

Benefits of Multiple Battery Levels for the Lifetime of Large Wireless Sensor Networks *

Mihail L. Sichitiu
Dept. of Electrical and Computer Eng.
North Carolina State University
Raleigh, NC 27695
Email: mlsichit@ncsu.edu

Rudra Dutta
Dept. of Computer Science
North Carolina State University
Raleigh, NC 27695
Email: rdutta@eos.ncsu.edu

Abstract

Large wireless sensor networks promise to revolutionize the way we sense and control the physical world. In the envisioned paradigm, large numbers of sensor nodes are scattered throughout the geographical area to be monitored. Due to the limited transmission range, data sensed by each sensor has to be forwarded in a multi-hop fashion before being delivered to a base station. The sensors closer to the base station have to forward comparatively more messages than sensors at the periphery of the network. Assuming that this increased activity of the sensors near the base station translates into a higher power consumption and each node has the same initial battery power, the nodes closer to the base station will deplete their batteries earlier. Besides the loss of the sensing capabilities of the nodes close to the base station, a more serious consequence of the death of the first tier of sensor nodes is the loss of connectivity between the nodes at the periphery of the network and the base station. Depending on the size of the network and the actual power consumed for transmission and reception, it is quite possible that a large percentage of the nodes will have almost full batteries, and yet be incapable of reaching the base station due to this loss in connectivity. Effectively, the lifetime of the network has expired. To alleviate this undesired effect, we investigate the effect of multiple battery levels to maximize the useful lifetime of the network. We demonstrate a powerful, yet attractively simple scheme to redistribute the total energy budget in multiple (but few) battery levels. We show by theoretical analysis, as well as simulation, that this substantially improves the network lifetime.

1 Introduction

Wireless sensor networks [1–3] will fundamentally change our ability to sense the physical world. Large numbers of inexpensive nodes will be deployed over the geographical area to be monitored. The large number of nodes offer a high degree of redundancy and increase the likelihood of some sensor being close enough to reliably sense interesting phenomena anywhere in the sensor field. The nodes send the sensor data to a base station, typically in a multi-hop fashion, using other sensor nodes in the process. Therefore, each node transmits its own traffic, as well as the traffic it forwards towards the base station on behalf of other nodes.

The sensor network can work in *continuous monitoring* mode, or be *event-driven*. A sensor node in continuous monitoring mode will take a sample at periodic intervals and send it (multi-hop) to the base station. Weather monitoring is perhaps best modeled by a continuous monitoring mode. An event-driven sensor node will only react to an event (e.g., a query or a pre-established trigger).

*This work was supported in part by the CACC and by NSF grant #0332271.

For example, a sensor deployed to detect forest fires may activate only when smoke or heat is detected. The scheme proposed in this paper applies to a continuous monitoring sensor network, where each of the sensor nodes sends approximately the same amount of traffic to the base station.

Like any other electronic device, sensor nodes have to be powered. A combination of several sensor network peculiarities makes the power problem especially challenging. First, the large number of cheap sensors precludes equipping each sensor with very large and expensive batteries. Second, in most situations, it is impossible to replace the batteries (especially in large, randomly scattered networks). Finally, many sensor networks are expected to work for days, weeks, and, sometimes, years.

Confirming the importance of the problem, many aspects of the problem have been extensively studied [4–24]. Some approaches [4–6] aim to minimize the transmission power expenditure, by reducing the power level of the transmitter, while still maintaining network connectivity. Medium access control (MAC) layer techniques [7–15] aim to conserve battery power by turning the receiver off whenever it is not needed (e.g., when a node starts overhearing a conversation between two other nodes, it can power down its transceiver until the medium becomes idle again). Several ad hoc power aware routing [16,17] algorithms choose routes in such a way as to maximize the lifetime of the network. It is clear that the power problem cannot be completely solved at any one single layer [23, 24].

In this paper we will approach a problem complementary to the ones considered in the literature. The motivation for this work stems from the observation that in a sensor network, regardless of the routing strategy, the sensor nodes closer to the gateway have to forward more packets than the ones at the periphery of the network. Assuming that this increase in workload results in an increase in power consumption (a reasonable assumption if the energy wasted in idle receiving is minimized using one of the existing approaches [7–14]), the nodes close to the gateway will die first, leading to a premature loss of connectivity in the sensor network. Note that this effect occurs regardless of the routing strategy, MAC layer, physical layer considerations, etc.

To alleviate this undesirable effect, we propose to use multiple levels of batteries, placed concentrically around the base station. We will show that given a total energy budget for the sensor network, the lifetime of the network can be significantly improved over the case where all sensors have a uniform battery level. The optimal solution is formulated theoretically and validated via simulations. The proposed scheme can be used *in addition* to other power conservation schemes proposed in literature.

For hand-placed sensors, the proposed multi-level battery strategy poses no significant problems. We consider the discussion of the practical aspects for other deployment options outside the scope of this paper.

The same effect of increasing the battery available at the inner tiers of the sensor network can be achieved by deploying a higher density of sensor near the base station (thus, resulting in a non-uniform node density); however, this is likely a more expensive solution (because in addition to the battery cost, there is also the cost of the extra sensors nodes), and may have its own deployment problems. Furthermore, depending on the application, beyond a certain node density there is no significant improvement in the monitoring abilities of the network.

However, if the higher density of nodes around the base station is combined with a scheme that keeps awake only a minimum number of sensors (e.g. [22, 25–28]) (as many as necessary to satisfy the monitoring requirements) the non-uniform density approach becomes equivalent to the proposed multiple battery level approach proposed in this paper.

2 Uniform Battery Level: The “Doughnut” Effect

In considering the batteries available to each node, it is most natural to think of the case of a unique battery level; that is, each node has an identical battery and, therefore, very closely the same battery level to start with. This can also be viewed as a uniform provisioning of the total energy budget available to the entire set of sensor nodes. This view is reasonable since every node is, after all, identical in other senses (in keeping with the ad hoc model). It is also the most practically useful model; since each node remains identical, there is no need to distinguish between two or more kinds of nodes when deploying them, say, by scattering from an aircraft.

This similarity of nodes extends to the roles they play in networking terms: each node is expected to originate traffic destined for the monitoring station as well as forward traffic originated by other nodes. However, after the deployment, there is a clear quantitative variation among the nodes in terms of the relative sending and forwarding they have to do. This is a consequence of the unique traffic pattern in sensor networks. All nodes source data to be sent to the monitoring station, which is the only sink node: an *egress* traffic pattern; hence, nodes closer to the monitoring station forward more data in a given period of time, whereas all nodes source data at the same rate (assuming identical sensing responsibilities for all nodes). Therefore, nodes closest to the monitoring station deplete their batteries at the highest rates and die first.

For the purpose of what follows, we assume that the area covered by the sensors is a circular area with the monitoring station in the center. This is a reasonable assumption and useful as an idealization because it is attractively simple and symmetric which eases analysis. However, note that the general conclusions of this paper are not restricted to this specific scenario. In fact, in any realistic spatial distribution of sensors, it is likely that a small fraction of the sensor nodes will be within direct radio range of the monitoring station, and increasingly larger fractions will be further and further away. Indeed, most such distributions can be thought of as a connected sub-area (usually regular) of the circular area we describe. For example, the field of sensors in a bridge embedded with sensors, in which the monitoring station is at one end, is a rectangle around a radius of the circular area. Even three-dimensional fields of sensors can be considered as extensions of this model if one of the dimensions is very much smaller than another, e.g., a building with

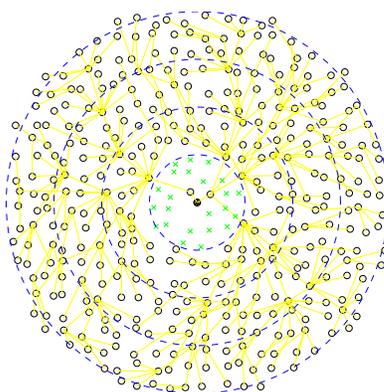


Figure 1: The “doughnut” effect: many nodes are alive; but, as soon as the last two surviving nodes in the first tier die, the network will be disconnected.

the monitoring station in the basement. Finally, scenarios where there are multiple monitoring stations can be decomposed into sets of sensor nodes such that each node reports to one monitoring station, and each such set treated as a field with a single monitoring station.

In such a circular field, consider the set of nodes close to the monitoring station at the center that can communicate directly with it. We refer to these one-hop neighbors of the monitoring station as the *first tier of nodes*. Since no other node can reach the monitoring station directly, traffic from every other node will have to be forwarded, in the last hop, by one of these first tier nodes (the traffic generated by these nodes themselves is sent directly to the sink and not forwarded). Similarly, the two-hop neighbors of the monitoring station will forward data for all nodes except the one-hop neighbors and themselves, etc. It is clear that the first tier nodes perform the largest amount of forwarding and will die first. If the spatial distribution of nodes is close to uniform, then the traffic load is equally distributed spatially. Each first tier node will forward roughly the same amount of traffic, and all first tier nodes will die at times very close to each other, after the network is first put into operation. Once all of the first tier nodes are dead, no other node will be able to send data to the monitoring station, and the lifetime of the network will be over by any sensible definition. Figure 1 shows the “doughnut-like” distribution of dead and live nodes in such a situation. Increasing the number of nodes in the network accentuates this effect, since there is more traffic to forward and the first tier of nodes has a smaller share of the total energy budget. For very large areas, most of the total energy budget could remain unexpended when disconnection takes place.

Our contributions in this paper are as follows. We obtain a precise quantitative description of this phenomenon, and thereby obtain the lifetime of the network when each node starts with the same battery. We show that the concept of network lifetime must be carefully considered in this context, because many intuitively obvious definitions of lifetime are not robust to special cases, and we supply such a robust definition. We consider various practical design problems of lifetime maximization: when battery budget is unconstrained but battery levels are constrained and vice versa, and also with constraint on the number of distinct battery levels but not the levels themselves. Our numerical results validate our theoretical work, and the magnitudes of the lifetime differences also go to show the importance of dealing with this phenomenon.

3 Quantitative Model and Analysis

3.1 Assumptions

In attempting to develop an analytical model to quantify the doughnut effect, we make several assumptions. These are broadly of two types: basic ones that reflect likely scenarios in the sensor network context, and idealizations to simplify the analysis.

First, we continue with the circular field of sensors, although our work is readily applicable to fields of other shapes as explained above.

We assume that we have a continuous monitoring sensor network [29, 30] with each node *periodically* sensing the environment and generating a data packet. That is, we assume that each node sources packets of one standard size at regular intervals P .

The routing is assumed to be fixed, by shortest hops, and unvarying. Realistically, even if routing is unvarying, the routes will be recomputed when a node that was acting as a gateway for traffic from some other nodes dies. This is what we assume in both our modeling and in our simulations, as we explain in Section 4. We also assume that the bandwidth of the network (characterized by the bandwidth of the medium close to the monitoring station) is sufficient to allow the entire generated traffic to be delivered to the monitoring station, again, reasonable, since sensor traffic is not expected to be voluminous. The delay seen by source to sink traffic is assumed to be irrelevant to our analysis.

We assume that all of the power consumed by a sensor node belongs to one of these three classes:

- **Power for generating a packet** includes the power necessary to sense the environment, i.e., powering the sensor, processing the sensed data, and transmitting the sensed data;
- **Power for forwarding a packet** includes the power used to receive and transmit a packet from another sensor node;
- **Power unrelated to data flows** includes all other network operations, e.g., routing, time synchronization, localization, idle listening, etc.

The above model is very general and applies to sensor networks employing a variety of networking protocols (MAC, routing, transport, etc.). For example, the MAC idle listening and the routing overhead adds to the power unrelated to data flows. Transmitting a packet (including the backoffs and retransmissions) contributes to the power for forwarding packets. The more efficient networking protocols have a reduced power unrelated to data flows.

The power unrelated to data flows can easily dominate the overall power consumption in a real implementation (especially due to idle listening), and several methods to save this power have been proposed [13, 23, 30, 31]. We assume that the power consumed transmitting and forwarding packets dominates the power consumed for the other tasks. If this assumption does not hold (i.e., the nodes consume approximately the same energy regardless of the number of flows forwarded), all nodes will deplete their batteries at approximately the same time (in which case there is no doughnut effect and no problem to solve).

For simplicity, we assume a fixed and known transmission radius for all nodes in the network. We ignore transmission errors (and, hence, retransmissions), although we can easily take those into account by adjusting the energy necessary to generate a packet and the power to forward a packet by the factor $\frac{1}{1-\rho}$ where ρ is the average packet error ratio.

We assume that the sensor nodes are distributed with roughly uniform density over the entire sensing area. This is a common assumption resulting from the need to assure a uniform level of monitoring quality through the sensor network, and the underlying assumption that all regions of this overall area are of equal interest to the sensing application, or are equally desired to be monitored.

We believe that these assumptions accurately reflect a large and important class of sensor networks, and consequently, the work we present in this paper is of high relevance and importance.

3.2 Notation

We first introduce some notation for quantities which are available from the specification of the problem instance. Some notation defined above is repeated for ease of reference.

N The initial number of nodes in the network.

R_M The radius of the network (i.e., the distance between the base station located at the center and the periphery nodes - notice the assumption of a circular network).

R_{TX} The transmission radius of the sensor nodes (assumed to be fixed).

B_i The sum of the battery powers available at the nodes of tier i .

P The period between the generation of two consecutive packets.

e_F The energy required to forward a packet.

e_G The energy required to generate a packet.

e_I The energy expended by each node in each period P that is independent of the data traffic.

L The lifetime of the network (explained in detail in Section 3.4).

3.3 Model

We define **tier** i as the set of the nodes that can reach the base station in i hops ($1 \leq i \leq T$).

We further define the following quantities which we will use in our analysis:

T The number of tiers in the network

$$T \approx \frac{R_M}{R_{TX}}; \quad (1)$$

T is typically larger than $\frac{R_M}{R_{TX}}$, but the approximation holds well for large, dense networks. Without a loss of generality, we can assume that T is an integer number (the results hold if T is real as well, but the explanations are easier to follow with the assumption of an integer T).

B The total energy budget

$$B = \sum_{i=1}^T B_i. \quad (2)$$

We consider the battery budget B as the main cost function. For simplicity we assumed the cost of the battery directly proportional with its size. The results in this paper can be easily extended to a more general cost function for each battery size.

N_i The number of nodes at tier i . With the assumption of a uniform density

$$N_i = \frac{(2i-1)N}{T^2}. \quad (3)$$

F_i The total number of flows carried by *all* nodes at tier i ; Each flow corresponds to the data generated by one sensor node (that has to be forwarded to the base station). All nodes at tier i will carry the flows for all nodes at tiers $i+1 \dots T$.

$$F_i = N - \sum_{j=1}^i N_j = \frac{N(T^2 - i^2)}{T^2} \quad (4)$$

L_i The lifetime of the nodes at tier i can be determined by considering the consumption components that consume the total available battery at level i :

$$B_i = \underbrace{\frac{L_i F_i e_F}{P}}_{\text{forwarding packets}} + \underbrace{\frac{L_i N_i e_G}{P}}_{\text{generating packets}} + \underbrace{\frac{L_i N_i e_I}{P}}_{\text{independent of packets}}. \quad (5)$$

Thus

$$L_i = \frac{B_i P}{F_i e_F + N_i (e_G + e_I)}, \quad (6)$$

$$= \frac{T^2 B_i P}{N (T^2 - i^2) e_F + (2i - 1) N (e_G + e_I)}. \quad (7)$$

3.4 Network Lifetime

The lifetime of the sensor network is less than trivial to define. The general essence of network lifetime is clear but some reflection is required to obtain a precise description that is at the same time simple and yet accurately captures this essence for a large variety of network parameters. Clearly, the sensor network should be considered dead once it cannot satisfy its monitoring requirements. With each sensor node that cannot report its measurements (due to battery depletion, failure, or lack of a usable route), the monitoring capability of the sensor network is decreased.

A commonly used monitoring requirement is that every point in the coverage area of the sensor network has to be in the monitoring range of at least X sensors, with $X \geq 1$. When deployment is not completely controlled, the possibility of random variation in the sensor positions must be admitted, and a more appropriate requirement is that with a probability greater than $p_c < 1$, every point in the area of coverage of the network should be in the monitoring range of no less than X nodes.

As the number of sensor nodes in the networks decreases, the quality of the coverage decreases. At some point the monitoring quality will be under a minimum (application dependent) threshold, and the network can be considered dead.

A reasonable definition of the lifetime of the network, therefore, would probably be: the time when the fraction of network nodes that cannot communicate data to the monitoring station (either because they are dead or because all nodes which could forward data from them are dead) exceeds some predetermined fraction $x\%$. The value of x would depend on the particular circumstances. But how is a general design method to proceed independent of x ? Fortunately, because of the doughnut effect, the actual lifetime is quite insensitive to x . Consider what happens when all nodes in tier i are dead: the monitoring station receives no data from any nodes in that tier *or further out*. Thus whenever a complete tier dies, there is a very large jump in the number of dead or disconnected nodes, far more than occurs due to random variation in the lifetimes of different nodes in individual tiers. Now consider that the number of nodes is the smallest in tier 1, which is most likely to die first due to the doughnut effect. Hence the lifetime of the network is virtually the same whatever the value of x , as long as $x > N_1/N$. This is graphically clear in Figure 2, which was obtained by direct numerical simulation of this effect, and which we discuss further in Section 4. One moment, more than 90% of the nodes are reporting data, in the next moment 0% are reporting. This effect would be less pronounced if by reallocation of batteries we allowed tier 1 to live longer and some other tier to die completely first. But the moment when the first complete tier to die did so, there would still be a significant jump in the fraction of non-reporting nodes, and very likely it would go from less than x to more than x , for reasonable choices of x .

A final possibility remains, that the death of nodes is not radially symmetric around the monitoring station. Thus all nodes in an asymmetric but large area might die or be disconnected, but still fall short of the fraction x . For example, all the nodes in a semicircular area of the circular sensor field shown in Figure 1 may die, but may still constitute only 49% of the original number of nodes deployed. By the above definition the network would still be alive, and our argument

about the tiers would fail. We disregard this possibility because this is directly at odds with the reasonable assumption that all regions of the area initially covered by sensors are equally desired to be sensed. A decrease in the quality of sensing over the entire area may be acceptable up to a given degree represented by x , such as would happen if nodes die randomly so as to reduce the density of nodes around any given region of the sensing area. But the complete loss of sensing in large regions of the area would not be acceptable, even if perfect sensing was still available from other regions.

In accordance with all the above considerations, in the remainder of this paper, we define the lifetime of the network as the minimum of the lifetimes of its tiers:

$$L = \min_{i=1,\dots,T} L_i. \quad (8)$$

3.5 Design Problems

Using the model formulated in Section 3.3, we will formulate several problems that are likely to be of practical importance to network designers, and address them in the context of this model. Our first problem is the obvious one of determining the lifetime of the network under the simple uniform battery provisioning.

Problem 1 *Given an energy budget B , **predict the lifetime** L of a sensor network with the parameters $N, R_M, R_{TX}, P, e_F, e_G$ and e_I , when the total battery budget is distributed uniformly among all network nodes.*

In this case, each node will start with the same level of battery, namely B/N ; we denote this unique level of battery by b_u . Then, the battery level at each tier is

$$B_i = \frac{N_i B}{N} = N_i b_u. \quad (9)$$

Hence, (7) becomes

$$L_i = \frac{BP}{N \frac{T^2-i^2}{2i-1} e_F + N(e_G + e_I)}. \quad (10)$$

Since the term $\frac{T^2-i^2}{2i-1}$ is monotonously decreasing with i for $1 \leq i \leq T$, L_i the lifetime of tier i is monotonously increasing with i . In other words, the lifetime of the network L is equal to the lifetime of the first tier (thus confirming the intuition in Section 2):

$$L = L_1 \quad (11)$$

$$= \frac{BP}{N(T^2 - 1)e_F + N(e_G + e_I)}. \quad (12)$$

As we can see from the above, the lifetime of the network increases proportionally to the available battery budget B and with the period between consecutive measurements P . A reasonable alternative problem specification would provide, instead of the total budget B , the uniform battery level b_u for each node. In that case, the lifetime of the network would be independent of the total number of nodes N . The only topology relevant quantity would be the number of tiers T . This makes sense, considering that if the number of tiers is kept constant, the number of nodes in each tier will increase proportionally, and, hence, the increased traffic flow from the higher tiers will be

distributed over a larger number of nodes in the lower tiers. We verify this assertion by numerical simulation in Section 4.

From the above, we also deduce a set of quantities that will be useful in later problems. It is obvious that when the lifetime of the first tier and, therefore, the whole network has expired, the nodes in other tiers have still not expended their battery levels. The power consumption for nodes at different tiers can be easily obtained from a consideration of (10) and (11). The hypothetical lifetime of tier i is L_i ; but, it is actually expending power for only the lifetime L of the network. After that, communication stops, and there is no more expenditure of power. Thus, the ratio of battery power b_{e_i} actually consumed by a node in tier i to the uniform battery level b_u that it started with is:

$$C_i = \frac{b_{e_i}}{b_u} = \frac{L_1}{L_i} = \frac{\frac{T^2-i^2}{2i-1}e_F + e_G + e_I}{(T^2-1)e_F + e_G + e_I}. \quad (13)$$

We call the ratios $\{C_i\}$ the *Ideal Allocation Ratios*, because ideally each node would be allocated only the amount of battery it would actually consume during the lifetime. The farther out a node, the lower the fraction C_i of its battery that it has consumed. The unconsumed energy is a wastage of the total energy budget. In all of the following problems, the broad design goal is to redistribute this energy budget non-uniformly so as to increase the lifetime of the whole network.

In considering these problems, it is important to keep in mind questions of scaling of any solution algorithm. Any such algorithm would be executed not by the nodes themselves but offline during network design. The number of nodes N is expected to be very high in these networks. Any algorithm requiring superpolynomial computation time in N is practically not useful. Even polynomial computation may be impractical unless of low order. The number of tiers T is related to N and the density of deployment. At its highest, it is of the order of the square root of N (when the spacing between nodes is roughly equal to the transmission range R_{TX}). Alternatively, we can think of T as the span of the overall sensing area in units of the transmission range. Clearly, T may also be significantly large and superpolynomial computations in T should be avoided. However, the number of distinct battery levels employed in the network (which we indicate by k below) is likely to be a quite small integer, because larger numbers will require more problematic deployment. Thus k will not usually pose a scalability problem.

In practice, for many applications, a designer would have several known battery capacities to choose from (e.g., AAA, AA, C and D). Then, the following problem can be formulated.

Problem 2 *Given k available battery levels $b_1 > b_2 > \dots > b_k > 0$, assign the battery level for each tier of nodes in a sensor network with parameters $N, R_M, R_{TX}, P, e_F, e_G$ and e_I , such that the total battery budget B is minimized subject to maximizing the lifetime L of the network.*

Two alternative flavors of the above problem can be articulated: (a) all nodes in any given tier are constrained to be assigned the same battery level, and (b) different nodes of the same tier can be assigned different battery levels. We refer to these two alternative versions of the problem as Problem 2A and Problem 2B respectively. Thus, we are required to obtain one unique battery level b_i for tier i in Problem 2A; every node in this tier should be assigned this battery level. For Problem 2B, we can provide more than one battery level for each tier, accompanied by the *proportions* of the total number of nodes in that tier that should be assigned each battery level. Typically, we would provide two battery levels to a tier to change over from some battery level b_{i_1} to the next lower one b_{i_2} in the middle of the tier. If we specified proportions f_1 and f_2 for the two levels (with $f_1 + f_2 = 1$), then the *effective battery level* of tier i would be given by $b_{i_1}f_1 + b_{i_2}f_2$. Intuitively, if the nodes with the two different battery levels are mixed randomly within a tier, the

tier as a whole will behave as if each node in it had the effective battery level. We have verified this intuition by numerical simulation. Below, we address Problem 2A first.

The maximum network lifetime is obtained when all nodes have the maximum battery level b_1 . Thus, the first tier of nodes should have batteries of capacity b_1 , and the maximum lifetime of the network will be

$$L = L_1(b_1) \tag{14}$$

where by $L_i(b)$ we denote the lifetime of tier i (given by (7)) if each node in tier i initially has a battery capacity b (i.e., the energy budget of tier i is given by $B_i = N_i b$).

The same maximum lifetime as in (14) may be obtained with a lower budget by assigning lower battery levels to nodes in the higher tiers (in tiers i with $1 < i < T$); however, if the next smallest capacity size is too small, then the lifetime of the network would be reduced because nodes in higher tiers will deplete their batteries *before* the nodes in the first tier.

For example, if $L_1(b_1) > L_2(b_2)$, the nodes in the second tier deplete their batteries faster than the nodes in the first tier, and the lifetime of the network is smaller than the maximum (14).

Given b_1 , the ideal level of battery that each node of tier i should have is obviously $C_i b_1$, where C_i are the Ideal Allocation Ratios given by (13), so that tier i will have exactly the same lifetime as tier 1. It remains to choose the smallest level from among the available battery levels that is at least this much. In short, for each tier i , in order to maximize the lifetime of the network and then conditionally minimize the total battery budget, we need to make sure that we assign the battery size b_{j_i} to tier i such that $b_{j_{i-1}} < C_i b_1 \leq b_{j_i}$.

Now, if we consider Problem 2B, we have the added flexibility of mixing the provided battery levels. The solution is along the same lines. First, we pick the highest battery level b_1 for tier 1 as before: this maximizes the network lifetime. Now we predict the ideal battery levels for each tier i as before using (13); but now, instead of picking the next highest available level from the available ones, we aim at attaining this ideal level exactly as the *effective battery level* as defined above by mixing the two provided battery levels in the appropriate ratio. In case the desired battery level is exactly equal to one of the provided battery levels, no mixing is needed. If the ideal level is less than the minimum provided battery level b_k , then that minimum level must be used in that tier instead; in this case some battery will indeed remain unconsumed at the end of the lifetime. ■

Note that in the above problem and later ones, we study the case where nodes with the different desirable levels are deployed in precise annular rings. This is of direct interest when the deployment is controlled (such as in a large bridge or other civil structure). When deployment is less controlled, such as by scattering from an aircraft, it is likely that some intermixing of nodes between different rings will occur. The extreme case of this is completely random deployment. We show later in this paper that in such a case lifetime will be drastically reduced. In all realistic cases of deployment, most nodes will have positions in the appropriate annular regions, but some randomness will be introduced; lifetime will be then somewhat reduced from that achieved in the ideal case, but the ideal lifetime is a good indicator of actual lifetime in such cases.

In the above problem, minimizing the battery budget is a secondary goal of the optimal design. It is also possible that the battery budget may be strictly a constraint for the design. If the total cost of the batteries used for a particular network is upper bounded by a total battery budget B and the battery capacities are fixed and given (as before), we can formulate the following problem:

Problem 3 *Given k available battery levels $b_1 > b_2 > \dots > b_k > 0$, and a total energy budget B , assign the battery level for each tier of nodes in a sensor network with parameters $N, R_M, R_{TX}, P, e_F, e_G$ and e_I , such that the **total lifetime L of the network is maximized.***

As with Problem 2, we can conceive of two alternate problems, Problems 3A and 3B, in which the nodes in any tier are constrained to have the same battery level, or mixing is allowed, respectively. We consider Problem 3A first.

Initially, assume that $L_i(b_j)$ is precomputed and available for each tier i and each battery level j . Set a “current level” p_i for each tier i , initializing to $p_1 = 1$ for the first tier, and p_i to the maximum value such that $L_i(b_{p_i}) \geq L_1(p_1)$. Initialize “current lifetime” L_c to $L_1(p_1)$.

Whenever the total battery $\sum_i N_i b_{p_i}$ becomes less than or equal to the battery budget B , the algorithm is terminated.

If the current total battery exceeds B , repeatedly perform the following. Find the tier i such that the difference $L_c - L_i(b_{p_i+1})$ is minimized. That is, we find the tier i such that assigning this tier the next lower battery level from the one it is currently assigned will result in the *minimum reduction of the lifetime*. Note that this minimum difference may well be zero at some iterations. Ties are broken arbitrarily. Increment the current level for tier i by 1, and update the current lifetime to the new lifetime of tier i . Recalculate the total battery used.

We claim that this algorithm will stop with an optimal solution. The initialization obviously does not rule out any optimal solutions since it merely results in the algorithm starting with the maximum lifetime unconstrained by total battery budget (i.e. the solution to Problem 2). In each following step, the battery level of some tier i is irreversibly decreased. To show that this choice does not exclude the optimal, we note that if this step is performed at all then the current total battery budget is still larger than the constraint. Then at least one tier of nodes must have a lower battery level in the optimal solution. By construction of this step, any such lower battery level at any tier would result in a final lifetime no higher than the lifetime $L_i(b_{p_i+1})$ that is achieved at the end of this step. If the lifetime is no higher, then nothing can be gained by leaving the battery level of tier i at b_{p_i} . In other words, reducing the battery level at this tier cannot prevent us from reaching the optimal.

Note that when the total battery first falls below the budget B , an optimal solution has been generated because an assignment is found that results in the maximum possible lifetime (as proven above) while the total battery used is below the budget. However, it may be the case that further battery level reductions can be made at some tiers without lowering the lifetimes any further. This would happen if we broke a tie in choosing the tier whose battery level is to be reduced. Practically, it might be preferable to obtain this reduction in total battery actually used in design. It is straightforward to continue the algorithm to reduce the battery levels as far as possible at other tiers, but this time accepting only zero decreases in lifetimes. Equivalently, (13) may be used to determine the battery levels for other tiers.

The complexity of this algorithm is low. Precomputing the lifetimes, the initialization step, and the main iteration are all of the order of $O(kT)$. A dynamic programming formulation to compute each lifetime as needed instead of precomputing them all would eliminate part of the computation but leaves the order unchanged; while such an algorithm may be preferable in actual practice we do not go into the details here.

Problem 3B turns out to be a modification of Problem 2B, where the total battery budget B is now a hard constraint. Accordingly, we take the following approach to solve it: first solve Problem 2B on the same parameters, ignoring the battery budget. If the total battery budget B_t of this solution does not exceed B , then this is the desired solution. If instead $B_t > B$, then obtain a new set of effective battery levels for each tier by scaling the battery levels by a factor of B/B_t : these are the new desired battery levels. A special case arises when some of these new desired battery levels are less than the minimum provided battery level b_k . In this case the battery levels for the inner tiers have to be reduced to allow the outer tiers to have battery level b_k , further reducing the lifetime. The procedure to do so is fairly straightforward and we do not detail it here.

If the available battery levels are not fixed, but no more than k distinct levels can be used, then the following problem arises. This problem is realistic because for some networks the designer may have the option of manufacturing custom batteries; or a wide selection may be available so that effectively any desired battery level may be picked. Yet for logistic reasons in deployment or other design policy, the number of distinct levels to be used in design may be limited. ■

Problem 4 *Given a total battery budget B and a maximum of k distinct battery levels, determine the optimal battery levels b_1, \dots, b_k , $b_1 > b_2 > \dots > b_k > 0$, and their assignment for each tier of nodes in a sensor network with parameters $N, R_M, R_{TX}, P, e_F, e_G$ and e_I , such that **the total lifetime L of the network is maximized.***

Once again, two alternative problems 4A and 4B can be posed.

If $k \geq T$, i.e. there are at least as many battery sizes as tiers in the network, the problem can easily be solved by setting $k = T$ and then obtaining the battery level b_i of each tier i in terms of the battery level b_1 to be assigned to the first tier by using the Ideal Allocation Ratios of (13). It remains to find b_1 , which can be obtained from the total battery budget using the following equation:

$$B = \sum_{i=1}^T b_i N_i = b_1 \sum_{i=1}^T C_i N_i. \quad (15)$$

Note that Problem 4B is indistinguishable from Problem 4A in this special case; since we can choose our own battery levels, mixing does not provide us any further flexibility.

However, if $k < T$, especially by a large margin, then the problem becomes one of integer optimization, and accordingly a hard one if T is large. We address this problem in a different way. In the next section, we present a continuous asymptotic form of the Ideal Allocation Ratios for the discrete case corresponding to (13). We then show how to use this asymptotic case as an approximation to obtain the optimal solution to Problem 4 under the idealizations of the asymptotic case. Finally, we show how this solution can be applied to the realistic case. ■

3.5.1 Continuous Asymptotic Case

We consider the asymptotic situation when the density of nodes per unit area remains the same, but the number of nodes per unit area increases sufficiently that the density can be considered to hold for any area, however small. Then, the traffic forwarded by a node in each unit of time will vary continuously with its radial distance r from the monitoring station. Initially, we develop the theory only considering the forwarding energy e_F and setting $e_G = e_I = 0$. Later, we show how to take these parameters into account.

For ease of discussion, we define:

τ The traffic generated by the nodes in a unit area, in unit time.

β The amount of energy required to transmit one unit of traffic, once.

n The (uniform) density of node placement over the sensing area.

Thus, $\tau\beta$ represents the energy required to forward, once, the traffic generated by the nodes in a unit area in unit time, and is equal to ne_F/P , by our earlier definitions. Then, the energy $e(r)$ expended by the nodes in a unit area around radius r in unit time is given by:

$$2\pi r R_{TX} e(r) = \pi(R_M^2 - r^2)\tau\beta \quad (16)$$

and, hence, the battery power $b(r)$ expended by each individual node in that area is given by:

$$b(r) = e(r)/n = \frac{(R_M^2 - r^2)\tau\beta}{2rR_{TX}n}. \quad (17)$$

The approximation in using the continuous model arises from the assumption that the above is valid for all values of r . In reality, this is valid for the radii $r = (2m + 1)R_{TX}/2$, that is the middle of tiers, where the nodes “around” radius r are correctly represented as being homogeneous.

Using the above, the energy E expended in the whole network in unit time can be easily obtained by integrating $e(r)$ over the entire area:

$$E = \frac{2}{3}\pi R_M^3 \tau\beta / R_{TX} \quad (18)$$

and multiplying E by a lifetime of L seconds gives the minimum amount of total energy budget that can achieve a total lifetime of L , when distributed according to (17), such that every node exhausts its battery at the same time. Naturally, this requires continuously varying battery levels.

Now, if we are realistically constrained to using only k distinct battery levels, as required by Problem 5, then the best lifetime will be obtained by approximating this continuously varying battery distribution as far as possible. The nodes in an annular area will have the same battery level, and, as before, annular areas closer to the center should have higher battery levels. Let all nodes from radius $\rho_0(=0)$ to ρ_1 have a battery level of b_1 , from ρ_1 to ρ_2 have a battery level of b_2 , and so on, until all nodes from radius ρ_{k-1} to $\rho_k(=R_M)$ have a battery level of b_k . Call these the “rings” $1 \dots k$. The problem is to determine $\rho_1 \dots \rho_{k-1}, b_1 \dots b_k$.

However, we can reduce the number of unknowns using the following observation. Consider the i -th ring. The lifetime of this ring is determined by the lifetime of the innermost tier in this ring, because inside each ring the batteries are uniform and the usual doughnut effect applies. For the longest lifetime of the network, each ring should have the same lifetime; that is, the tiers $t_i, i = 1 \dots k$ should have the same lifetime where t_i is the innermost tier of ring i . Clearly, $t_1 = 1$. The assumption made in the statement above is that each ring starts exactly at a tier boundary, so that “innermost tier” makes sense. This is in keeping with the approximation in the continuous model as noted above. Later in this section, we show how to reconcile this with a realistic solution of the problem. Now, we know that $b(r)$ as given by (17) correctly reflects the relationship between battery power and node position for a given lifetime *at the radii in the middle of tiers*. We can use this to eliminate the unknowns b_i above by expressing them in terms of ρ_i . Note that there are only $k - 1$ of the latter, and we shall obtain $k - 1$ equations to solve for them. Since we started with k unknowns b_i , we need one more equation to completely eliminate them. This last equation is obtained by using the total battery budget B which is given as part of the problem instance. We carry out this development below.

Briefly, we consider the battery power midway through the innermost tier of each ring. For ease of notation, we define the variables $r_i, i = 1 \dots k$, as $r_i = \rho_{i-1} + R_{TX}/2$, and $r_{k+1} = R_M + R_{TX}/2$. We already have $r_1 = R_{TX}/2$. For equal lifetimes of the tiers t_i , and, hence, the rings i , the battery levels b_i must be in the same ratio as $b(r_i)$. Letting L stand for the equal lifetime thus achieved, we can assert:

$$b_i = b(r_i)L = \frac{(R_M^2 - r_i^2)\tau\beta L}{2r_i R_{TX}n}. \quad (19)$$

The total battery in the network is given by $\sum_{i=1}^k b_i N_i$. With the help of the above equation and a little manipulation, we obtain:

$$B = \frac{\pi\tau\beta L}{2R_{TX}} S, \quad (20)$$

where

$$S = \sum_{i=1}^k \frac{R_M^2 - r_i^2}{r_i} (r_{i+1}^2 - R_{TX}r_{i+1} - r_i^2 + R_{TX}r_i). \quad (21)$$

Naturally, the lifetime L will be maximized for a given battery budget B if r_2, r_3, \dots, r_k are chosen so that the sum S is minimized. Clearly, we can obtain $k - 1$ equations by taking the partial derivative of S with respect to each of r_2, r_3, \dots, r_k , and setting each of the resulting expressions to zero. Simultaneously solving these equations allows us to determine the optimal values for the $\{r_i\}$.

It might appear that this gives us quite complicated equations, but, actually, the particular structure of S is conducive to this procedure. This is because r_i only appears in $i - 1$ -th and i -th terms of the sum S . For $k = 2$, there is only one equation which is straightforward to obtain. Denoting the i -th term of S by S_i , for larger values of k , the above observation allows us to obtain:

$$\frac{\partial S}{\partial r_i} = \frac{\partial S_{i-1}}{\partial r_i} + \frac{\partial S_i}{\partial r_i} \quad (22)$$

$$\begin{aligned} & 2R_M^2 r_i / r_{i-1} - R_M^2 R_{TX} / r_{i-1} - 2r_{i-1} r_i \\ = & + R_{TX} r_{i-1} - R_M^2 r_{i+1}^2 / r_i^2 + R_M^2 R_{TX} r_{i+1} / r_i^2 \\ & - R_M^2 - r_{i+1}^2 + R_{TX} r_{i+1} + 3r_i^2 - 2R_{TX} r_i. \end{aligned} \quad (23)$$

A closed form solution for the above system appears difficult to obtain and also unlikely to provide much insight. Fortunately, they are also unnecessary for practical purposes since they can be solved symbolically in MATLAB for given values of the parameters R_M and R_{TX} , providing numerical solutions for the problem instance defined by these parameter values. We have repeatedly verified, by direct numerical minimization of S , that such a solution, while taking only a fraction of the time direct numerical minimization requires, gives us the desired $\{r_i\}$ for the minimum value of S correctly. We use this method to obtain $\{r_i\}$ values in all numerical results presented in Section 4.

Substituting these $\{r_i\}$ in (20) gives us the maximum lifetime corresponding to a given total battery budget. From the point of view of the practical network designer, it remains to find the actual battery levels $\{b_i\}$ to use, and the realistic regions over which to deploy them. We can find $\{b_i\}$ by noting that there is only one unknown b_1 , since all other b_i can be expressed in terms of b_1 by using the proportions given by (17) and knowing $\{r_i\}$. Finally, summing $b_i N_i$ over the k rings and equating to B gives us b_1 , and, hence, all $\{b_i\}$.

We now show how to take the effect of nonzero e_G and e_I into account. Naturally, since the overall battery budget and the maximum obtainable life are related linearly in (20), depending only on the ratio of the sensor field radius and the transmission range, the $\{r_i\}$ obtained by the above method will still give us the optimal radii to change battery levels; but, now they cannot be directly used to find $\{b_i\}$, because each node also spends $(e_G + e_I)$ in each period. Thus, the battery budget for forwarding only has to be reduced. Let the optimal lifetime under these conditions be L_0 . Then, we have:

$$B = \frac{L_0}{P} N(e_g + e_I) + \frac{\pi \tau \beta L_0}{2R_{TX}} S_{\min} \quad (24)$$

$$= \frac{L_0}{P} N(e_g + e_I) + \frac{e_F N L_0}{2R_{TX} P R_M^2} S_{\min} \quad (25)$$

where S_{\min} is obtained by substituting the $\{r_i\}$ already obtained in the expression for S . The first term in the above represents the constant energy expenditure by each node, and the second term represents the different forwarding load as before. We can solve the above for L_0 . Then, the battery

budget for forwarding only can be obtained as $B^f = \pi\tau\beta L_0 S_{\min}/2R_{TX}$, from which we can obtain the battery component at each node at various rings $\{b_i^f\}$ by the procedure already detailed, finally obtaining $b_i = b_i^f + (L_0/P)(e_G + e_I)$, the actual design parameters desired.

A final step remains to convert the values obtained to actual design parameters. In general, the $\{r_i\}$ obtained may not result in rings that start at tier boundaries. As previously with the solutions of Problems 3 and 4, we can randomly mix the two required battery levels to obtain the effective desired battery level at such tiers where the $\{r_i\}$ obtained prescribe a changeover of battery levels.

4 Simulation Results

To validate the results presented in the previous section and to quantify the effect of various parameters, we decided to simulate the wireless sensor network in several scenarios. It quickly became apparent that any of the existing network simulators (ns-2, OPNET, Glomosim, OMNET++, Qualnet, etc.) are inadequate for simulating networks of hundreds of nodes with lifetimes on the order of several months (or even years). The common problem of these simulators is their relatively high level of detail at the physical and MAC layers. Therefore, we wrote a dedicated simulator, that abstracts many of the lower layer issues, to focus on the energy consumption problem.

4.1 Base Case

We will first study the effects of using multiple battery levels on the lifetime of the network for a base case corresponding to one set of parameters. In the next section, we will vary each of these parameters and comment on the effects on the network lifetime.

For the base case, we considered a rectangular grid in a circular area centered around the gateway. We assume a transmission radius that is slightly larger than three times the diagonal of the grid. Any two nodes at a distance less than the transmission radius can communicate with no errors, any nodes at a distance larger than the transmission radius cannot communicate. For the base case, we considered a network of $N = 905$ nodes, which corresponds to a ten-hop route for the peripheral nodes.

Since we are simulating a power aware network, we consider a very simple power aware routing protocol: one of the shortest paths is always used, and among the nodes advertising the same shortest paths, the one with the highest power is always used. Any further ties in choosing the next hop node is broken arbitrarily. The same routes are used until the batteries of a node are completely exhausted. New routes are formed after the death of each node.

We used power consumption values similar to the ones available for the Berkeley notes [13], with the transmission current $I_{Tx} = 17mA$, reception current $I_{Rx} = 10mA$ and idle current $I_{idle} = 10\mu A$. We assumed that each transmitter has only one transmission power level, and, thus, the transmission power does not depend on the distance between the source and the destination. A sensing current $I_{sense} = 5mA$ is drawn for $10ms$ for each collected sample.

Each node is powered by two high quality AA batteries ($2000mAh$). For the lifetime of the network, each node generates one 50ms sample (packet) every minute.

Figure 2 depicts (with a continuous line) the percentage of nodes that depleted their batteries and (with a dotted line) the percentage of nodes that can no longer reach the gateway, either because they have depleted their batteries, or are disconnected from the gateway. In the beginning, all the nodes that cannot longer reach the gateway depleted their batteries. After the batteries of the nodes in the first tier are depleted, there is a sudden increase in the nodes that cannot reach

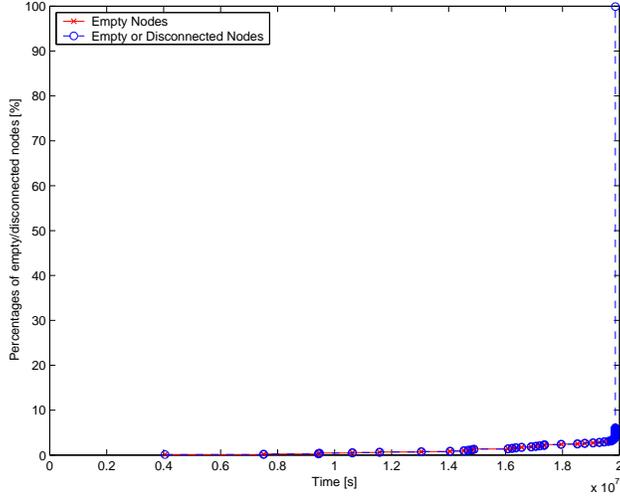


Figure 2: Evolution in time of the percentages of nodes with empty batteries (marked with an 'o') and the percentage of nodes that cannot reach the gateway (represented with an 'x') if a uniform battery is used.

the gateway. This is the moment the network becomes disconnected, and, effectively, the end of the lifetime of the network.

In Section 3, we assumed that the forwarding load is shared *equally* by all of the nodes in the same tier. A consequence of this assumption is that all of the nodes in the same tier should deplete their batteries at the same time. However, in reality (and in the simulations), in the beginning (when all nodes have the same amount of battery left) any of the nodes advertising the same shortest distance to the gateway can be used for forwarding. Therefore, some nodes, by chance, will forward more traffic than others. Depending on how many flows the most “unlucky” node (the one that forwards most flows in the network) has to forward, the time of the death of the first node can vary significantly from one simulation run to the other. Therefore, the time of the death of the first node is not a significant measure of the lifetime of the network. In contrast, the time of the death of the first tier of nodes (equal to the time the network becomes disconnected) has a very low variance, and matches very well the theoretical predictions.

4.1.1 The Effect of Multiple Battery Levels

If the same battery budget, as in the case of a uniform battery, is considered, but the batteries are redistributed on two levels, the corresponding network death/disconnection percentages are depicted in Fig. 3. The nodes with the largest batteries are distributed near the gateway. The battery levels and the number of nodes at each battery level are optimized using the continuous method presented in Section 3.5.1. For the base case and two battery levels, the optimum is achieved when 13% of the nodes (the nodes closest to the gateway) have 5718 mAh batteries, and 87% of the nodes have batteries with 1442 mAh capacity (the same total battery budget as for the single battery level case is used). Incidentally, the difference between the two battery levels is similar to the difference between the Mica2 and Mica2dot Berkeley nodes. Figure 3 already shows that the disconnection curve, while shifted toward longer lifetimes, has a similar characteristic as the curve in Fig. 2. Indeed, in both cases, the lifetime of the network is limited by the lack of connectivity, rather than by the battery depletion of the nodes.

Considering the results in Figs. 2 and 3, we define the lifetime of the network as the time it takes for 50% of the nodes to lose connectivity to the gateway. For most cases, the time it takes to

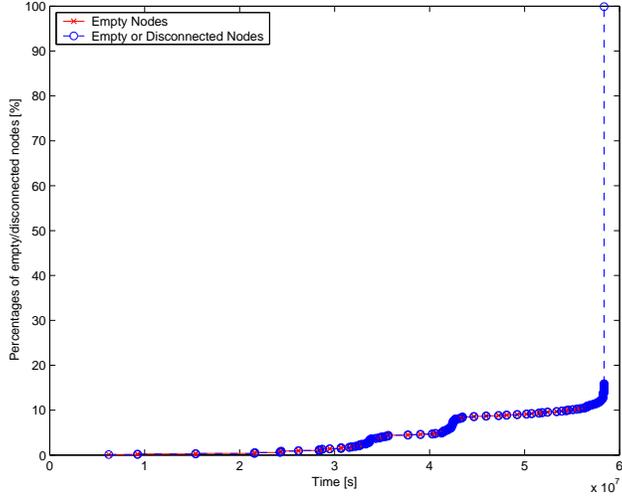


Figure 3: Evolution in time of the percentages of nodes with empty batteries (marked with an 'o') and the percentage of nodes that cannot reach the gateway (represented with an 'x') if two battery levels are used.

lose 50% of the nodes is equal to the time it takes to lose the first tier of nodes, which, in turn, is equal to the network disconnection time.

The lifetime of the network for the two battery level case (Fig. 3) is increased by comparison with the one battery level case (depicted in 2). The lifetime is further increased if more than two battery levels are considered, as seen in Figure 4. For each battery level, the optimization in Section 3.5.1 was performed, and optimum number of nodes and battery levels were used for the simulation. The same total energy budget was used for each simulation. The nodes of the same battery level were placed in concentric rings, with the nodes with the largest batteries closer to the gateway. Figure 4 shows that the increase in network lifetime tapers off with the increase in the

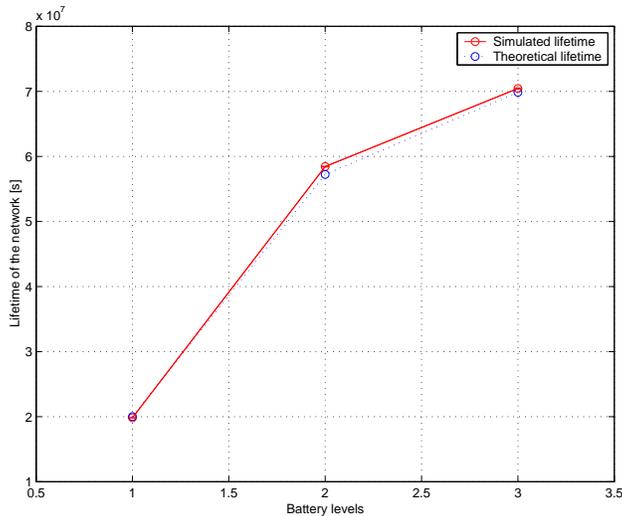


Figure 4: The increase in network lifetime as a function of the increase in the battery levels. The network has the same energy budget at each battery level.

number of battery levels. This indicates that most of the increase in the lifetime of the network can be achieved with a relatively small number of battery levels. To increase the confidence in

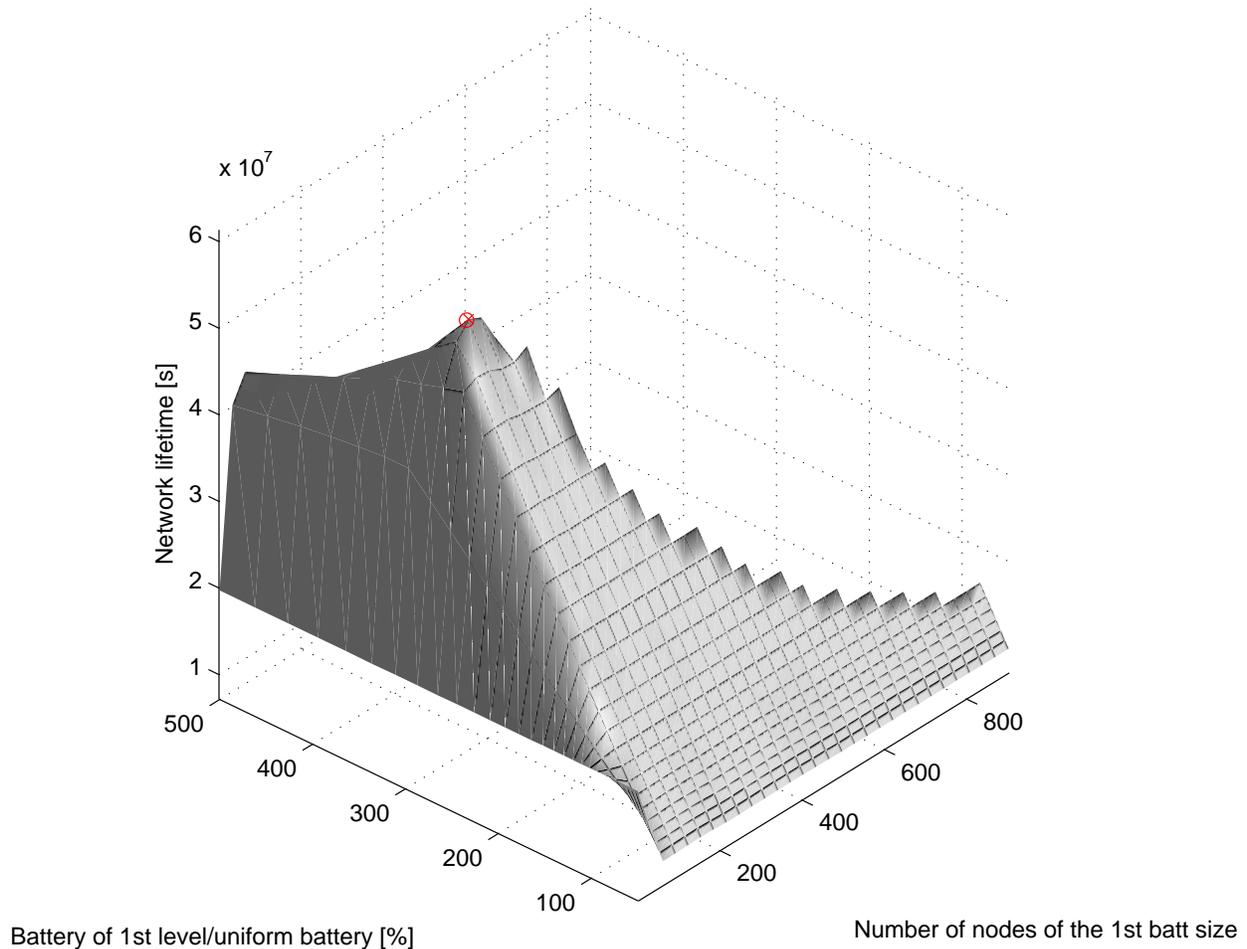


Figure 5: Experimental determination of the best placement and battery level of two levels of batteries.

the simulation results, for each point in Fig. 4 (and in all subsequent simulations), we performed 30 simulations with the same configuration. The mean and standard deviation are depicted. The standard deviation for Fig. 4 is so small that it cannot be observed on the graph. The theoretical network lifetime was computed as L_1 in (7).

To assess the effect of the idealizations in Section 3.5.1 on the optimality of the battery level choices, we performed a brute-force search for the optimal values for the case of two battery levels. Figure 5 depicts the lifetime of the network as a function of the battery level of the first level of nodes and the number of nodes in the first level. The maximum lifetime is marked with the symbol 'x'. The maximum, as predicted by theory in Section 3.5.1, is marked by the symbol 'o'. Once again, the continuous model matches the simulation results very well.

4.2 Departures from the Base Case

In this section, we consider the effect that various parameters have on the lifetime of the network. To separate the effects of each parameter we vary, we will keep the rest of the parameters in the base case constant.

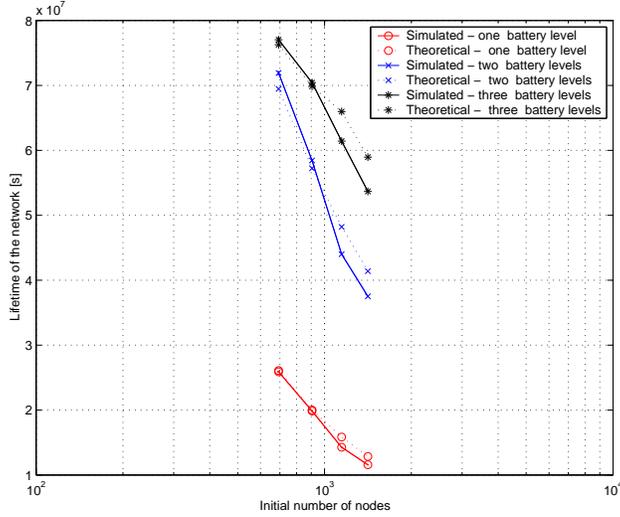


Figure 6: Dependency of the network lifetime on the initial number of nodes for three battery levels (constant density).

4.2.1 Dependency on Number of Nodes

The number of nodes for the base case $N = 905$ was chosen such that the network is large enough to satisfy the assumptions in Section 3.5.1 and yet run in a reasonable time (the entire batch of simulations takes a few days to complete). Figure 6 shows the dependency of the network lifetime on the initial number of nodes. The density of the network was kept constant, (the area of the network was proportionally increased with the number of nodes).

The lifetime of the network decreases with the increase in the initial number of nodes (hence the maximum number of hops). This is expected; as we already know a larger number of tiers results in more battery wastage. The small discrepancies between theory and simulation results are the effect of quantization (e.g., in theory 85.3 nodes should be in the first tier, but since simulated network topology is symmetric, only multiples of four nodes can be in the first tier).

4.2.2 Dependency on Node Density

In this section, we study the effect of increasing the density of the network (while keeping the network size constant) on the lifetime of the network. We expect the density to have no influence on the network lifetime as long as it remains uniform. The reason behind this prediction is that as long as the density is uniform, the number of nodes in each tier will increase in the same proportion, and, hence, the number of flows carried by each node does not change. The simulation results depicted in Fig. 7 confirm our intuition: with a very good precision, the lifetime of the network remains constant even if the density of the network doubles or triples.

4.2.3 Dependency on Grid Regularity

In the base case we assumed a regular placement of the nodes in a rectangular grid. The choice is meant to reduce the degree of randomness in the network, to be able to isolate the effects of other parameters. To evaluate the effect of small perturbations on the lifetime of the network, we gradually perturbed the grid, by an additive perturbation on each axis. The perturbation has a uniform distribution around zero.

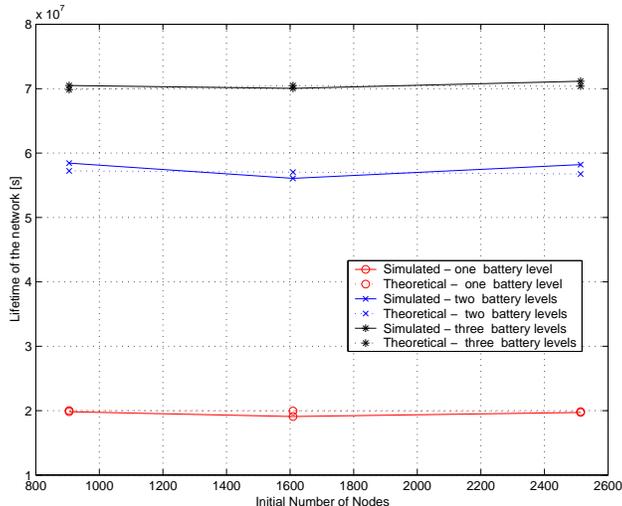


Figure 7: Dependency of the network lifetime on the network density for three battery levels (constant network area).

The average results, shown in Fig. 8 do not vary significantly with the increase in the grid perturbation. Even in the extreme case, where the limits of the perturbation are larger than the grid size, the lifetime of the network is similar to the one obtained for a regular grid. The standard deviation of the lifetime increases notably with the perturbation from the ideal grid. These simulation results considerably enhance the generality of the results we deduced for uniform grids and points toward their applicability in a scenario where the nodes are randomly placed, rather than in a grid.

4.2.4 Random vs Concentric Placement

Finally, we study the effect of placing the nodes with large batteries at random positions, rather than in concentric rings around the gateway (with larger nodes closer to the gateway). To this end, after computing the optimal number of nodes and the optimal battery levels using the optimization procedure described in Section 3.5.1, we placed the nodes with larger batteries in random positions in the regular grid, rather than closer to the gateway. The results, depicted in Fig. 9 confirm the expectations: the lifetime of the network is significantly less if the nodes are placed randomly, rather than in concentric rings.

5 Conclusion

In this paper, we identified a problem that we expect to occur in large, multihop wireless sensor networks. Under the assumption that larger traffic transmission load causes faster battery depletion, the nodes closer to a gateway will die before the nodes at the periphery of the network. The main disadvantage of the expiration of the nodes close to the gateway is that the network becomes disconnected while most of the nodes still have a considerable amount of battery left.

To alleviate this undesirable effect, we propose to use multiple levels of batteries, placing the nodes with larger batteries closer to the gateway. Under a total battery budget, we show a method to compute the optimal battery levels and number of nodes for each battery level. With this strategy, we show that the lifetime of the network can be significantly improved, even if a small

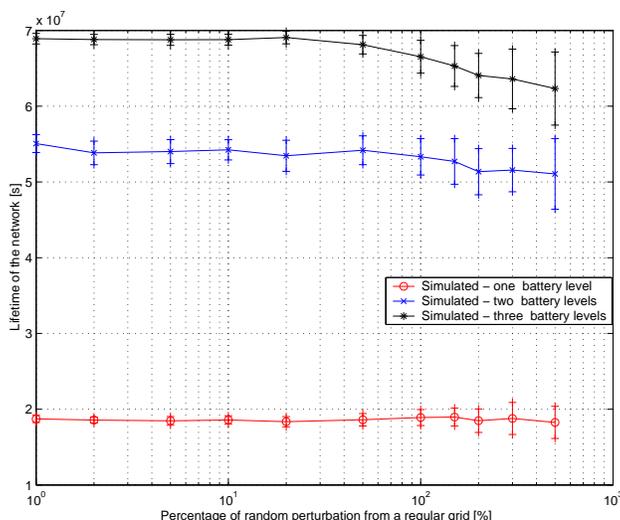


Figure 8: The effect of perturbations in the regularity of the grid. The perturbation on each axis is expressed as a percentage of the grid size.

number of battery levels is used.

References

- [1] I. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, “A survey on sensor networks,” *IEEE Communication Magazine*, vol. 40, no. 8, pp. 102–116, Aug. 2002.
- [2] —, “Wireless sensor networks: A survey,” *IEEE Computer*, vol. 38, no. 4, pp. 393–422, Mar. 2002.
- [3] G. Asada, T. Dong, F. Lin, G. Pottie, W. Kaiser, and H. Marcy, “Wireless integrated network sensors: Low power systems on a chip,” in *Proc. of the 24th European Solid-State Circuits Conference*, The Hague, Netherlands, 1998. [Online]. Available: citeseer.nj.nec.com/278712.html
- [4] R. Ramanathan and R. Hain, “Topology control of multihop wireless networks using transmit power adjustment,” in *INFOCOM (2)*, 2000, pp. 404–413. [Online]. Available: citeseer.nj.nec.com/ramanathan00topology.html
- [5] R. Wattenhofer, L. Li, P. Bahl, and Y.-M. Wang, “Distributed topology control for power efficient operation in multihop wireless ad hoc networks,” in *Proc. of INFOCOM*, 2001.
- [6] S. Narayanaswamy, V. Kawadia, R. S. Sreenivas, and P. R. Kumar, “Power control in ad-hoc networks: Theory, architecture, algorithm and implementation of the COMPOW protocol,” in *Proc. of European Wireless 2002. Next Generation Wireless Networks: Technologies, Protocols, Services and Applications*, Florence, Italy, Feb. 25–28 2002, pp. 156–162. [Online]. Available: black.csl.uiuc.edu/~prkumar/ps_files/compow_ewc_2002.pdf
- [7] W. Ye, J. Heidemann, and D. Estrin, “An energy-efficient MAC protocol for wireless sensor networks,” in *Proceedings of the IEEE Infocom*, USC/Information Sciences Institute. New York, NY, USA: IEEE, June 2002, pp. 1567–1576. [Online]. Available: <http://www.isi.edu/johnh/PAPERS/Ye02a.html>
- [8] S. Singh and C. Raghavendra, “Power efficient MAC protocol for multihop radio networks,” in *The Ninth IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, 1998, pp. 153–157.
- [9] R. Rozovsky and P. R. Kumar, “SEEDEX: A MAC protocol for ad hoc networks,” in *Proc. of The ACM Symposium on Mobile Ad Hoc Networking & Computing, (MobiHoc)*, Long Beach, CA, Oct. 2001, pp. 67–75.
- [10] Y.-C. Tseng, C.-S. Hsu, and T.-Y. Hsieh, “Power-saving protocols for IEEE 802.11-based multi-hop ad hoc networks,” in *Proc. of INFOCOM*, vol. 1, 2002, pp. 200–209.
- [11] M. Stemm and R. H. Katz, “Measuring and reducing energy consumption of network interfaces in hand-held devices,” *IEICE Transactions on Communications*, vol. E80-B, no. 8, pp. 1125–1131, 1997.
- [12] L. M. Feeney and M. Nilsson, “Investigating the energy consumption of a wireless network interface in an ad hoc networking environment,” in *Proc. of IEEE INFOCOM*, 2001, pp. 1548–1557.

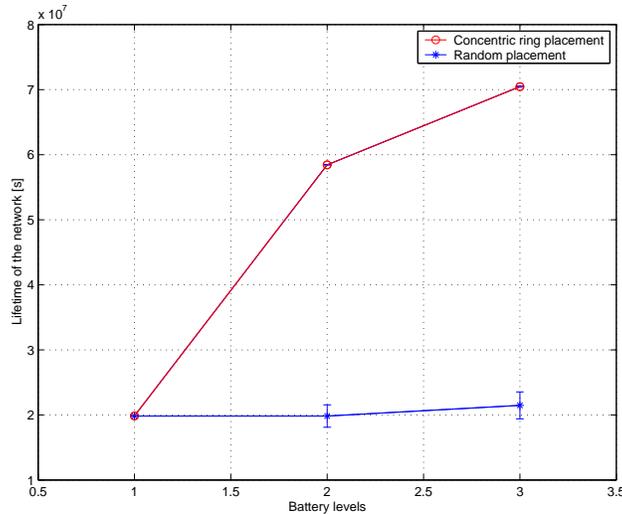


Figure 9: The effect of the placement of the nodes with large batteries at random positions rather than in concentric rings centered at the gateway.

- [13] D. E. Culler, J. Hill, P. Buonadonna, R. Szewczyk, and A. Woo, “A network-centric approach to embedded software for tiny devices,” in *Proc. of the First International Workshop on Embedded Software (EMSOFT)*, Oct. 2001.
- [14] S. Singh and C. Raghavendra, “PAMAS: Power aware multi-access protocol with signalling for ad hoc networks,” *ACM Computer Communications Review*, 1999. [Online]. Available: citeseer.nj.nec.com/460902.html
- [15] E.-S. Jung and N. Vaidya, “A power saving mac protocol for wireless networks,” Technical Report, UIUC, July 2002.
- [16] S. Singh, M. Woo, and C. S. Raghavendra, “Power-aware routing in mobile ad hoc networks,” in *Mobile Computing and Networking*, 1998, pp. 181–190. [Online]. Available: citeseer.nj.nec.com/singh98poweraware.html
- [17] M. W. Subbarao, “Dynamic power-conscious routing for MANETS: an initial approach,” in *IEEE Vehicular Technology Conference*, 1999, pp. 1232–1237.
- [18] T. A. ElBatt, S. V. Krishnamurthy, D. Connors, and S. Dao, “Power management for throughput enhancement in wireless ad-hoc networks,” in *IEEE International Conference on Communications*, 2000, pp. 1506–1513.
- [19] J. P. Monks, V. Bhargavan, and W.-M. W. Hwu, “A power controlled multiple access protocol for wireless packet networks,” in *Proceedings of INFOCOM*, 2001, pp. 219–228.
- [20] J.-C. Cano and P. Manzoni, “Evaluating the energy-consumption reduction in a manet by dynamically switching-off network interfaces,” in *Proc. of the Sixth IEEE Symposium on Computers and Communications*, 2001, pp. 186–191.
- [21] R. Kravets, K. Schwan, and K. Calvert, “Power-aware communication for mobile computers,” in *Proc. of the International Workshop on Mobile Multimedia Communications(MoMuc’99)*, 1999. [Online]. Available: citeseer.nj.nec.com/kravets99poweraware.html
- [22] B. Chen, K. Jamieson, H. Balakrishnan, and R. Morris, “SPAN: An energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks,” in *MOBICOM 2001*, July 2001.
- [23] R. Min, M. Bhardwaj, N. Ickes, A. Wang, and A. Chandrakasan, “The hardware and the network: Total-system strategies for power aware wireless microsensors,” in *Proc. of the IEEE CAS Workshop on Wireless Communications and Networking*, Pasadena, CA, USA, Sep. 2002. [Online]. Available: <http://www.mit.edu/~rmin/research/min-cas02.pdf>
- [24] E. Shih, S.-H. Cho, N. Ickes, R. Min, A. Sinha, A. Wang, and A. Chandrakasan, “Physical layer driven protocol and algorithm design for energy-efficient wireless sensor networks,” in *Proc of Mobicom 2001*, 2001.
- [25] A. Cerpa and D. Estrin, “ASCENT: Adaptive self-configuring sensor networks topologies,” in *Proc. of the Twenty First International Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2002)*, New York, NY, USA, June 2002. [Online]. Available: citeseer.ist.psu.edu/559481.html

- [26] W. Heinzelman, J. Kulik, and H. Balakrishnan, "Adaptive protocols for information dissemination in wireless sensor networks," in *Proc. MOBICOM*, 1999, pp. 174–185. [Online]. Available: cite-seer.ist.psu.edu/kulik99adaptive.html
- [27] W. R. Heinzelman, A. Chandrakasan, and H. Balakrishnan, "Energy-efficient communication protocol for wireless microsensor networks," in *HICSS*, 2000. [Online]. Available: cite-seer.ist.psu.edu/rabinerheinzelman00energyefficient.html
- [28] Y. Xu, J. Heidemann, and D. Estrin, "Geography-informed energy conservation for ad hoc routing," in *Proc. of ACM MobiCom 2001*, July 2001.
- [29] S. R. Madden, M. J. Franklin, J. M. Hellerstein, and W. Hong, "The design of an acquisitional query processor for sensor networks," in *In Proc. of (SIGMOD'03)*, San Diego, CA, June 2003. [Online]. Available: <http://www.cs.berkeley.edu/madden/acqp.pdf>
- [30] M. L. Sichitiu, "Cross-layer scheduling for power efficiency in wireless sensor networks," in *Proc. of Infocom 2004*, Hong Kong, PRC, Mar. 2004.
- [31] A. Woo and D. E. Culler, "A transmission control scheme for media access in sensor networks," in *Proc. ACM MOBICOM*, July 2001.