Hierarchical Power-aware Routing in Sensor Networks

Qun Li, Javed Aslam, Daniela Rus
Department of Computer Science
Dartmouth College
{liqun, jaa, rus}@cs.dartmouth.edu

ABSTRACT

This paper discusses online power-aware routing in large sensor networks. We seek to optimize the lifetime of the network. We develop an approximation algorithm called *max-min* zP_{min} that has a good empirical competitive ratio. To ensure scalability, we introduce a hierarchical algorithm, which is called zone-based routing.

1. INTRODUCTION

The proliferation of low-power analog and digital electronics has created huge opportunities for the field of wireless computing. It is now possible to deploy hundreds of devices of low computation, communication and battery power. They can create ad hoc networks and be used as distributed sensors to monitor large geographical areas, as communication enables for field operations, or as grids of computation. These applications require great care in the utilization of power. The power level is provided by batteries and thus it is finite. Every message sent and every computation performed drains the battery.

We focus on a global metric by maximizing the time to the partition of the network. We model this as the time to the failure of the first node. This metric is very important for ad-hoc networks where messages have to be delivered at high rates.

In this paper, we propose an online approximation algorithm for power-aware message routing that optimizes the lifetime of the network. Our algorithm, called the *max-min* zP_{min} algorithm, combines the benefits of selecting the path with the minimum power consumption and the path that maximizes the minimal residual power in the nodes of the network. We show that the *max-min* zP_{min} algorithm has a good competitive ratio in practice, approaching the performance of the optimal off-line routing algorithm under realistic conditions.

We propose another online algorithm called *zone-based routing* that relies on *max-min* zP_{min} and is scalable for large scale networks (Section 4). Our experiments show that the performance of zone-base routing is very close to the performance of *max-min* zP_{min} with respect to optimizing the lifetime of the network.

Zone-base routing is a hierarchical approach where the area covered by the (sensor) network is divided into a small number of zones. Each zone has many nodes and thus a lot of redundancy in routing a message through it. To send a message across the entire area we find a "global" path from zone to zone and give each zone control over how to route the message within itself.

Most previous research on ad-hoc network routing [17] focused on the protocol design and performance evaluation in terms of the message overhead and loss rate. To improve the scalability of routing algorithms for large networks, many hierarchical routing methods have been proposed in [11, 5, 12, 2, 7, 15, 13, 10]. This previous work focused on how to find the correct route efficiently, but did not consider optimizing power while sending messages. Singh et al. [18] proposed power-aware routing and discussed different metrics in power-aware routing. Minimal energy consumption was used in [16]. Chang and Tassiulas [3] also proposed maximizing the lifetime of a network when the message rate is known. The work presented in this paper is different from these previous results in that we develop online, hierarchical, and scalable algorithms that do not rely on knowing the message rate and optimize the lifetime of the network. Related results in sensor networks include [14, 1, 9, 6, 8, 4]. The high-level vision of wireless sensor networks was introduced in [14, 1]. Achieving energy-efficient communication is an important issue in sensor network design. Using directed diffusion for sensor coordination is described in [9, 6, 8].

2. FORMULATION OF POWER-AWARE ROUTING

Suppose a host needs power e to transmit a message to another host who is d distance away. We use the following formula to compute the power consumption for sending this message:

$$e = kd^c \ \ [3, 8],$$

where k and c are constants for the specific wireless system (usually $2 < c < 4$). We focus on networks where power is a finite resource. Only a finite number of messages can be transmitted between any two hosts. We wish to solve the problem of routing messages so as to maximize the battery lives of the hosts in the system. The lifetime of a network with respect to a sequence of messages is the earliest time when a message cannot be sent due to saturated nodes.

0. Find the path with the least power consumption, P_{min} by using the Dijkstra algorithm.
1. Find the path with the least power consumption in the graph.
    If the power consumption is greater than $z \cdot P_{min}$ or no path is found,
        then the previous shortest path is the solution, stop.
2. Find the minimal u_{tij} on that path, let it be u_{min}.
3. Find all the edges whose residual power fraction u_{tij} is no greater than u_{min}, remove them from the graph.
4. Goto 1.

Figure 1: *max-min* zP_{min}**-path algorithm**

3. AN ONLINE MAX-MIN ALGORITHM POWER-AWARE ROUTING

In this section we develop an approximation algorithm for online power-aware routing and show experimentally that our algorithm has a good empirical competitive ratio and comes close to the optimal.

Intuitively, message routes should avoid nodes whose power is low because overuse of those node will deplete their battery power. Thus, we would like to route messages along the path with the maximal minimal fraction of remaining power after the message is transmitted. We call this path the *max-min path*. Another concern with the max-min path is that going through the nodes with high residual power may be expensive as compared to the path with the minimal power consumption. Too much power consumption decreases the overall power level of the system and thus decreases the life time of the network. There is a tradeoff between minimizing the total power consumption and maximizing the minimal residual power of the network. We propose to enhance a max-min path by limiting its power consumption.

The two extreme solutions to power-aware routing for one message are: (1) compute a path with minimal power consumption P_{min}; and (2) compute a path that maximizes the minimal residual power in the network. We look for an algorithm that optimizes *both* criteria. We relax the minimal power consumption for the message to be zP_{min} with parameter $z \geq 1$ to restrict the power consumption for sending one message to zP_{min}. We propose an algorithm we call *max-min* zP_{min} that consumes at most zP_{min} while maximizing the minimal residual power fraction. The rest of the section describes the *max-min* zP_{min} algorithm, presents empirical justification for it, a method for adaptively choosing z and describes some of its theoretical properties.

The following notation is used in the description of the *max-min* zP_{min} algorithm. Given a network graph (V, E), let $P(v_i)$ be the initial power level of node v_i, e_{ij} the weight of the edge v_iv_j, and $P_t(v_i)$ is the power of the node v_i at the time t. Let $u_{tij} = \frac{P_t(v_i)-e_{ij}}{P(v_i)}$ be the residual power fraction after sending a message from i to j.

Figure 1 describes the algorithm. In each round we remove at least one edge from the graph. The algorithm runs the Dijkstra algorithm to find the shortest path for at most $|E|$ times where $|E|$ is the number of edges.

We conducted an experiment for evaluating the performance of the *max-min* zP_{min} algorithm.

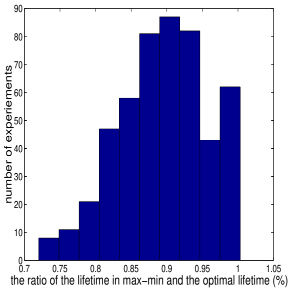

Figure 2: The figure shows the histogram that compares *max-min* zP_{min} **to optimal for 500 experiments. In each experiment the network consists of 20 nodes randomly placed in a 10*10 network space. The cost of messages is given by** $e_{ij} = 0.001 * d_{ij}^3$**. The hosts have the same initial power and messages are generated for hosts to one gateway host. The horizontal axis is the ratio between the lifetime of the** *max-min* zP_{min}**max-min algorithm and the optimal lifetime, which is computed off-line.**

Figure 2 shows the data that compares the *max-min* zP_{min} algorithm to the optimal routing strategy. We computed the optimal strategy by using a linear programming package[1]. We ran 500 experiments. In each experiment a network with 20 nodes was generated randomly in a $10 * 10$ network space. The messages were sent to one gateway node repeatedly. We computed the ratio of the lifetime of the *max-min* zP_{min} algorithm to the optimal lifetime. Figure 2 shows that $max - min$ zP_{min} performs better than 80% of optimal for 92% of the experiments and performs within more than 90% of the optimal for 53% of the experiments. Since the optimal algorithm has the advantage of knowing the message sequence, we believe that *max-min* zP_{min} is practical for applications where there is no knowledge of the message sequence.

4. ZONE-BASED ROUTING

The *max-min* zP_{min} algorithm requires accurate power level information for all the nodes in the network. For large scale sensor networks this is not a feasible assumption. Instead we propose to cluster together groups of sensors and estimate the overall routing power of the cluster for the purpose of the *max-min* zP_{min} algorithm. More specifically we propose to organize the network structurally in geographical zones, and hierarchically to control routing across the zones. The idea is to group together all the nodes that are in geographic proximity as a zone, treat the zone as an entity in the network, and allow each zone to decide how to route a message across[2]. The hosts in a zone autonomously

[1]To compute the optimal lifetime, the message rates are known. The max-min algorithm does not have this information.

[2]This geographical partitioning can be implemented easily using GPS information from each host.

direct local routing and participate in estimating the zone power level. Each message is routed across the zones using information about the zone power estimates. In our vision, a global controller for message routing manages the zones. This may be the node with the highest power, although other schemes such as round robin may also be employed.

If the network can be divided into a relatively small number of zones, the scale for the global routing algorithm is reduced. The global information required to send each message across is summarized by the power level estimate of each zone. We believe that in sensor networks this value will not need frequent updates because observable changes will occur only after long periods of time.

The rest of this section discusses (1) how the hosts in a zone collaborate to estimate the power of the zone; (2) how a message is routed within a zone; and (3) how a message is routed across zones.

4.1 Zone Power Estimation

The power estimate for each zone is controlled by a node in the zone. This estimation measures the number of messages that can flow through the zone. Since the messages come from one neighboring zone and get directed to a different neighboring zone, we propose a method in which the power estimation is done *relative to the direction* of message transmission.

The protocol employed by the controller node consists of polling each node for its power level followed by running the *max-min* zP_{min} algorithm. The returned value is then broadcasted to all the zones in the system. The frequency of this procedure is inversely proportional to the estimated power level. When the power level is high, the power estimation update can be done infrequently because a small number of messages routed through the zone will not change the overall power much. When the power level is low, message transmission through the zone is likely to change the power distribution significantly.

Without loss of generality, we assume that zones are square so that they have four neighbors pointed to the North, South, East, and West[3]. We assume further that it is possible to communicate between the nodes that are close to the border between two zones, so that in effect the border nodes are part of both zones. In other words, neighboring zones that can communicate with each other have an area of overlap (see Figure 3 upper).

The power estimate of a zone can be approximated as follows. We can use the *max-min* zP_{min} algorithm to evaluate the power level, find the *max-min* zP_{min} path, simulate sending Δ messages through the path, and repeat until the network is saturated. Δ is chosen to be proportionate to the power level of the zone.

More precisely, consider Figure 3 left. To estimate the power of zone B with respect to sending messages in the direction from A to C, let the left part of the overlap between A and C be the source area and the right part of the overlap between B and C the sink area. The power of zone B in the direction from A to C is the maximal number of messages that can flow from the source nodes to the sink nodes before a node in B gets saturated. This can be computed with the *max-min* zP_{min} algorithm. We start with the power graph of zone B and augment it. We create an imaginary source node S and connect it to all the source nodes. We create an imaginary sink node T and connect all the sink nodes to it. Let the weights of the newly added edges be 0. The *max-min* zP_{min} algorithm run on this graph determines the power estimate for zone B in the direction of A to C.

[3]this method can easily be generalized to zones with finite number of neighboring zones.

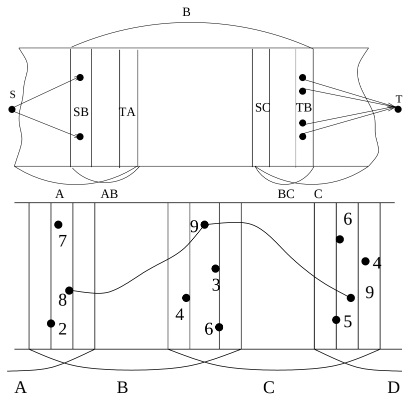

Figure 3: Three zones, A, B, and C. SB, SC are the source areas of B and C, and TA, TB are the sink areas of A and B. AB and BC are overlap border areas. The lower figure shows how to connect the local path in zone B with the local path in zone C. The number next to each node is the number of paths passing through that node in the power evaluation procedure.

---

choose Δ for the message granularity. $P = 0$;
repeat{
  Find the *max-min* zP_{min} path for Δ messages
  send the Δ messages through the zone
  $P = P + \Delta$
} **until** (some nodes are saturated)
return P

Figure 4: An approximation algorithm for zone power evaluation.

4.2 Global Path Selection

Given power-levels for each possible direction of message transmission, it is possible to construct a small zone-graph that models the global message routing problem. Figure 5 shows an example of a zone graph. A zone with k neighbors is represented by $k + 1$ vertices in this graph[4]. One vertex labels the zone; k vertices correspond to each message direction through the zone. The zone label vertex is connected to all the message direction vertices by edges in both direction.

[4]For square zones $k = 4 + 1$ as shown in Figure 5.

In addition, the message direction vertices are connected to the neighboring zone vertices if the current zone can go to the next neighboring zone in that direction. Each zone vertex has a power level of ∞. Each zone direction vertex is labeled by its estimated power level computed with the procedure in Section 4.1. Unlike in the model we proposed in the previous algorithm, the edges in this zone graph do not have weights. Thus, the global route for sending a message can be found as the max-min path in the zone graph that starts in the originator's zone vertex and ends in the destination zone vertex for the message. We would like to bias towards path selection that uses the zones with higher power level. We can modify the Bellman-Ford algorithm to accomplish this.

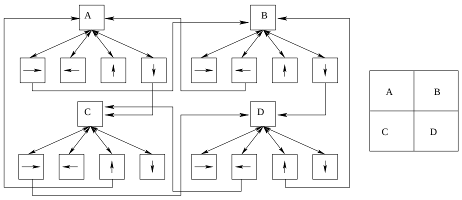

Figure 5: Four zones are in a square network field. The power of a zone is evaluated in four directions, left, right, up, and down. A zone is represented as a zone vertex with four direction vertices. The power labels are omitted from this figure.

4.3 Local Path Selection

Given a global route across zones, our goal is to find actual routes for messages within a zone. The *max-min* zP_{min} algorithm is used directly to route a message within a zone.

If there are multiple entry points into the zone, and multiple exit points to the next zone, it is possible that two paths through adjacent zones do not share any nodes. These paths have to be connected.

The following algorithm is used to ensure that the paths between adjacent zones are connected (see Figure 3 lower). For each node in an overlap region, we compute how many paths can be routed locally through that node when zone power is evaluated. In order to optimize the message flow between zones, we find paths that go through the nodes that can sustain the maximal number of messages. Thus, to route a message through zone B in the direction from A to C we select the node with maximum message weight in the overlap between A and B, we select the node with maximum message weight in the overlap between B and C, and compute the *max-min* zP_{min} paths between these nodes.

4.4 Performance Evaluation for Zone-based Routing

The zone-based routing algorithm does not require as much information as would be required by *max-min* zP_{min} algorithm over the entire network. By giving up this information, we can expect the zone-based algorithm to perform worse than the *max-min* zP_{min} algorithm. We designed a large experiment to measure how the zone-based algorithm does relative to the *max-min* zP_{min} algorithm.

We disperse $1,000$ nodes randomly in a regular network space (see Figure 6). The zone partition is described in the figure. Each zone has averagely 40 nodes. Each node sends one message to a gateway node in each round. The zone power evaluation protocol is executed after each round. By running the *max-min* zP_{min} algorithm, we ran the algorithm for about 41000 messages before one of the hosts got saturated. By running the zone-based routing algorithm, we got about 39000 messages before one of the nodes got saturated. The performance ratio between the two algorithms in terms of the lifetime of the network is 94.5%. Without the zone structure, the number of control messages on the power of each node in every information update is 1000, and they need to be broadcasted to 1000 nodes. In zone-based algorithm, the number of control messages is just the number of the zones, 48 here, and they are broadcasted to 24 zones after the zone power evaluation. And the zone-based routing dramatically reduces the running time to find a route in our simulation. More experiments on some other network spaces are ongoing.

5. CONCLUSION

We developed an online algorithm called the *max-min* zP_{min} algorithm and showed that it had a good empirical competitive ratio to the optimal off-line algorithm that knows the message sequence. We also proposed a hierarchical. Zone-based power-aware routing partitions the ad-hoc network into a small number of zones. Each zone can evaluate its power level with a fast protocol. A global path for each message is determined across zones. Within each zone, a local path for the message is computes so as to not decrease the power level of the zone too much.

6. REFERENCES

[1] Jon Agre and Loren Clare. An integrated architeture for cooperative sensing networks. *Computer*, pages 106 – 108, May 2000.

[2] A.D. Amis, R. Prakash, T.H.P. Vuong, and D.T. Huynh. Max-min d-cluster formation in wireless ad hoc networks. In *Proceedings IEEE INFOCOM 2000. Conference on Computer Communications*, March 2000.

[3] Jae-Hwan Chang and Leandros Tassiulas. Energy conserving routing in wireless ad-hoc networks. In *Proc. IEEE INFOCOM*, Tel Aviv, Israel, Mar. 2000.

[4] Yu Chen and Thomas C. Henderson. S-NETS: Smart sensor networks. In *Seventh International Symposium on Experiemental Robotics*, Hawaii, Dec. 2000.

[5] B. Das, R. Sivakumar, and V. Bharghavan. Routing in ad hoc networks using a spine. In *Proceedings of Sixth International Conference on Computer Communications and Networks*, Sept. 1997.

[6] Deborah Estrin, Ramesh Govindan, John Heidemann, and Satish Kumar. Next century challenges: Scalable coordination in sensor networks. In *ACM MobiCom 99*, Seattle, USA, August 1999.

[7] M. Gerla, X. Hong, and G. Pei. Landmark routing for large ad hoc wireless networks. In *Proceedings of IEEE GLOBECOM 2000*, San Francisco, CA, Nov. 2000.

[8] W. Rabiner Heinzelman, A. Chandrakasan, and H. Balakrishnan. Energy-efficient routing protocols for wireless microsensor networks. In *Hawaii International Conference on System Sciences (HICSS '00)*, Jan. 2000.

[9] Chalermek Intanagonwiwat, Ramesh Govindan, and Deborah Estrin. Directed diffusion: A scalable and robust communication paradigm for sensor networks. In *Proc. of the Sixth Annual International Conference on Mobile Computing and Networks (MobiCOM 2000)*, Boston, Massachusetts, August 2000.

[10] Mario Joa-Ng and I-Tai Lu. A peer-to-peer zone-based two-level link state routing for mobile ad hoc networks. *IEEE Journal on Selected Areas in Communications*, 17, Aug. 1999.

[11] P. Krishna, N.H. Vaidya, M. Chatterjee, and D.K. Pradhan. A cluster-based approach for routing in dynamic networks. *Computer Communication Review*, 27, April 1997.

[12] A.B. McDonald and T.F. Znati. A mobility-based framework for adaptive clustering in wireless ad hoc networks. *IEEE Journal on Selected Areas in Communications*, 17, Aug. 1999.

[13] M.R. Pearlman and Z.J. Haas. Determining the optimal configuration for the zone routing protocol. *IEEE Journal on Selected Areas in Communications*, 17, Aug. 1999.

[14] G. J. Pottie and W. J. Kaiser. Wireless integrated newtork sensors. *Communications of the ACM*, 43(5):51–58, May 2000.

[15] S. Ramanathan and M. Steenstrup. Hierarchically-organized, multihop mobile networks for multimedia support. *ACM/Baltzer Mobile Networks and Applications*, 3(1):101–119, June 1998.

[16] Volkan Rodoplu and Teresa H. Meng. Minimum energy mobile wireless networks. In *Proc. of the 1998 IEEE International Conference on Communications, ICC'98*, volume 3, pages 1633–1639, Atlanda, GA, June 1998.

[17] Elizabeth Royer and C-K. Toh. A review of current routing protocols for ad hoc mobile wireless networks. In *IEEE Personal Communication*, volume 6, pages 46 – 55, April 1999.

[18] S. Singh, M. Woo, and C. S. Raghavendra. Power-aware routing in mobile ad-hoc networks. In *Proc. of Fourth Annual ACM/IEEE International Conference on Mobile Computing and Networking*, pages 181–190, Dallas, TX, Oct. 1998.

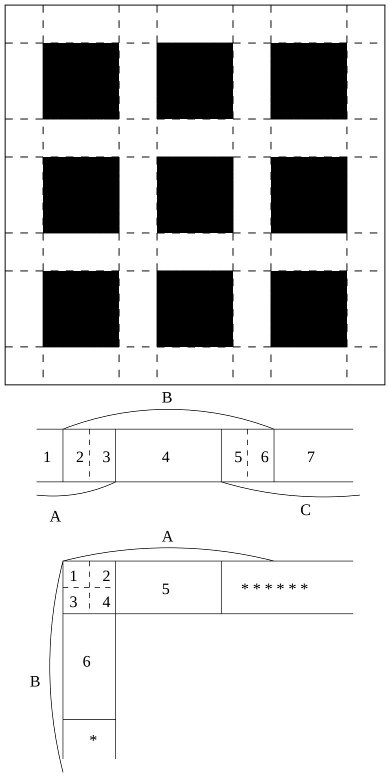

Figure 6: The scenario used for the zone-based experiment. The network space is a $10*10$ square with nine buildings blocking the network. Each building is of size $2*2$, and regularly placed at distance 1 from the others. The sensors are distributed randomly in the space nearby the buildings. Each sensor has an initial power of 4000. The power consumption formula is $e_{ij} = 10 * d_{ij}^3$. We partition the network space into 24 zones, each of which is of size $1*4$ or $4*1$, depending on its layout. For each zone, there is another corresponding zone with the same nodes but with opposite direction. For example, in the second figure, area $2,3,4,5,6$ constitute a zone, with 2 and 6 its source and sink areas; and $6,5,4,3,2$ constitute another zone with 6 and 2 its source and sink areas. We have a total of 48 zones. The lower figures show the layout of the neighboring zones. In the second figure, 3 is the sink area of the zone A, and 5 is the source area of zone C. The border area of A and B is $2,3$; and the border area of B and C is $5,6$. The third figure shows two perpendicular zones. The source area of B is $1,2$. The border area of A and B is $1,2,3,4$.