

Maintenance Efficient Routing in Wireless Sensor Networks

Andre Barroso, Utz Roedig and Cormac Sreenan
Mobile & Internet Systems Laboratory, University College Cork, Ireland
Email: {a.barroso u.roedig c.sreenan}@cs.ucc.ie}

Abstract

This paper presents an analysis framework of routing protocols that can be applied to produce sensor fields that are much less expensive to maintain. The framework is based on a maintenance model that is simple, yet flexible enough to capture real world deployment scenarios of sensor networks. As an illustration, the framework is used to assess the impact of different forwarding techniques for a known geographical routing protocol on the overall maintenance costs of different sensor fields. The results obtained indicate that an one-size-fits-all approach for the design of maintenance efficient routing protocols does not hold in large deployments of wireless sensor networks. However, savings of up to 50% in maintenance cost were observed through simple modifications of the forwarding strategy.

1. Introduction

As sensor nodes become cheaper and smaller the dominating cost factors of wireless sensor networks will lie in deployment and maintenance. Deployment will generally involve trained personal and specialized equipment that may include airplanes to drop sensors over areas that cannot be accessed otherwise. Some sensor units may instead require careful placement in the field thus consuming many hours of qualified labor. Furthermore, in long lived systems it is necessary to keep the network operational for a period of time that surpasses the lifetime provided by the batteries when the network was first deployed. Maintenance will thus be required and will involve periodic recharging or replacement of batteries in sensor fields that cannot harvest enough energy from the environment to remain operational. The use of large batteries or redundant nodes are solutions that may defer but cannot indefinitely prevent the need for manual intervention. Given the potential high costs, an appropriate design of sensor networks must take deployment and maintenance needs into consideration.

The frequency of maintenance operations in a sensor field is essentially dependent on the way nodes are de-

pleted. Given the impact of communication in the energy consumption of sensor nodes, the field depletion profile can be greatly influenced by the traffic flow inside the network. Protocols in the network layer can therefore help the shaping of favorable depletion profiles according to some appropriate metric that captures the concept of maintenance efficiency.

This paper presents an analysis framework of routing protocols that can be applied to produce sensor fields that are much less expensive to maintain. The framework is based on a maintenance model that is simple, yet flexible enough to capture real world deployment scenarios of sensor networks. To the best of our knowledge, this is the first work in the literature to propose a deeper analysis on the costs of deploying and operating wireless sensor networks by way of improving their design.

2. Maintenance in Wireless Sensor Networks

2.1. Concepts and Assumptions

Although the term maintenance may encompass a vast number of different activities, this study adopts the restrictive definition of a maintenance operation as the replacement of batteries from one or more nodes in the sensor field. Human intervention in the field required by equipment damage or malfunctioning is intentionally omitted from this definition. Since the study aim is the design of protocols able to affect maintenance costs by altering how batteries are depleted, sensor nodes are assumed to be reliable. If necessary, sensors with limited expected lifetime can be replaced together with batteries.

The purpose of evaluating the impact of routing protocols in the maintenance operations also requires the assumption that batteries depletion is dominated by communication activity. Clearly, if battery depletion is independent on the traffic generated or forwarded in a node, routing algorithms cannot affect maintenance. This assumption holds in systems which are kept in sleeping states and wake up only to collect/generate/forward data, a mode of operation typical for the motivation scenarios presented next.

2.2. Motivation Scenarios

The problem of maintenance in wireless sensor networks can be better understood by centering on examples of long term deployments where human intervention is periodically needed to keep the field operational. Two such scenarios are presented in the following paragraphs. The scenarios are later used in the discussion of maintenance models.

Precision Agriculture. A promising application of wireless sensor networks is in the instrumentation of farming areas to collect data which can be used to enhance crop yields.

Crop fields may cover extensive areas and the overlaid sensor network will require periodic battery recharging for continuous operation in situations where energy cannot be extracted from the environment or the extraction is insufficient. The access to individual sensors are likely to be unproblematic as crops are laid for easy reaping. Nevertheless, sensors may be spread over large areas and great distances might be covered to access nodes requiring battery replacement.

Habitat Monitoring. Wireless sensor networks have also been deployed successfully for the monitoring of living creatures in their habitat. The wireless feature of the technology ensures the system is minimally intrusive.

The structure and characteristics of a sensor field used in habitat monitoring can vary greatly depending on the type of habitat considered. As an example, the monitoring network may cover a small area in which sensors are placed in locations with very distinct access requirements. Some sensors may be located on the top of trees, others underwater and still other may be on the ground. In each case, the access to a node requiring battery replacement may involve different equipment, personnel, times and consequently different costs.

2.3. Maintenance Model

In order to formalize the study of the impact of how battery depletion affects maintenance costs, a model is required that describes when batteries should be replaced and the costs involved in the replacement. In the following paragraphs the policies and cost structure (the cost model) for the maintenance of wireless sensor networks are described. The maintenance policy and cost model define a maintenance model for the network. Later, we show how the maintenance model is used to assess maintenance aware routing protocols.

Maintenance policy

The specifics of the maintenance operations and their frequency are defined by the maintenance policy. A simple policy might have the following structure:

A maintenance operation is triggered every time a node has less than 10% of its initial

battery charge remaining. During the maintenance operation, the battery of the depleted node is recharged/replaced.

Different policies are possible depending on the particularities of the sensor field, resources available and application requirements.

Although battery recharging of a single node per maintenance operation might be reasonable under certain circumstances, several scenarios justify multiple battery replacements per maintenance operation. This is specially true in fields where the cost to access the depleted sensor is dominant. In these cases, it makes sense to recharge batteries of all nodes in the vicinity of the depleted sensor even though they are not technically considered depleted.

The concept of vicinity is variable from sensor field to sensor field. In the habitat monitoring scenario described in Section 2.2, nodes placed in a same tree could be considered as part of the same vicinity. The same can be said of sensors placed in a same pond. We refer to a group of nodes in the same vicinity as a maintenance zone. More formally:

A maintenance zone is a set of sensor nodes S such that for every pair of sensors $s_1, s_2 \in S$, the cost of accessing s_2 from s_1 is negligible.

As the definition implies, sensors in a field are primarily grouped into maintenance zones according to the maintenance cost model. Indeed, if the access cost to sensors is defined by the geometric distance to reach them, then all sensors physically close to each other can be grouped into a single zone. If access costs to sensors are defined entirely by their absolute position in the field then two sensors physically close to each other, one on a tree and the other in a pond, may not belong to the same zone.

In this study, for simplicity, maintenance zones are always a partition of the area covered by the sensor field.

Cost Model

Sensor fields may contain nodes underneath water, on the top of hills or spread over a large flat area. In each of these situations, the equipment, the personnel and the effort necessary to perform a maintenance operation have different characteristics that will affect the maintenance cost.

The cost of servicing a sensor in a sensor field can be divided in many different ways. It suffices for the purpose of this study to decompose the total cost in two factors:

- *Access Cost* : one-time resources spent while accessing the sensor to be serviced.
- *Recharging Cost* : one-time resources spent while servicing an individual sensor in its current location in the sensor field. It includes the battery and hardware replaced.

These cost components are added to produce the maintenance cost $C_m(s)$ of servicing a single node s in the sensor field:

$$C_m(s) = \text{Access Cost} + \text{Recharging Cost} \quad (1)$$

A concrete example of a cost model helps the understanding of the concepts just presented. Consider the deployment of wireless sensor nodes for environmental monitoring in redwood trees at University of California Botanical Garden's Mather Redwood Grove [10]. In this deployment, several sensor nodes are attached in different positions of trees that can be hundreds of feet tall. Climbing equipment is used to deploy such sensors. The access cost to reach an individual node may include the vehicle/fuel used to reach the sensor field and the labor cost of the people involved. Such costs are proportional to the distances involved and the difficulty involved in climbing the trees. Once the sensors are reached, the recharging cost involves the batteries replaced.

2.4. Applying the Maintenance Model

The model previously defined can be used to quantify the total cost of maintaining a sensor field. Indeed, every maintenance operation incurs a cost defined by equation 1. During the lifetime of a sensor field, I maintenance operations will take place. The total maintenance cost C_t of the sensor field, maintained according with policy P , is then given by:

$$C_t(P) = \sum_{i=0}^I C_{m_i} \quad (2)$$

This cost can be improved by reducing either C_m or I . Reduction of C_m cannot be achieved by altering the profile of battery depletion in the sensor field. The value of this parameter is dependent on elements such as transportation costs, labor costs, etc. Therefore, the goal of reducing total maintenance cost by minimizing C_m is out of the scope of this study. Reduction of I , however, can be achieved by a combination of factors including the choice of maintenance policy, design and operation of sensor networks as discussed next.

2.5. Towards Maintenance Efficiency in Wireless Sensor Networks

A maintenance policy impacts total maintenance costs of a sensor field by affecting the frequency in which nodes are recharged. As mentioned in Section 2.3, sensor fields where the cost of a maintenance operation is solely defined by the access cost should opt for a policy in which nodes in

the same maintenance zone are replaced concurrently. This procedure reduces the frequency I of operations in the field without adding any extra cost (the only cost is incurred in accessing the zone).

The consequence of adopting the strategy of recharging the battery of several nodes at a single access cost is the possible presence of nodes in a zone that are far away from depletion when recharged. From the perspective of a system designer, this situation means that individual maintenance zones unequally depleted at the moment of maintenance contain energy that should have been used to delay the maintenance operation, increase the reliability of communication or any other use. The energy was "wasted" since using it would bring benefits without incurring additional costs.

A system designer affects the maintenance cost by choosing to deplete the sensor field in accordance with a profile that shapes not only the number I of maintenance operations but also where these operations take place in the case of sensor fields where different zones incur different access costs. Some level of load balancing should be applied for uniform depletion of the sensor field as a whole and across the zones. On the other hand, depletion should take place more slowly in zones with high access costs.

Finally, maintenance costs are affected by the way applications use the network. As an example, an application may operate different sections of the network more intensely at given periods while other applications may continuously operate the field in a uniform way.

3. Maintenance Efficient Routing Protocols

Routing protocols hold a great potential for affecting maintenance costs since battery depletion can be shaped at a network level by controlling the flow path of packets. This section deals with the process of evaluating and designing routing protocols that try to incorporate maintenance efficiency in the system. The first step in this process is the definition of an appropriate metric able to capture this concept. Variations of a geographical routing protocol are used to exemplify evaluation and design for maintenance efficiency.

3.1. Maintenance Metrics

The design of maintenance efficient routing protocols requires the existence of appropriate metrics for their evaluation. Research in sensor networks, however, has focused primarily on energy efficiency as the main design goal of data dissemination protocols. Besides total energy consumption, another metric commonly used is network lifetime, defined as the time for the first node in the network to deplete or another similar condition. These metrics alone are nonethe-




Figure 1. Maintenance metric for precision agriculture scenario

less inappropriate for the design of maintenance efficient sensor networks since they oversimplify maintenance costs. Energy efficiency overstates the importance of energy consumed in the system by ignoring the fact that the cost of the Joules injected into batteries may be irrelevant to the overall maintenance cost. Maximizing network lifetime is an objective more aligned with the problem of achieving maintenance efficiency. However, it also neglects the costs differences of depleting the sensor field in different areas.

A suitable metric for the performance assessment of maintenance efficient routing protocols can be derived directly from the maintenance model presented in Section 2.3 in the form condensed by equation 2. The problem of instantiating a metric for a particular system requires knowledge of the maintenance policy P and the appropriate maintenance operation costs for the field. As an illustration, consider the definition of a maintenance metric for routing protocols on the two wireless sensor network scenarios described in Section 2.2. In both scenarios it is assumed a policy P that replaces nodes of an entire maintenance zone in each operation. Furthermore, the recharging cost is considered negligible.

Precision Agriculture Scenario. The sensor field is partitioned in artificial zones for maintenance purposes as depicted in Figure 2. Since recharging costs are assumed negligible, the maintenance operation cost is defined by the access costs to each zone. The crop field is assumed to cover an extensive area and the access to individual zones is unproblematic. Therefore the access cost can be characterized solely by the euclidean distance from the zone to a geographical point where the personnel involved in the maintenance starts their journey (all operations are assumed to begin in this point). In the figure, the maintenance cost for nodes in zone B is thus higher than for nodes in zone A.

Habitat Monitoring Scenario. As in the previous scenario, the sensor field is partitioned in artificial zones and maintenance operation cost is defined by the access costs to each zone. The field is assumed to cover a small area and zones lie in very different terrains. In Figure 2, the mainte-




Figure 2. Maintenance metric for habitat monitoring scenario

nance cost for nodes in zone A is higher than for nodes in zone B since the former contain nodes in a lake.

3.2. Design of Maintenance Efficient Routing Protocols

The maintenance model and respective metric constitute the starting point in the process of designing routing protocols for maintenance efficiency. In this section, variations of a geographical routing protocol are proposed, each exploring different aspects of the system and maintenance model. The efficiency of each approach is evaluated in Section 4. The choice of a geographical routing protocol as the basis of this study was based on its flexibility in finding alternative paths towards the destination when forwarding decisions are local. In general, as long as the geographical location of the destination is known, a path towards it can be found. It is therefore easy to generate variations of routes towards the desired node. Nonetheless, the reasoning presented in the study can be extended to other types of protocols, such as directed diffusion[3] and broadcast built trees rooted at the sink. The protocol variations presented next explore only the design space of packet forwarding based only on local information. No attempt is made to achieve global optimality on the solutions given the cost of obtaining the needed data in large scale networks.

GPSR. The Greedy Perimeter Stateless Routing (GPSR) protocol is a well known geographic routing protocol described in [5]. All nodes in GPSR must be aware of their position within a sensor field. Each node communicates its current position periodically to its neighbors through beacon packets. Upon receiving a data packet, a node analyzes its geographic destination. If possible, the node always forwards the packet to the neighbor geographically closest to the packet destination. If there is no neighbour geographically closer to the destination, the protocol tries to route around the "hole" in the sensor field.

GPSR-R. A GPSR variation in which a packet is randomly forwarded to any neighbour closer to the destination. This approach seeks to reduce the frequency of maintenance

operations in the sensor field by load balancing the traffic in the network.

GPSR-ME. Packets are forwarded to the neighbour node closer to the destination with the maximum battery energy level. This approach seeks to reduce the frequency of maintenance operations in the sensor field through a guided load balancing as opposed to the random approach of GPSR-R.

GPSR-MZ. Messages are forwarded to a neighbour node closer to the destination within a maintenance zone of minimum cost among the candidates. In case multiple candidates exist conforming to this criteria, one of them is chosen randomly. The approach tries to reduce the total cost of maintenance by avoiding expensive zones.

GPSR-MZME. Messages are forwarded to a neighbour node closer to the destination within a maintenance zone of minimum cost among the candidates. In case multiple candidates exist conforming to this criteria, the one with maximum remaining energy is chosen. The approach tries to reduce the total cost of maintenance by avoiding expensive zones. As a subordinate goal, the approach seeks to reduce the frequency of maintenance operations in the sensor field by load balancing the traffic in the network.

GPSR-MEMZ. Messages are forwarded to a neighbour node closer to the destination with the maximum battery energy level. In case multiple candidates exist conforming to this criteria, the one within a maintenance zone of minimum cost among the candidates is chosen. The approach tries to reduce the frequency of maintenance operations in the sensor field through load balancing. As a subordinate goal, the approach seeks to reduce the total cost of maintenance by avoiding expensive zones.

4. Evaluation

4.1. Experiment Setup

A comparative analysis of the merits of each variation of the geographical routing protocol was conducted through simulations in the precision agriculture and habitat monitoring scenarios presented in Section 2.2. In each scenario, the goal is to understand the impact of different forwarding criteria in the overall maintenance cost and energy consumption of a long-lived battery powered system. A detailed characterization of the experiment is presented next.

Simulator. For the experiments, a lightweight event-driven simulator was written in C++. Its main design objectives were simplicity and scalability with the size of the network. Given the level of abstraction required in this study, the option for a fast tool over more complex simulators such as ns-2 is justifiable. A node is able to transmit/receive packets to other nodes inside a well defined transmission range and each node incorporates a battery with a maxi-




Figure 3. Sensor field structure for the experiment

imum energy capacity. The energy level of a battery is reduced by a fixed amount with every packet transmitted. The access to the shared media is collision-free.

Sensor Field. The sensor field is a grid of wireless sensor nodes organized as depicted in Fig. 3. All sensors have the same specification and are equally spaced from each other. A full battery allows for 1000 packet transmissions. Maintenance zones were dimensioned to be much smaller than the total number of nodes in the field number. The grid is comprised of 324 nodes organized in a 18x18 arrangement and it is partitioned in 36 zones, each including 9 nodes. The radio range is adjusted so that only immediate adjacent nodes can communicate (i.e., 8 neighbors or less per node). A base station collecting reports from other sensors in the field is placed at the bottom left-most corner of the grid.

Maintenance Model. In both scenarios analyzed, the cost model was chosen to reflect the assumption that access costs will dominate the total maintenance cost of the sensor network. The maintenance policy is such that a maintenance operation is triggered every time a node has less than 10% of its initial battery charge remaining. During the maintenance operation, the batteries of ALL nodes in the same zone are recharged/replaced.

For the precision agriculture scenario, zone access costs are characterized solely by the euclidean distance from the zone to a geographical point where the personnel involved in the maintenance start their journey (all operations are assumed to begin at this point). For the habitat monitoring scenario, each zone is assigned an access cost from a fixed set of values. Figure 4 depicts the access costs for different sensor fields used in the experiments. Access costs in fields (a), (b) and (c) are proportional to the euclidean distance from each zone to points *A*, *B* and *C* respectively. In fields (d), (e) and (f) zone access costs can assume one of two values: 1 or 5. Zones with access cost 5 are scattered in the field in different ways to test how protocols can be af-

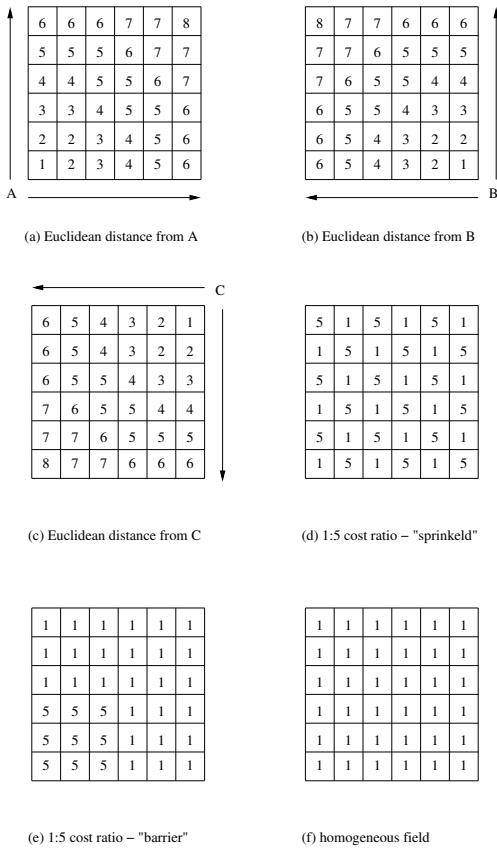


Figure 4. Access costs for simulated scenarios

ected by such variation.

Operation Model. It is assumed that, at any time, exactly one sensor is actively sending data to the base-station. This sensor is selected randomly within the sensor field. A node sends n messages before a new node is selected. Each message sent is separated from the previous one by an interval of 5 seconds. In this setup, the network generates the same number of data packets per unit of time for every run of the simulator.

4.2. Comparative Evaluation

The different forwarding criteria described in 3.2 are compared according to the following metrics:

- *Total Maintenance Cost*, computed as defined in eq. 2;
- *Total Number of Zone Accesses*;
- *Total Dissipated Energy*;
- *Average Number of Hops* defined as the ratio between the sum of hops traveled by each report message gen-

erated in the field and the total number of report messages.

The total maintenance cost indicates the impact of the forwarding technique on the actual cost of maintaining the sensor field. The total number of zone accesses indicates the frequency of accesses to maintenance zones during the experiment. For two protocols with similar maintenance costs, this metric reveals whether many low access cost or a few high access cost zones were depleted. The total dissipated energy indicates the number of Joules injected in the system. Finally, the average number of hops exposes the potential impact of the forwarding techniques on latency.

Experiment 1 - Precision Agriculture Scenario In the first experiment, routing protocols are simulated in fields where zone access costs are proportional to the euclidean distance of a single geographical point. Figure 5 shows the results for a field where the reference geographical point for computing costs coincides with the base station location (see fig. 4(a)). The field is operated in a way that each node sends a single message before a new node is randomly selected to report (i.e., $n = 1$).

In this configuration, according to fig 5(a) the most maintenance efficient forward technique is GPSR-MZME. The maintenance cost of this scheme is about 50% less than the cost presented by the canonical GPSR. A combination of two factors contributes to the good performance of GPSR-MZME. First, the average number of hops per message tend to be small with this scheme. Indeed, according to fig. 5(c), only GPSR presents a lower average number of hops. GPSR has lower average because it greedily forwards packets through the shortest geographical path available. GPSR-MZME, on the other hand, forwards packets preferably to nodes within zones of minimal access cost. Notice, however, that in the simulation arrangement, the sensor field presents a gradient of zone access costs that leads to the base-station. In other words, the decision of moving to a zone with lower access cost coincides with the goal of approaching the final destination. Short paths result in less energy consumption in the field as can be seen in figs. 5(c) and 5(d). If this extra present energy in the field is consumed properly, then energy efficiency will also translate into maintenance efficiency. This observation leads to the second contributing factor for the performance of GPSR-MZME, i.e., its subsidiary objective of forwarding packets to nodes with maximum energy level. In order to observe the benefit of adding this subsidiary goal, consider the number of zone accesses for both GPSR-MZ and GPSR-MZME in figure 5(b). The only difference between these two approaches is the presence of the mentioned subsidiary goal in GPSR-MZME. This feature greatly reduces the number of maintenance operations in the field which ultimately reduces the total maintenance cost.




Figure 5. Simulation results for sensor field of fig. 4(a) and $n = 1$

As a final remark, note that the most energy efficient scheme according to fig. 5(d) is GPSR, but this fact does not directly translate into maintenance efficiency as can be seen in fig. 5(a).

Impact of Euclidean Distance Reference. The performance of GPSR-MZME and GPSR-MZ in fields where zone access costs are proportional to the euclidean distance of a single geographical point is greatly dependent on the relative position of this point and the base-station. In situations where the reference geographical point and the base-station are approximately in the same direction, the goals of minimizing maintenance costs and delivering a packet to the destination are aligned. Indeed, forwarding packets to zones of minimum costs also means approaching the base-station. This synergy contributes to energy savings in the field since the path traveled by messages tend to be shorter. In scenarios where the base-station and reference geographical point are not in the same direction a conflict arises in the forwarding process. The effort of maintaining packets in zones of low access cost hinders the movement towards the base station resulting in longer paths. As a consequence of this conflict, GPSR-MZME and GPSR-MZ perform worse than the other forwarding approaches where zone access costs are not taken into consideration or are only subsidiary in the forwarding decision process. Figure 6 shows the total maintenance cost of the sensor field when the geographical reference point is respectively located at corners B and C of the sensor field (see fig 4). The graphs indicate the poor performance of GPSR-MZME and GPSR-MZ in comparison to the other techniques.




Figure 6. Total maintenance cost when base-station and euclidean distance reference are in different locations

Impact of Operation Dynamics. In the simulation, parameter n described in 4.1 controls the operation dynamics of the network. When $n = 1$, there is no correlation between the origin of two consecutive messages generated in the sensor field. Each message is generated at a random node and forwarded to the base-station according to a chosen scheme. As the value of n increases, the behaviour of the traffic assumes a hot-spot profile where traffic is originated from a sole location for extended periods of time before moving to a different point in the network. The impact of this parameter is specially noticeable in the comparative performance of GPSR with the other techniques. When the value of n is high, consecutive messages generated at the same node and forwarded through GPSR will follow always the same path. This behaviour causes the fast depletion of nodes along this path and an increase of the number of maintenance operations required. When n is low, this negative effect is counter-balanced by the randomization of the packet sources. This randomization creates a form of load balancing in the field that is independent of the forwarding technique. Other techniques already incorporate load balancing schemes and therefore are less affected than GPSR. Figure 7 depicts the total maintenance cost of the forwarding techniques when $n = 100$. Compare the values obtained with those shown in figure 5 where $n = 1$. GPSR performance is greatly reduced by the absence of randomization in the selection of forwarding paths with the higher value of n .

Experiment 2 - Habitat Monitoring Scenario In the second experiment, routing protocols are simulated in fields where zone access costs can assume one of two values, where one value is several times larger than the other. Experiments were conducted in three different arrangements of this structure as depicted in figures 4(d)-(f). In the organization of figure 4(d), expensive zones are “sprinkled” in the field. As shown in graph of figure 8(a), forwarding techniques GPSR-MZ and GPSR-MZME, tuned to avoid expensive zones, present the lowest maintenance costs under this scenario. Specifically, the maintenance cost of GPSR-




Figure 7. Total maintenance cost for sensor field of fig. 4(a) and $n = 100$




Figure 9. Total maintenance cost for sensor field of fig. 4(e) and $n = 1$




Figure 8. Simulation results for sensor field of fig. 4(d) and $n = 1$




Figure 10. Total maintenance cost for sensor field of fig. 4(f) and $n = 1$

MZ and GPSR-MZME are respectively 65% and 55% of the cost incurred by GPSR. Nonetheless, the overall number of maintenance accesses to the field are highest when these techniques are applied (figure 8(b)). This result occurs because bypassing expensive zones reduces the number of alternative routes in the fields creating hot spot points that deplete faster and require battery replacement more often. As expected, the same phenomena occurs in the sensor field configuration of figure 4(e). However, because expensive zones cannot be really bypassed, GPSR-MZ and GPSR-MZME do not perform as well as in the previous case as shown in figure 9. In both scenarios, the performance of GPSR-MZ and GPSR-MZME relative to the other techniques is very dependent on the cost ratio between the expensive and inexpensive zones. In general, the higher the ratio, the better the relative performance. On the other hand, if this ratio is close to one, then these protocols will under perform all the others since they require more accesses to the field. The simulation result when all zones have the same access cost is shown in figure 10. Note that in this case, GPSR-MZ degenerates into GPSR-R and GPSR-MZME degenerates into GPSR-ME. It is also worth mentioning that in this scenario, the objective of maximizing the network lifetime (defined as the time to first depletion) coincides with the objective of minimizing maintenance costs.

5. Related work

Research related to maintenance in wireless sensor networks has been generally restricted to the development of remote network reprogramming schemes and energy-efficient hardware/software. Tools allowing remote network programming have received special attention in recent years because of the recognition of the difficulties associated with accessing large number of nodes deployed in possibly hostile environments. Schemes such as Deluge[2] and Maté[6] were thus proposed to disseminate code in wireless sensor networks without human intervention.

Regarding the development of operational software for such networks, the focus has been on the extension of the lifetime of the system through the use of techniques that conserve batteries as much as possible. Several data dissemination protocols were designed to meet this objective. In [1], a flow augmentation algorithm is proposed to define paths that maximize the system lifetime in ad-hoc networks, where lifetime is defined as the length of time until the first battery drain-out among all nodes. The authors in [7] use a variation of this metric and develop an approximation algorithm called $max - min zP_{min}$ to solve the problem of routing messages so as to maximize the time until a message cannot reach its destination due to depleted nodes. Other papers addressing the problem of energy efficiency for data dissemination in wireless sensor networks are [9], [8] and [11].

Sensor networks can be powered primarily or secondarily by extracting energy from the environment [4]. This mechanism is known as energy scavenging or harvesting. Although energy scavenging is a promising technique in many applications for wireless sensor networks, there will be deployments where the only form of energy available is battery provided or the harvested energy is insufficient to cover the energy budget. In these scenarios, the need for periodic battery recharging will still exist for long term systems. To our knowledge, this is the first work in the literature to propose a deeper analysis of the costs of deploying and operating wireless sensor networks with means for improving their design.

6. Conclusions

Although energy efficiency is an essential aspect for the practical deployment of large scale wireless sensor networks, this attribute represents only one aspect of a multi-dimensional problem. Clearly, by solely extending the lifetime of batteries it is possible to extend the lifetime of a system, all other conditions being equal. This approach however does not address the cost of recharging individual batteries. As a consequence, systems that require less maintenance interventions may incur higher costs in practice. Furthermore, this approach largely ignores the effects of “repairing” the system after depletion. Indeed, different maintenance policies will affect the way batteries can be depleted in the system while preserving low costs. A group of batteries assigned to always be replaced together is for instance maintained at the rate of its fastest depleting node. In such a case, other nodes in this group can speed up their depletion rate without affecting the maintenance cost of the group. Nodes can use this “free” energy for different purposes, such as increased fault-tolerance or quality of service.

The experiments conducted in this study demonstrated that an one-size-fits-all approach for the design of maintenance efficient routing protocols does not hold in large deployments of wireless sensor networks. The variability of scenarios makes it impossible for a single technique to perform consistently well. Indeed, applying techniques that explicitly take into consideration the cost of accessing certain zones in the field might be counterproductive in scenarios where the objective of minimizing such costs is entirely at odds with the objective of delivering packets to their destination. This observation benefits protocol designer as much as the architect of the sensor field which should make sure conflict of objectives are minimized (for instance, through a proper placement of the base station). In the general case, the model and analysis technique introduced in this paper can be applied to produce fields that are much less expensive to operate. In the experiments conducted, savings of up

to 50% in maintenance cost were observed in the long run.

Acknowledgments

The support of the Informatics Research Initiative of Enterprise Ireland is gratefully acknowledged.

References

- [1] J.-H. Chang and L. Tassiulas. Energy Conserving Routing in Wireless Ad-Hoc Networks. In *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies*, volume 1, pages 22–31, 2000.
- [2] J. W. Hui and D. Culler. The Dynamic Behavior of a Data Dissemination Protocol for Network Programming at Scale. In *SenSys '04: Proceedings of the 2nd international conference on Embedded networked sensor systems*, pages 81–94. ACM Press, 2004.
- [3] C. Intanagonwivat, R. Govindan, D. Estrin, J. Heidemann, and F. Silva. Directed Diffusion for Wireless Sensor Networking. *Networking, IEEE/ACM Transactions on*, 11(1):2–16, Feb 2003.
- [4] A. Kansal and M. B. Srivastava. An Environmental Energy Harvesting Framework for Sensor Networks. In *Proceedings of the 2003 international symposium on Low power electronics and design*, pages 481–486. ACM Press, 2003.
- [5] B. Karp and H. T. Kung. GPSR : Greedy Perimeter Stateless Routing for Wireless Networks. In *Mobile Computing and Networking*, pages 243–254, 2000.
- [6] P. Levis and D. Culler. Maté: A Tiny Virtual Machine for Sensor Networks. In *ASPLOS-X: Proceedings of the 10th international conference on Architectural support for programming languages and operating systems*, pages 85–95. ACM Press, 2002.
- [7] Q. Li, J. Aslam, and D. Rus. Hierarchical Power-Aware Routing in Sensor Networks. In *DIMACS Workshop on Pervasive Networking*, Rutgers University, May 2001.
- [8] S. Lindsey and C. Raghavendra. PEGASIS: Power-Efficient gathering in sensor information systems. In *IEEE Aerospace Conference*, volume 3, pages 3–1125 – 3–1130. IEEE, 2002.
- [9] A. Manjeshwar and D. P. Agrawal. TEEN: A Routing Protocol for Enhanced Efficiency in Wireless Sensor Networks. In *Proceedings of the 15th International Parallel & Distributed Processing Symposium*, page 189. IEEE Computer Society, 2001.
- [10] S. Yang. Redwoods go high tech. UC Berkeley News, July 2003. Press Release.
- [11] Y. Yu, R. Govindan, and D. Estrin. Geographical and Energy Aware Routing: A Recursive Data Dissemination Protocol for Wireless Sensor Networks. Technical Report UCLA/CSD-TR-01-0023, UCLA, 2001.