. In *ASPLOS*, 2000.
- [5] Chalermek Intanagonwiwat, Ramesh Govindan, and Deborah Estrin. Directed diffusion: A scalable and robust communication paradigm for sensor networks. In *Proc. ACM Mobicom*, Boston, MA, 2000.
- [6] Asis Nasipuri and Samir R. Das. On-Demand Multipath Routing for Mobile Ad Hoc Networks. In *Proceedings of the 8th Int. Conf. on Computer Communications and Networks (IC3N)*, Boston, MA, 1999.
- [7] G. Pottie and W. Kaiser. Wireless Sensor Networks. *Communications of the ACM*, 43(5):51–58, May 2000.