BIPAR: BImodal Power-Aware Routing Protocol For Wireless Sensor Networks

HANY MORCOS IBRAHIM MATTA AZER BESTAVROS
{hmorcos, matta, best}@cs.bu.edu
Computer Science Department, Boston University

Abstract

Wireless Sensor Networks (WSN) have the potential to change the way we perform many tasks today. Examples include military applications, agriculture applications and medical applications. Routing protocols in WSN have to operate in challenged environments. In these environments, packet losses and node failures are common. One other challenge is the limited power supply of sensors since they are battery-powered, which makes power saving a crucial feature of any WSN protocol in order to increase the lifetime of the whole network.
In this paper, we present ,BIPAR, a new routing protocol that counteracts the effects of the environment on sensors and, at the same time, tries to minimize its power consumption. The design of BIPAR is very intuitive. It is a semi-reliable protocol that tries to use the least amount of power to deliver data packets, i.e., it routes packets on the least-power path to the sink. If successful, this behavior saves as much battery power as possible. On the other hand, if the first transmission on the least-power path is not successful, BIPAR switches to the max-power path to the sink. This behavior consumes more energy than the first transmission, but maximizes the probability of successful communication.

We simulated BIPAR in ns2[4] and evaluated it under different modes of node failure models. We compared it against GRAB[1], min-power routing scheme, and max-power routing scheme. Our simulations show that BIPAR delivers at least 30% more reports than GRAB, when node failures are spread all over the routing field. BIPAR delivers as much as 50% more reports than min-power routing protocol under the same conditions.

1 Introduction

Recent technological advances in electronics made viable the deployment of large number of small sized sensors composing wireless sensor networks (WSN). WSN have the potential to change the way we perform many tasks today. Examples of potential applications include environmental applications and habitat monitoring, military applications and medical applications[3].

Due to the nature of their potential hosting environments and their technological structure, WSN are considered challenged networks. Routing protocols in WSN have to operate effectively in such environments. For example routing protocols should expect and successfully react to node failures and/or packet losses. As routing conditions in such environments change over time, adaptation to the current situation should be inherent in any routing protocol for WSN. Moreover, as sensors are battery powered, power conservation is a very crucial feature of any protocol to be deployed in WSN, as the lifetime of any sensor network is inversely proportional to the amount of power consumed during its operation.

Many routing protocols have been recently proposed for WSN (e.g., [1][2]). The proposed protocols succeed in realizing and reacting to the environmental effects on the performance of WSN, however they either lack in the adaptation requirement, or they react globally to local changes in routing conditions. GRAB [1] is an example of the second type. Global reactions have the advantage of global view of the network but they lack the details of the local view. When a central processing point reacts globally to a local change, it is not easy for that point to precisely affect the local area. Hence this approach may incur higher overhead on parts of the network that were not affected by the local changes in the first place.

In this paper we present the design of a new routing protocol for WSN, BIPAR. BIPAR has two modes of operation (hence the name bimodal). The two modes are min-power and max-power routing. Min-power routing is a routing scheme that delivers packets over the minimum-power path from the source to the destination. In doing so, min-power routing favors paths that have physically shorter hops to those of longer hops (as explained later in Section 2.) Min-power routing is more adequate to environments with low packet loss ratio and node failure rate. The other mode is max-power routing. Contrary to min-power, max-power routing uses more power to route packets and it favors paths of physically longer hops to those of shorter hops. Using max-power routing maximizes the probability of success of communication in the presence of node failure and packet losses which make it suitable for environments with high node failure rate and packet loss ratio. Since max-power routing favors physically longer hops, it minimizes the number of hops traversed to reach the final destination.

Trying to save as much power as possible, any BIPAR node first routes packets in the min-power mode. In case communication is successful, BIPAR saves as much power as possible. Otherwise, nodes would switch to max-power mode to route packets. The second transmission would consume more power than the first one but it would have much higher success probability as well. This model suits massive node failures where a number of nearly located nodes fail simultaneously. Switching to max-power routing in this case can be the only solution to deliver packets successfully.

The main contribution of BIPAR is the ability to switch between the two modes of operation to adapt to the current routing conditions in the environment.

We have simulated BIPAR in ns-2[4] and compared it to GRAB[1], min-power routing, and to max-power routing. As expected, BIPAR gets the best of the two worlds: it achieves high throughput with the minimum amount of power. We evaluated the four protocols in different scenarios of node failure and packet loss. BIPAR maintained high throughput with minimum amount of power. Achieving the lowest power per successful report among the four protocols, BIPAR delivers as many as 30% more successful reports than GRAB.

The rest of the paper is organized as the follows; Section 2 describes the characteristics of the routing protocols that we use as a basis for comparison. We, then, give the design details of BIPAR in Section 3. in Section 4, we describe the evaluation methodology and the performance results of BIPAR compared to GRAB, min-power and max-power routing. Then Section 5 surveys related work. Section 6 gives future venues to explore. Finally, Section 7 concludes the paper.

2 Evaluated Protocols

In this section we describe the different protocols that we evaluate against BIPAR. Specifically, we describe GRAB[1], min-power routing and max-power routing. Recall that min-power and max-power are the two modes of operation of BIPAR. For all protocols we assume that sensors are distributed in a routing field. Sensor nodes can sense some stimulus of interest, and communicate with each other using radio communication. Sensor nodes have limited wireless range in which they can communicate. There is a special point for collecting sensed data from the sensors in the field. We call this point *the sink*. Some stimulus of interest occurs at some point of the field. Upon receiving a signal from that stimulus, the closest sensor node assumes responsibility to forward its sensed data back to the sink, we call this node *the source*. If the sink is within transmission range of that node, the data is sent directly to the sink. Otherwise, data has to go through multiple hops to the sink.

2.1 Min-power and Max-power schemes

As we mentioned above, min-power and max-power routing schemes have different characteristics as routing protocols. Max-power routing minimizes the number of hops to the destination and prefers physically longer hops to shorter ones. While min-power routing minimizes the total consumed power and prefers shorter hops to longer ones. The main difference between the two routing schemes is in the cost function we assign to different links in the routing algorithm. In max-power routing, since it minimizes the number of hops, the cost for all links is equal (e.g., 1). While in min-power routing the cost function for a link between any two nodes is the least amount of power needed to transmit a packet between the two nodes. Here, we note that, while optimizing the number of hops, max-power routing also optimizes the physical distance over which packets are transmitted [1].
The relation between the amount of power (P) needed to transmit packets between any two nodes X and Y, and D, the distance between X and Y, is given by the relation

$$P = D^{\alpha} \tag{1}$$

where α is a constant that depends on the signal propagation model and the value of D. In most cases α varies between 2 and 4. We have three observations on Equation 1

- As we can notice from Equation 1, power, as a cost function, does not obey the triangular inequality. A direct consequence of that is, when using power as a cost function, the total cost of a path that consists of hops that are physically short (but more in number) would be less than that of a path that consists of hops that are physically longer (but less in number);
- Max-power routing can be related to the model in equation 1. By setting α to 1 we get the cost function for a link under the max-power scheme;
- In any cost function that follows equation 1, as the value of α gets higher, the resulting routing protocol would tend to choose routes that consist of hops that are physically short. This property is well-suited for environments with low packet loss rate and node failure ratio.

[1] This is true since physical distance obeys the triangular inequality.

To illustrate the last point consider the probability of losing a packet over an N-hop path ($P_x^{fail}(N)$), where x is the node failure rate, is given by

$$P_x^{fail}(N) = 1 - (1 - x)^N \tag{2}$$

It is obvious that the more hops packets traverse, the higher the probability of loss. We conclude that, if a cost function follows Equation 1, the higher the value of α the less suitable the resulting routing protocol would be for routing in environments with high packet loss ratio or node failure rate.

To illustrate the difference between min-power routing and max-power routing we give this example. Consider the scenario given in Figure 1. Node X needs to forward a data packet. Under min-power routing, X sends the packet to node Y, since Y is the closest neighbor to X. As obvious from the graph, this transmission consumes the least amount of power but reaches only node Y. Min-power routing reaches the physically closest neighbor of a node. If this neighbor has failed for any reason, or if there was some packet loss probability around this neighbor, the communication fails.

While under max-power routing, X sends the packet to node Z, the furthest node in X neighbors' list. Max-power routing consumes more power than min-power, but maximizes the probability of delivering the packet to a node closer to the sink than X. To see that, note that the second transmission reaches nodes: Y, A, B, C, D, and Z. If min-power routing failed due to packet loss at Y, now Y has another chance to get the packet. Even if Z has failed (for some reason) not only Y, but also nodes, A, B, C, and D have a chance to receive the packet and to send it towards the sink, which maximizes the probability of successfully sending the packet to the sink. To consider X successful in forwarding packets, it is enough that one node (at least) closer to the sink than X receives the packet.

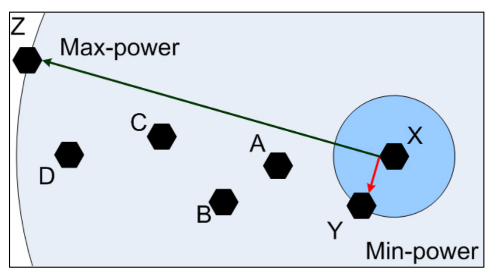

Figure 1: The effect of the first and the second transmission from X on the consumed energy and the success probability

2.2 GRAB

GRAB is a routing protocol for WSN. GRAB has two phases of operation: cost establishment phase and data forwarding phase. In the cost establishment phase, GRAB assigns a cost to each forwarding node that is proportional to the minimum amount of power needed to forward packets from this node to the sink. The cost field constitutes a monotonically increasing value as nodes get further from the sink. In the data forwarding phase, to forward packets from the source to the sink, only nodes with cost less than that of the sender of any packet can forward this packet. This restriction ensures that GRAB is loop-free. Moreover, the source assigns each packet a fixed budget, which is the total amount of energy that may be used to forward this packet to the sink. This budget is not to be exceeded otherwise the packet is dropped. When receiving a packet, any node X checks if it has enough credit. Credit is calculated as a function that involves the cost of this node and that of the source as well as the amount of power consumed so far to forward the packet. If a node has enough credit, it sends the packet to 3 of its closest neighbors. Otherwise, it sends the packet to only the next neighbor on the least-power path to the sink. This credit distribution function is shown in [1] to allot more credit to nodes closer to the source, which is important to establish a forwarding mesh as fast as possible to overcome node failures or packet losses.

GRAB depends on the redundancy in the forwarding mesh to overcome unreliability in the routing environment. GRAB also calculates the average throughput in a fixed-size window of reports. When this average falls below a certain threshold, the sink reassigns the cost to the forwarding nodes to restore the throughput. The underlying assumption is that throughput decreases as a result of node failures which may create holes in the old routing tables. Reestablishing the cost field in the remaining nodes would restore throughput. Note that the actual process of assigning cost values to forwarding nodes is done as follows: the sink broadcasts an advertisement (ADV) packet to all its neighbors. Every node that gets this ADV packet uses it to calculate its own cost and then rebroadcasts the packet. Having *every* node in the forwarding field broadcast a packet using maximum transmission power is a high overhead. This behavior is what we mentioned above about the global reaction of GRAB to local changes in routing conditions by having the *whole* network broadcast an ADV packet. Some parts of the network that are not affected by these local changes by they still have to broadcast this ADV packet which is a significant overhead in terms of power consumption.

3 Protocol Design

In this section we present the design details of BIPAR. The operation of BIPAR has two phases: cost establishment phase and data forwarding phase.

Cost Establishment Phase The point of this phase is to set the routing status in the forwarding sensor nodes. In this phase, the sink sends *Advertisement packet* (ADV). The ADV packet serves two purposes:

- It assigns costs to each node. A node's cost is the least amount of power needed to transmit packets from this node to the sink. The ADV packet has a cost field. When the sink first broadcasts it, the ADV packet has a cost of 0. When node X receives an ADV packet from node Y, it sets its own cost as the sum of the cost field in the packet and the amount of power needed to transmit packets from Y to X [1]. Then, X sets the cost of the ADV packet to its own cost and rebroadcasts the packet. Being set in the forwarding node this way, the cost field comprises a monotonically increasing field as we get further from the sink. After forwarding the first ADV packet, X will not forward any other ADV packet of the same or of higher cost than the first one. X will only forward another ADV packet of lower cost or of higher sequence number [2].
- During the cost establishment phase, nodes get a chance to build a list of neighbors toward the sink. This list is the routing table of each node. By relating the cost field in the ADV packet to the sender of the packet, nodes can build the neighbors' list. The invariant of that list is: *nodes in the neighbors' list of any node X have cost strictly less than that of X*. X keeps its neighbors' list sorted based on the physical distance between itself and the respective neighbor.

After this phase every intermediate node (including the source) would have a cost assigned to it. BIPAR shares this phase with GRAB[1].

Data Forwarding Phase In this phase sensor nodes sense the environment and send their measured data back to the sink. When a stimulus of interest occurs, a number of sensors detect it with different signal amplitudes. The sensor that has the highest amplitude of all nodes assumes responsibility to forward the measured data to the sink. This node becomes *the source*. The source assigns a power budget to each data packet it sends. This budget is the total amount of power to be used to forward this packet from the source to the sink. If any packet exceeds its budget it will be dropped. The budget usually takes the form of

$$bugdet = \alpha \times Cost(source) \tag{3}$$

where α is an integer greater than 1. Note that setting α to 1, would mandate that the packet be forwarded along the minimum-power path to the sink. This is equivalent to setting the total budget to the cost of the source. Since the cost of the source is the cost of the ADV packet delivered to it along the minimum-power path, setting α to 1 mandates that the packet take the exact same least-power path back to the sink. As the value of α gets higher, it allows packets to be forwarded on paths that deviate more from the least-power path. Hence higher α means less restrictions and more paths to use to forward the packet. Consequently higher probability of successfully delivering packets and higher consumed energy to forward them.

Along with the budget, the source sends the following fields in the data packets

- Sender's cost: which is the cost of the sending node.
- Consumed power so far: which is the amount of energy consumed so far in forwarding this packet.

Upon receiving any data packet from node Y, node X compares its own cost to the cost of the sender. Node X can only rebroadcast the packet if its cost is less than that of Y, otherwise X drops the packet. This insures that BIPAR is loop-free.

If X decides to rebroadcast the packet, it calculates the power needed to send the packet from Y to itself and update the consumed power so far field of the packet. The latter is

[1] We assume that nodes can estimate the power needed to communicate with any direct neighbor using the signal to noise ratio (SNR). If the forwarding node includes the power level used to transmit the packet, and the receiving node can measure the power level at which the packet was received, then nodes can estimate the physical distance to direct neighbors based on some signal propagation model.

[2] ADV packets have sequence numbers as an indication of the time of sending them.

checked against the budget allowed for this packet. If the packet has exceeded its budget, X drops it.

X then consults its neighbors' list and picks its closest neighbor to forward this packet to. This behavior basically forwards packets on the least-power path to the sink which is equivalent to the min-power routing scheme. X then waits for an acknowledgement (ACK) for a predefined timeout interval. If X gets an ACK for its packet during this timeout interval, then X's job is done concerning this packet. Otherwise, X would consult its neighbors' list again, this time picking its furthest neighbor to forward the same packet to. The second transmission is equivalent to the behavior of max-power routing. It is obvious that the second transmission consumes more power than the first one but it maximizes the probability of successfully sending this packet to a node that is closer to the sink than X.

Finally, we explain the mechanism of ACK's in BIPAR. Upon receiving a packet from node Y, node X calculates the distance between itself and Y(say D1). X, then, looks up its neighbors' list and decides on D2, the distance to send the packet over. If D2 is larger than D1, then Y can overhear X's transmission and consider it as an *implicit* ACK. Otherwise, X sends a small ACK packet to Y, which is an *explicit* ACK.

4 Performance Analysis

We simulated BIPAR in ns-2[4] and compared it to GRAB[1], min-power routing, and max-power routing protocols. We next give the network model, performance metrics, and then the results of the simulation.

4.1 Network Model

The field size is 2000mx2000m, with 400 nodes uniformly distributed in the field. The sink is located at the left hand of the field at (10,1000) and the stimulus at the right hand of the field at (1990,1000). The initial battery of forwarding sensors is 50 Joules while we assumed that the source and the stimulus have infinite supply of energy. The stimulus generates a new report every 5 seconds. The simulation runs for 5000 seconds which was enough to drain the battery of many forwarding nodes. We assumed stationary sensor nodes, with maximum transmission range of 250m (the default value in ns wireless extension.) The power consumption rates are set to 0.66W for transmitting, 0.395W for receiving, and 0.035W for idle (the same values assumed in GRAB[1] and Diffusion[2]).

4.2 Performance Metrics

The performance metrics we used to evaluate the protocols are

- The delivery ratio of the network which is the number of packets received at the sink divided by the number of packets generated at the source.
- Delay of delivering the data packets. The point of this metric is to study the effect of the BIPAR's timeout interval on the total delay of successfully delivered reports.
- The time of delivering the last report. This is a measure of the network's lifetime. As we mentioned above, we allowed our simulations to run long enough to drain the power of forwarding nodes, so the network fails as a result of limited battery power. Since different routing protocols manage power consumption differently, the later the time of sending the last report the longer the lifetime of the network is.
- The amount of energy consumed per successful data report. This measure is an indication of the cost paid to realize the achieved performance.

4.3 Simulation Results

The point of this experiment is to test the reaction of the four routing protocols to node failures. Our node failure model in this experiment is as follows. We assume that nodes stay up and functioning for an exponentially-distributed amount of time with an average of 300 seconds. Then nodes go to an exponentially-distributed temporary blackout period of average T. We varied the value of T from 60 to 210 (increase step of 30 seconds.) The X-axis in the following graphs is ratio between the average downtime (T) and the average uptime (300 seconds) periods. The points to the left of the X-axis have an average downtime of 60 seconds (ratio of 0.2), while those to the right have an average of 300 seconds (ratio of 1).

Figure 2 shows the throughput of the four protocols. BIPAR achieves higher throughput than the other three protocols. This proves that the semi-reliability of BIPAR yields higher throughput. Min-power routing delivers relatively high throughput for low failure rate of nodes, however, its throughput goes down for higher failure rates. The reason for that is, for low failure rate, nodes are up most of the time, hence min-power routing tables (i.e., the neighbors' list) correctly reflect the topology of the network, hence min-power routing achieves high throughput. While for higher failure rates, nodes are down for longer time, hence the routing status of forwarding nodes is stale. Since, min-power routing does not have a notion of reliability, failed nodes cause path failure that is not detected by min-power, hence, the low throughput. On the other hand, max-power routing achieves very low throughput. There are two reasons for that. First, high powered transmissions cause many collisions which decrease the total throughput. Second, using only high-powered transmissions to reach the sink consumes much power and causes packets to exceed their credit very fast. Hence nodes drop packets because of the budget constraint. To verify that claim, we tried the same experiments with much looser constraint on the allowed budget. Max-power routing achieved very high throughput compared to that of BIPAR, but using much more power. Figure 3 shows a ratio between the number of successful data packets under GRAB, max-power routing and min-power routing to the number of successful data packets under BIPAR. BIPAR delivers as much as 30% more packets than GRAB, and up to 50% on average more than min-power routing.

Figure 4 shows the total delay of successfully delivered data packets. As expected the timeout interval of BIPAR affects the total delay of its delivered reports. However, the delay of BIPAR is less than twice that of GRAB. Also note that since min-power routing delivers packets on the least-power path advancing one hop at a time, it has higher delay than GRAB and max-power but still lower than BIPAR. Ensuring that the delay of routing protocols is bounded is important for some applications where delivering data too late is as good as not delivering it at all. Figure 4 shows that the delay of BIBAR is still very small compared to the time between successive reports generated at the source (5 seconds).

Figure 5 shows the average time of delivering the last report in the simulations. Recall that simulations run for 5000 seconds to drain the power of sensors that participate in the forwarding process. The time of delivering the last report is an indication of the network lifetime under the different protocols. It is not simply the time until the first node dies, as

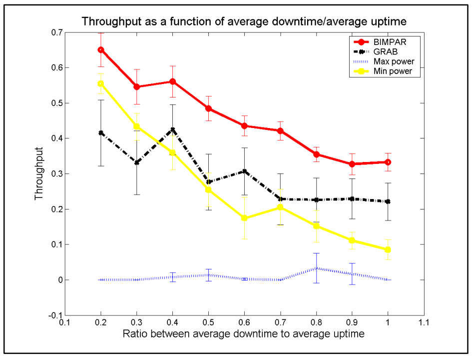

Figure 2: Throughput of the four protocols in the presence of node failures

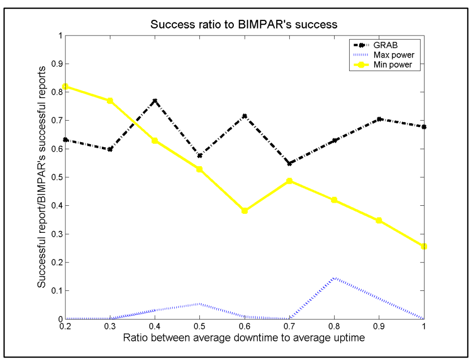

Figure 3: Ratio between the number of successful data packets under GRAB, min-power routing, and max-power routing to the number of successful data packets under BIPAR

the network still functions for sometime after that. This measure should be considered along the delivery ratio of each protocol. For example, if a routing protocol does not deliver any data reports in the beginning of the simulation and starts delivering packets at the end will have "longer" lifetime but with lower delivery ratio. BIPAR has the latest time of delivering the last report. This indicates that BIMPAR has the potential to increase the lifetime of the whole network. The reason for that is BIPAR's measures to maintain high throughput saves the overhead of resending ADV packets through the whole network which saves energy that BIPAR uses in sending useful data packets.

Figure 6 shows the consumed energy per successful report. This metric is the cost paid by each routing protocol to achieve its performance. There is always a compromise between performance and cost. If the cost for high performance is very high as well, then that performance may not be totally justified. This graph is obtained by dividing the total consumed energy by the number of successfully delivered data packets. It is obvious from Figure 2 that the throughput of max-power is small. This is the reason we did not include it in this graph as it has very high cost (dividing consumed power by a very small number of delivered reports.)

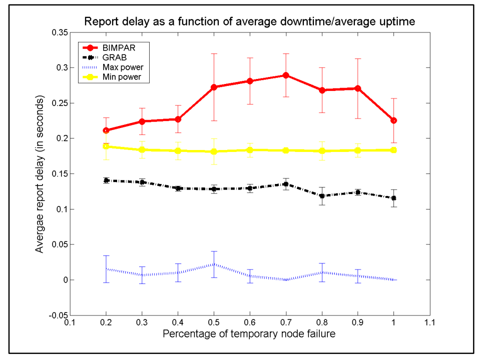

Figure 4: Average delay of delivering reports

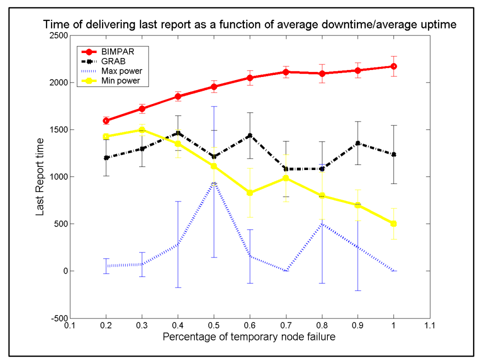

Figure 5: Time of delivering the last report before nodes run out of power

BIPAR has less consumed power per successful data packet than GRAB and min-power routing. The message of Figure 6 is that BIPAR's performance is very well justified by a low cost. The reason min-power routing consumes a lot of energy is not obvious. The high throughput that min-power routing achieves (Figure 2) is helped by refreshing the cost field in the forwarding nodes very often to update the routing status of the nodes. This behavior while enables min-power routing to achieve high throughput, consumes power during the ADV broadcast process. Min-power routing does not consume a lot of power to forward data reports, as it uses the least-power path to send its packets, but it consumes much more power to forward control packets (i.e., ADV packets). That is why the total amount of power consumed by min-power is high.

5 Related Work

There has been a lot of work recently in routing protocols in sensor networks. Diffusion [3] is a sensor networks routing protocol that establishes multiple single-path routes between the source and the sink. Based on some performance metrics (e.g. delay), the sink evaluates the performance of each route and selects one or more of them to be the primary route(s). The sink reinforces the selected routes by assigning

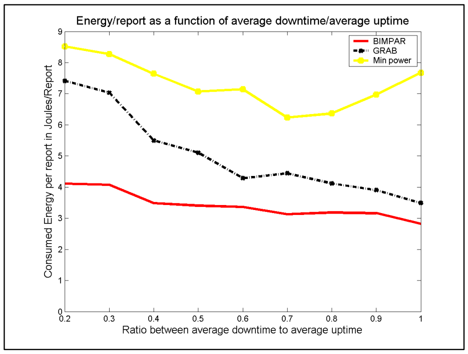

Figure 6: Consumed energy per report under BIPAR and GRAB

them higher reporting rate. Diffusion still refreshes the other routes with low reporting rate to keep them alive and evaluate their quality. Based on path performance, the sink (or any intermediate node on a previously reinforced path) may switch its primary route. This allows diffusion to react locally to route failure or degradation. One inherent problem with this scheme is the speed of switching to and establishing a high-quality, loop-free path. Another issue is its dependance on single paths. It has been shown that the performance of multiple single paths is inferior to that of multiple braided paths of the same number[1][6]. In its first transmission attempt, BIPAR uses single paths reliably. In case a single path fails, BIPAR switches to braided-path forwarding. Thus BIPAR avoids the problems of sticking to single path forwarding.

GRAB is another sensor networks routing protocol. We described GRAB in Section 2. BIPAR shares with GRAB the main cost establishment phase. During the data forwarding phase, BIPAR has some notion of reliability, in the sense of waiting for ACK for data packets and resending the same packet in case the first transmission was not successful. One more difference, in its first attempt, BIPAR uses the minimum-power path to save power and credit for places where it is needed. The results in the previous section show that BIPAR delivers higher throughput than GRAB with less consumed power.

6 Future Work

As we mentioned above, BIPAR has two modes of operation, specifically, min-power and max-power routing. These two modes are the two extremes of communication. The former is sending packets only to the nearest neighbor, while the latter is sending packets to all neighbors. A potential technique to experiment with in BIPAR is to try to visit intermediate points between these two extremes. It is not very clear now how to move between points in that spectrum, but a strong candidate is to apply some controller that decides on the next operating point based on the current routing conditions.

7 Conclusion

We presented the design of BIPAR, a new routing protocol for sensor networks. We argued for the benefits of switching modes of operation in the presence of node failures or packet losses. We simulated BIPAR in ns-2 and compared its performance against GRAB, min-power, and max-power routing schemes. BIPAR achieves higher throughput than the other protocols with less consumed power per successful report.

8 References

[1] F. Ye, G. Zhong, S. Lu, and L. Zhang. GRAdient Broadcast: A Robust Data Delivery Protocol for Large Scale Sensor Networks. To appear in *ACM Wireless Networks (WINET)*, Vol. 11, No. 2, March 2005.
[2] C. Intanagonwiwat, R. Govindan, and D. Estrin. Directed Diffusion: A Scalable and Robust Communication Paradigm for Sensor Networks. In *Proc. of MobiCOM*, August 2000.
[3] Ning Xu. A Survey of Sensor Network Applications. Computer Science Department, University of Southern California.
[4] The Network Simulator - ns-2. http://www.isi.edu/nsnam/ns/.
[5] Sally Floyd. Congestion Control Principles. RFC 2914, September 2000.
[6] 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.