

SEP: A Stable Election Protocol for clustered heterogeneous wireless sensor networks *

Georgios Smaragdakis Ibrahim Matta Azer Bestavros
Computer Science Department
Boston University
Boston, MA 02215, USA

{gsmaragd, matta, best}@cs.bu.edu
Technical Report BUCS-TR-2004-022

ABSTRACT

We study the impact of heterogeneity of nodes, in terms of their energy, in wireless sensor networks that are hierarchically clustered. In these networks some of the nodes become cluster heads, aggregate the data of their cluster members and transmit it to the sink. We assume that a percentage of the population of sensor nodes is equipped with additional energy resources—this is a source of heterogeneity which may result from the initial setting or as the operation of the network evolves. We also assume that the sensors are randomly (uniformly) distributed and are not mobile, the coordinates of the sink and the dimensions of the sensor field are known. We show that the behavior of such sensor networks becomes very unstable once the first node dies, especially in the presence of node heterogeneity. Classical clustering protocols assume that all the nodes are equipped with the same amount of energy and as a result, they can not take full advantage of the presence of node heterogeneity. We propose SEP, a heterogeneous-aware protocol to prolong the time interval before the death of the first node (we refer to as *stability period*), which is crucial for many applications where the feedback from the sensor network must be reliable. SEP is based on weighted election probabilities of each node to become cluster head according to the remaining energy in each node. We show by simulation that SEP always prolongs the stability period compared to (and that the average throughput is greater than) the one obtained using current clustering protocols. We conclude by studying the sensitivity of our SEP protocol to heterogeneity parameters capturing energy imbalance in the network. We found that SEP yields longer stability region for higher values of extra energy brought by more powerful nodes.

1. INTRODUCTION

Motivation: Wireless Sensor Networks are networks of tiny, battery powered sensor nodes with limited on-board processing, storage and radio capabilities [1]. Nodes sense and send their reports toward a processing center which is called “sink.” Designing protocols and applications for such networks has to be energy aware in order to prolong the lifetime of the network, because the replacement of the embedded batteries is a very difficult process once these nodes have been installed. Classical approaches like Direct Trans-

mission and Minimum Transmission Energy [9] do not guarantee well balanced distribution of the energy load among nodes of the sensor network. Using Direct Transmission (DT), sensor nodes transmit directly to the sink, as a result nodes that are far away from the sink would die first [6]. On the other hand, using Minimum Transmission Energy (MTE), data is routed over minimum-cost routes, where cost reflects the transmission power expended. Under MTE, nodes that are near the sink act as relays with higher probability than nodes that are far from the sink. These former nodes tend to die fast. Under both DT and MTE, a part of the field will not be monitored for a significant part of the lifetime of the network, and as a result the sensing process of the field will be biased. A solution proposed in [7], called LEACH, guarantees that the energy load is well distributed by dynamically created clusters, using cluster heads dynamically elected according to a priori optimal probability. Cluster heads aggregate reports from their cluster members before forwarding them to the sink. By rotating the cluster-head role uniformly among all nodes, each node tends to expend the same energy over time.

Most of the analytical results for LEACH-type schemes are obtained assuming that the nodes of the sensor network are equipped with the same amount of energy—this is the case of *homogeneous* sensor networks. In this paper we study the impact of heterogeneity in terms of node energy. We assume that a percentage of the node population is equipped with more energy than the rest of the nodes in the same network—this is the case of *heterogeneous* sensor networks. We are motivated by the fact that there are a lot of applications that would highly benefit from understanding the impact of such heterogeneity. One of these applications could be the re-energization of sensor networks. As the lifetime of sensor networks is limited there is a need to re-energize the sensor network by adding more nodes. These nodes will be equipped with more energy than the nodes that are already in use, which creates heterogeneity in terms of node energy. Note that due to practical/cost constraints it is not always possible to satisfy the constraints for optimal distribution between different types of nodes as proposed in [8].

There are also applications where the spatial density of sensors is a constraint. Assuming that with the current technology the cost of a sensor is tens of times greater than the cost of embedded batteries, it will be valuable to examine whether the lifetime of the network could be increased by simply distributing extra energy to some existing nodes without introducing new nodes.¹

*This work was supported in part by NSF grants ANI-0095988, ANI-9986397, EIA-0202067, and ITR ANI-0205294.

¹We also study the case of uniformly distributing such extra energy over all nodes. In practice, however, it maybe difficult to achieve

Perhaps the most important issue is that heterogeneity of nodes, in terms of their energy, is simply a result of the network operation as it evolves. For example, nodes could, over time, expend different amounts of energy due to the radio communication characteristics, random events such as short-term link failures or morphological characteristics of the field (*e.g.* uneven terrain).

Our Contribution: In this paper we assume that the sink is not energy limited (at least in comparison with the energy of other sensor nodes) and that the coordinates of the sink and the dimensions of the field are known. We also assume that the nodes are uniformly distributed over the field and they are not mobile. Under this model, we propose a new protocol, we call SEP, for electing cluster heads in a distributed fashion in two-level hierarchical wireless sensor networks. Unlike prior work (reviewed throughout the paper and in Section 7), SEP is heterogeneous-aware, in the sense that election probabilities are weighted by the initial energy of a node relative to that of other nodes in the network. This prolongs the time interval before the death of the first node (we refer to as *stability period*), which is crucial for many applications where the feedback from the sensor network must be reliable. We show by simulation that SEP provides longer stability period and higher average throughput than current clustering heterogeneous-oblivious protocols. We also study the sensitivity of our SEP protocol to heterogeneity parameters capturing energy imbalance in the network. We show that SEP is more resilient than LEACH in judiciously consuming the extra energy of advanced nodes—SEP yields longer stability region for higher values of extra energy.

Paper Organization: The rest of the paper is organized as follows. Section 2 provides the model of our setting. Section 3 defines our performance measures. In Section 4 we address the problem of heterogeneity in clustered wireless sensor networks, and in Section 5 we provide our solution to the problem. Section 6 presents simulation results. We review related work in Section 7. Section 8 concludes with directions for future work.

2. HETEROGENEOUS WSN MODEL

In this section we describe our model of a wireless sensor network with nodes heterogeneous in their initial amount of energy. We particularly present the setting, the energy model, and how the optimal number of clusters can be computed.

Let us assume the case where a percentage of the population of sensor nodes is equipped with more energy resources than the rest of the nodes. Let m be the fraction of the total number of nodes n , which are equipped with α times more energy than the others. We refer to these powerful nodes as *advanced* nodes, and the rest $(1 - m) \times n$ as *normal* nodes. We assume that all nodes are distributed uniformly over the sensor field.

2.1 Clustering Hierarchy

We consider a sensor network that is hierarchically clustered. The LEACH (Low Energy Adaptive Clustering Hierarchy) protocol [6] maintains such clustering hierarchy. In LEACH, the clusters are re-established in each “round.” New cluster heads are elected in each round and as a result the load is well distributed and balanced among the nodes of the network. Moreover each node transmits to the closest cluster head so as to split the communication cost to the sink (which is tens of times greater than the processing and operation cost). Only the cluster head has to report to the sink and may expend a large amount of energy, but this happens periodically for

such uniform distribution because extra energy could be expressed only in terms of discrete battery units. Even if this is possible, we show in this paper that such fair distribution of extra energy is not always beneficial.

each node. In LEACH there is an optimal percentage p_{opt} (determined a priori) of nodes that has to become cluster heads in each round assuming uniform distribution of nodes in space [2, 3, 6, 7].

If the nodes are *homogeneous*, which means that all the nodes in the field have the same initial energy, the LEACH protocol guarantees that everyone of them will become a cluster head exactly once every $\frac{1}{p_{opt}}$ rounds. Throughout this paper we refer to this number of rounds, $\frac{1}{p_{opt}}$, as *epoch* of the clustered sensor network.

Initially each node can become a cluster head with a probability p_{opt} . On average, $n \times p_{opt}$ nodes must become cluster heads per round per epoch. Nodes that are elected to be cluster heads in the current round can no longer become cluster heads in the same epoch. The non-elected nodes belong to the set G and in order to maintain a steady number of cluster heads per round, the probability of nodes $\in G$ to become a cluster head increases after each round in the same epoch. The decision is made at the beginning of each round by each node $s \in G$ independently choosing a random number in $[0, 1]$. If the random number is less than a threshold $T(s)$ then the node becomes a cluster head in the current round. The threshold is set as:

$$T(s) = \begin{cases} \frac{p_{opt}}{1 - p_{opt} \cdot (r \bmod \frac{1}{p_{opt}})} & \text{if } s \in G \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where r is the current round number. The election probability of nodes $\in G$ to become cluster heads increases in each round in the same epoch and becomes equal to 1 in the last round of the epoch. Note that by round we define a time interval where all clusters members have to transmit to the cluster head once. We show in this paper how the election process of cluster heads should be adapted appropriately to deal with *heterogeneous* nodes, which means that *not* all the nodes in the field have the same initial energy.

2.2 Optimal Clustering

Previous work have studied either by simulation [6, 7] or analytically [2, 3] the optimal probability of a node being elected as a cluster head as a function of spatial density when nodes are uniformly distributed over the sensor field. This clustering is optimal in the sense that energy consumption is well distributed over all sensors and the total energy consumption is minimum. Such optimal clustering highly depends on the energy model we use. For the purpose of this study we use similar energy model and analysis as proposed in [7].




Figure 1: Radio Energy Dissipation Model.

According to the radio energy dissipation model illustrated in Figure 1, in order to achieve an acceptable Signal-to-Noise Ratio (SNR) in transmitting an L -bit message over a distance d , the energy expended by the radio is given by:

$$E_{Tx}(l, d) = \begin{cases} L \cdot E_{elec} + L \cdot \epsilon_{fs} \cdot d^2 & \text{if } d < d_0 \\ L \cdot E_{elec} + L \cdot \epsilon_{mp} \cdot d^4 & \text{if } d \geq d_0 \end{cases}$$

where E_{elec} is the energy dissipated per bit to run the transmitter or the receiver circuit, ϵ_{fs} and ϵ_{mp} depend on the transmitter amplifier model we use, and d is the distance between the sender and the




Figure 2: (top) Optimal number of clusters; and (bottom) Optimal probability of a node to become a cluster head, as a function of number of nodes in a $100m \times 100m$ field where the sink is located in the center.

receiver. By equating the two expressions at $d = d_0$, we have $d_0 = \sqrt{\frac{\epsilon_{fs}}{\epsilon_{mp}}}$. To receive an L -bit message the radio expends $E_{Rx} = L \cdot E_{elec}$.

Assume an area $A = M \times M$ square meters, and n the number of nodes that are uniformly distributed over that area. For simplicity, assume the sink is located in the center of the field, and that the maximum distance of any node to the sink is $\leq d_0$. Thus, the energy dissipated in the cluster head node during a round is given by the following formula:

$E_{CH} = L \cdot E_{elec} \left(\frac{n}{k} - 1\right) + L \cdot E_{DA} \frac{n}{k} + L \cdot E_{elec} + L \cdot \epsilon_{fs} d_{toBS}^2$
 where k is the number of clusters, E_{DA} is the processing (data aggregation) cost of a bit per signal, and d_{toBS} is the distance between the cluster head and the sink. The energy used in a non-cluster head node is equal to:

$$E_{nonCH} = L \cdot E_{elec} + L \cdot \epsilon_{fs} \cdot d_{toCH}^2$$

where d_{toCH} is the distance between a cluster member and its cluster head. Assuming that the nodes are uniformly distributed, it can be shown that:

$$E[d_{toCH}^2] = \iint (x^2 + y^2) \rho(x, y) dx dy = \frac{M^2}{2 \cdot \pi \cdot k}$$

where $\rho(x, y)$ is the node distribution.

The energy dissipated in a cluster per round is the following:

$$E_{cluster} \approx E_{CH} + \frac{n}{k} E_{nonCH}$$

The total energy dissipated in the network is equal to:

$$E_{tot} = L \cdot \left(2nE_{elec} + nE_{DA} + \epsilon_{fs} \left(k \cdot d_{toBS}^2 + n \frac{M^2}{2 \cdot \pi \cdot k} \right) \right)$$

By differentiating E_{tot} with respect to k and equating to zero, the optimal number of constructed clusters can be found:²

$$k_{opt} = \sqrt{\frac{n}{2\pi}} \frac{M}{d_{toBS}} = \sqrt{\frac{n}{2\pi}} \frac{2}{0.765} \quad (2)$$

because the average distance from a cluster head to the sink is given by [3]:

$$E[d_{toBS}] = \int_A \sqrt{x^2 + y^2} \frac{1}{A} dA = 0.765 \frac{M}{2}$$

The optimal probability of a node to become a cluster head, p_{opt} , can be computed as follows:

$$p_{opt} = \frac{k_{opt}}{n} \quad (3)$$

Figure 2 shows the values of k_{opt} and p_{opt} as a function of the number of nodes in a $100m \times 100m$ field where the sink is located in the center. The optimal construction of clusters (which is equivalent to the setting of the optimal probability for a node to become a cluster head) is very important. In [6], the authors showed that if the clusters are not constructed in an optimal way, the total consumed energy of the sensor network per round is increased exponentially either when the number of clusters that are created is greater or especially when the number of the constructed clusters is less than the optimal number of clusters. Our simulation results confirm this observation in our case where the sink is located in the center of the sensor field.

3. PERFORMANCE MEASURES

We define here the measures we use in this paper to evaluate the performance of clustering protocols.

- *Stability Period*: is the time interval from the start of network operation until the death of the first sensor node. We also refer to this period as “stable region.”
- *Instability Period*: is the time interval from the death of the first node until the death of the last sensor node. We also refer to this period as “unstable region.”
- *Network lifetime*: is the time interval from the start of operation (of the sensor network) until the death of the last alive node.
- *Number of cluster heads per round*: This instantaneous measure reflects the number of nodes which would send directly to the sink information aggregated from their cluster members.
- *Number of alive (total, advanced and normal) nodes per round*: This instantaneous measure reflects the total number of nodes and that of each type that have not yet expended all of their energy.

²It is interesting to notice that the optimal number of clusters is independent of the dimensions of the field and only depends on the number of nodes n .




Figure 3: (top) A wireless sensor network; (middle) An instance of the network where all the nodes are alive; (bottom) An instance of the network where some nodes are dead.

- *Throughput:* We measure the total rate of data sent over the network, the rate of data sent from cluster heads to the sink as well as the rate of data sent from the nodes to their cluster heads.

Clearly, the larger the stable region and the smaller the unstable region are, the better the reliability of the clustering process of the sensor network is. On the other hand, there is a tradeoff between

reliability and the lifetime of the system. Until the death of the last node we can still have some feedback about the sensor field even though this feedback may not be reliable. The unreliability of the feedback stems from the fact that there is no guarantee that there is at least one cluster head per round during the last rounds of the operation. In our model, the absence of a cluster head prevents any reporting about that cluster to the sink. The throughput measure captures the rate of such data reporting to the sink.

4. HETEROGENEOUS-OBLIVIOUS PROTOCOLS

The original version of LEACH does not take into consideration the heterogeneity of nodes in terms of their initial energy, and as a result the consumption of energy resources of the sensor network is not optimized. The reason is that LEACH depends only on the spatial density of the sensor network.

Using LEACH in the presence of heterogeneity, and assuming both normal and advanced nodes are uniformly distributed in space, we expect that the first node dies on average in a round that is close to the round where the first node dies in the homogeneous case wherein each node is equipped with the same energy as that of a normal node in the heterogeneous case. Furthermore, we expect the first dead node to be a normal node. We also expect that in the following rounds the probability of a normal node to die is greater than the probability of an advanced node to die. During the last rounds only advanced nodes are alive. Our expectations are confirmed by simulation results in Section 6. We next demonstrate how such heterogeneous-oblivious clustering protocol fails to maintain the stability of the system, especially when nodes are heterogeneous. This motivates our proposed SEP protocol presented in Section 5.

4.1 Instability of Heterogeneous-oblivious Protocols

In this section we discuss the instability of heterogeneous-oblivious protocols, such as LEACH, once some nodes die. In this case, the process of optimal construction of clusters fails since the spatial density deviates from the assumed uniform distribution of nodes over the sensor field.

Let us assume a heterogeneous ($m = 0.2, \alpha = 1$) sensor network in a $100m \times 100m$ sensor field, as shown in Figure 3(top). For this setting we can compute from Equation (2) the optimal number of clusters per round, $k_{opt} = 10$. We denote with \circ a normal node, with $+$ an advanced node, with \cdot a dead node, with $*$ a cluster head and with \times the sink. As long as all the nodes are alive, the nodes that are included in the same Voronoi cell will report to the cluster head of this cell; see Figure 3(middle).

At some point the first node dies; see Figure 3(bottom). After that point the population of sensors decreases as nodes die randomly. The population reduction introduces instability in the sensor network and the cluster head election process becomes unreliable. This is because the value of p_{opt} is optimal only when the population of the network is constant and equal to the initial population (n). When the population of the nodes starts decreasing the number of elected cluster heads per round is very unstable (lower than intended) and as a result there is no guarantee that a constant number of cluster heads (equal to $n \times p_{opt}$) will be elected per round per epoch. Moreover there are less alive nodes so the sampling (sensing) of the field is over less nodes than intended to be.

The only guarantee is that there will be at least one cluster head per epoch (cf. Equation 1). As a result at least in one round per epoch all alive nodes will report to the sink. The impact and qual-

ity of these reports highly depends on the application. For some applications even this minimal reporting is a valuable feedback, for others it is not. Clearly minimal reporting translates to significant under-utilization of the resources and the bandwidth of the application.

LEACH guarantees that in the homogeneous case the unstable region will be short. After the death of the first node, all the remaining nodes are expected to die on average within a small number of rounds as a consequence of the uniformly remaining energy due to the well distributed energy consumption. Even when the system operates in the unstable region, if the spatial density of the sensor network is large, the probability that a large number of nodes be elected as cluster heads is significant for a significant part of the unstable region (as long as the population of the nodes has not been decreased significantly). In this case, even though our system is unstable in this region, we still have a relatively reliable clustering (sensing) process. The same can be noticed even if the spatial density is low but the p_{opt} is large. On the other hand LEACH in the presence of node heterogeneity yields a large unstable region. The reason is that all advanced nodes are equipped with almost the same energy but, the cluster head election process is unstable and as a result most of the time these nodes are idle, as there is no cluster head to transmit.

In the next section, we introduce our new heterogeneous-aware SEP protocol whose goal is to increase the stable region and as a result decrease the unstable region and improve the quality of the feedback of wireless clustered sensor networks, in the presence of heterogeneous nodes.

5. OUR SEP PROTOCOL

In this section we describe SEP, which improves the stable region of the clustering hierarchy process using the characteristic parameters of heterogeneity, namely the fraction of advanced nodes (m) and the additional energy factor between advanced and normal nodes (α).

In order to prolong the stable region, SEP attempts to maintain the constraint of well balanced energy consumption. Intuitively, advanced nodes have to become cluster heads more often than the normal nodes, which is equivalent to a fairness constraint on energy consumption. Note that the new heterogeneous setting (with advanced and normal nodes) has no effect on the spatial density of the network so the apriori setting of p_{opt} , from Equation (3), does not change. On the other hand, the total energy of the system changes. Suppose that E_o is the initial energy of each normal sensor. The energy of each advanced node will be $E_o \cdot (1 + \alpha)$. The total energy of the new heterogeneous setting is equal to:

$$n \cdot (1 - m) \cdot E_o + n \cdot m \cdot E_o \cdot (1 + \alpha) = n \cdot E_o \cdot (1 + \alpha \cdot m)$$

So, the total energy of the system is increased by $1 + \alpha \cdot m$ times. The first improvement to the existing LEACH is to increase the epoch of the sensor network in proportion to the energy increment. In order to optimize the stable region of the system, the new epoch must become equal to $\frac{1}{p_{opt}} \cdot (1 + \alpha \cdot m)$ because the system has $\alpha \cdot m$ times more energy and virtually $\alpha \cdot m$ more nodes (with the same energy as the normal nodes).

We can now increase the stable region of the sensor network by $1 + \alpha \cdot m$ times, if (i) each normal node becomes a cluster head once every $\frac{1}{p_{opt}} \cdot (1 + \alpha \cdot m)$ rounds per epoch; (ii) each advanced node becomes a cluster head exactly $1 + \alpha$ times every $\frac{1}{p_{opt}} \cdot (1 + \alpha \cdot m)$ rounds per epoch; and (iii) the average number of cluster heads per round per epoch is equal to $n \times p_{opt}$ (the spatial density does not change). Constraint (ii) is very strict—If at the end of each epoch




Figure 4: Performance of naive solution.

the number of times that an advanced sensor has become a cluster head is not equal to $1 + \alpha$ then the energy is not well distributed and the average number of cluster heads per round per epoch will be less than $n \times p_{opt}$. This problem can be reduced to a problem of optimal threshold $T(s)$ setting (cf. Equation 1), with the constraint that each node has to become a cluster head as many times as its initial energy divided by the energy of a normal node.

5.1 The Problem of Maintaining Well Distributed Energy Consumption Constraints in the Stable Period

If the same threshold is set for both normal and advanced nodes with the difference that each normal node $\in G$ becomes a cluster head once every $\frac{1}{p_{opt}} \cdot (1 + \alpha \cdot m)$ rounds per epoch, and each advanced node $\in G$ becomes a cluster head $1 + \alpha$ times every $\frac{1}{p_{opt}} \cdot (1 + \alpha \cdot m)$ rounds per epoch, then there is no guarantee that the number of cluster heads per round per epoch will be $n \times p_{opt}$. The reason is that there is a significant number of cases where this number can not be maintained per round per epoch with probability 1. A worst-case scenario could be the following. Suppose that all normal nodes become cluster heads once within the first $\frac{1}{p_{opt}} \cdot (1 - m)$ rounds of the epoch. In order to maintain the well distributed energy consumption constraint, all the remaining nodes, which are advanced nodes, have to become cluster heads with probability 1 for the next $\frac{1}{p_{opt}} \cdot m \cdot (1 + \alpha)$ rounds of the epoch. But the threshold $T(s)$ is increasing with the number of rounds within each epoch and becomes equal to 1 only in the last round (all the remaining nodes in the last round become cluster head with probability 1). So the above constraint is not satisfied. Figure 4 shows that the performance of this naïve solution is very close to that of LEACH. In the next subsection, we introduce SEP where the extra energy of advanced nodes is forced to be expended within subepochs of the original epoch.

5.2 Guaranteed Well Distributed Energy Consumption Constraints in the Stable Period

In this section we propose a solution, we call SEP (Stable Election Protocol), which is based on the initial energy of the nodes.

This solution is more applicable compared to any solution which assumes that each node knows the total energy of the network in order to adapt its election probability to become a cluster head according to its remaining energy [5]. Our approach is to assign a weight to the optimal probability p_{opt} . This weight must be equal to the initial energy of each node divided by the initial energy of the normal node. Let us define as p_{nrm} the weighted election probability for normal nodes and p_{adv} the weighted election probability for the advanced nodes.

Virtually there are $n \times (1 + \alpha \cdot m)$ nodes with energy equal to the initial energy of a normal node. In order to maintain the minimum energy consumption in each round within an epoch, the average number of cluster heads per round per epoch must be constant and equal to $n \times p_{opt}$. In the heterogeneous scenario the average number of cluster heads per round per epoch is equal to $n \cdot (1 + \alpha \cdot m) \times p_{nrm}$ (because each virtual node has the initial energy of a normal node). The weighed probabilities for normal and advanced nodes are, respectively:

$$p_{nrm} = \frac{p_{opt}}{1 + \alpha \cdot m}$$

$$p_{adv} = \frac{p_{opt}}{1 + \alpha \cdot m} \times (1 + \alpha)$$

In Equation (1), we replace p_{opt} by the weighted probabilities to obtain the threshold that is used to elect the cluster head in each round. We define as $T(s_{nrm})$ the threshold for normal nodes and $T(s_{adv})$ the threshold for advanced nodes. Thus, for normal nodes, we have:

$$T(s_{nrm}) = \begin{cases} \frac{p_{nrm}}{1 - p_{nrm} \cdot (r \bmod \frac{1}{p_{nrm}})} & \text{if } s_{nrm} \in G' \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

where r is the current round, G' is the set of nodes that have not become cluster heads within the last $\frac{1}{p_{nrm}}$ rounds of the epoch, and $T(s_{nrm})$ is the threshold applied to a population of $n \cdot (1 - m)$ (normal) nodes. This guarantees that each normal node will become a cluster head exactly once every $\frac{1}{p_{opt}} \cdot (1 + \alpha \cdot m)$ rounds per epoch, and that the average number of cluster heads per round per epoch is equal to $n \cdot (1 - m) \times p_{nrm}$.

Similarly, for advanced nodes, we have:

$$T(s_{adv}) = \begin{cases} \frac{p_{adv}}{1 - p_{adv} \cdot (r \bmod \frac{1}{p_{adv}})} & \text{if } s_{adv} \in G'' \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

where G'' is the set of nodes that have not become cluster heads within the last $\frac{1}{p_{adv}}$ rounds of the epoch, and $T(s_{adv})$ is the threshold applied to a population of $n \cdot m$ (advanced) nodes. This guarantees that each advanced node will become a cluster head exactly once every $\frac{1}{p_{opt}} \cdot \frac{1 + \alpha \cdot m}{1 + \alpha}$ rounds. Let us define this period as *sub-epoch*. It is clear that each epoch (let us refer to this epoch as "heterogeneous epoch" in our heterogeneous setting) has $1 + \alpha$ sub-epochs and as a result, each advanced node becomes a cluster head exactly $1 + \alpha$ times within a heterogeneous epoch. The average number of cluster heads per round per heterogeneous epoch (and sub-epoch) is equal to $n \cdot m \times p_{adv}$.

The average number of cluster heads per round per heterogeneous epoch is equal to the average number of cluster heads that are normal nodes per round per heterogeneous epoch plus the average number of cluster heads that are advanced nodes per round per sub-epoch. This average number is given by:

$$n \cdot (1 - m) \times p_{nrm} + n \cdot m \times p_{adv} = n \times p_{opt}$$




Figure 5: A numerical example for a heterogeneous network with parameters $m = 0.2$ and $\alpha = 3$ and $p_{opt} = 0.1$. We define as $x = r \bmod \frac{1}{p_{opt}}$ and as $x' = r \bmod \frac{1}{p_{nrm}}$, where r is the current round.

which is the desired number of cluster heads per round per epoch. We next discuss the implementation of our SEP protocol.

5.3 SEP Deployment

As mentioned in Section 1, the heterogeneity in the energy of nodes could result from normal network operation. For example, nodes could, over time, expend different amounts of energy due to the radio communication characteristics, random events such as short-term link failures or morphological characteristics of the field (e.g. uneven terrain). To deal with such heterogeneity, our SEP protocol could be triggered whenever a certain energy threshold is exceeded at one or more nodes. Non-cluster heads could periodically attach their remaining energy to the messages they sent during the handshaking process with their cluster heads, and the cluster heads could send this information to the sink. The sink can check the heterogeneity in the field by examining whether one or a certain number of nodes reach this energy threshold. If so, then the sink could broadcast to cluster heads in that round the values for p_{nrm} and p_{adv} , in turn cluster heads unicast these values to nodes in their clusters according to the energy each one has attached earlier during the handshaking process.

If some of the nodes already in use have not been programmed with this capability, a reliable transport protocol, such as the one proposed in [10], could be used to program such sensors. Evaluating the overhead of such SEP deployment is a subject of our on-going work.

5.4 Numerical Example

Assume that 20% of the nodes are advanced nodes ($m = 0.2$) and equipped with 300% more energy than other (normal) nodes ($\alpha = 3$). Consider a population of a sensor network in a $100m \times 100m$ field of 100 nodes. The p_{opt} for this setting is approximately equal to 0.104325 (cf. Figure 2). For simplicity let us set $p_{opt} = 0.1$. This means that on average, 10 nodes must become cluster heads per round.

If we consider a homogeneous scenario where each node has initial energy equal to the energy of a normal node, then the epoch would be equal to $\frac{1}{p_{opt}} = 10$ rounds. In our heterogeneous case, the extended heterogeneous epoch is equal to $\frac{1 + \alpha \cdot m}{p_{opt}} = \frac{1}{p_{nrm}} = 16$ rounds, and each sub-epoch is equal to $\frac{1}{p_{opt}} \cdot \frac{1 + \alpha \cdot m}{1 + \alpha} = 4$ rounds, as illustrated in Figure 5. On average, $n \cdot (1 - m) \times p_{nrm} = 5$ normal nodes become cluster heads per round and all of them will become cluster heads exactly once within 16 rounds (one heterogeneous epoch). Furthermore, on average, $n \cdot m \times p_{adv} = 5$ advanced nodes become cluster head per round. The total number of sensors that become cluster heads (both normal and advanced) is equal to 10, which is the desired number. Moreover each advanced sensor becomes a cluster head exactly once every sub-epoch and becomes




Figure 6: Number of alive nodes using LEACH in the presence of heterogeneity: (top) $m = 0.1$ and $\alpha = 2$, and (bottom) $m = 0.2$ and $\alpha = 1$.

$(1 + \alpha)$ times a cluster head within a heterogeneous epoch, *i.e.* each advanced node becomes a cluster head 4 times within a heterogeneous epoch.

Operation	Energy Dissipated
Transmitter/Receiver Electronics	$E_{elec} = 50nJ/bit$
Data Aggregation	$E_{DA} = 5nJ/bit/signal$
Transmit Amplifier if $d_{max\ to\ BS} \leq d_0$	$\epsilon_{fs} = 10pJ/bit/m^2$
Transmit Amplifier if $d_{max\ to\ BS} \geq d_0$	$\epsilon_{mp} = 0.0013pJ/bit/m^4$

Table 1: Radio characteristics used in our simulations.

6. SIMULATION RESULTS

We simulate a clustered wireless sensor network in a field with dimensions $100m \times 100m$. The population of the sensors is equal to $n = 100$ and the nodes, both normal and advanced, are randomly (uniformly) distributed over the field. This means that the horizontal and vertical coordinates of each sensor are randomly selected between 0 and the maximum value of the dimension. The

sink is in the center and the maximum distance of any node from the sink is approximately $70m$ (the setting of Figure 3). This setting is realistic for most of outdoor applications. The initial energy of a normal node has been set to $E_0 = 0.5J$ (equal to one AA battery)—Although this value is arbitrary for the purpose of this study, this does not affect the behavior of our method. The radio characteristics used in our simulations are summarized in Table 1. The size of the message that nodes send to their cluster heads as well as the size of the (aggregate) message that a cluster head sends to the sink is set to 4000 bits.

In the next subsections we simulate the heterogeneous-oblivious LEACH and our SEP protocol, in the presence of heterogeneity in the initial energy of nodes. We evaluate the behavior of both protocols in terms of the performance measures defined in Section 3. We also examine the sensitivity of SEP to the degree of heterogeneity in the network. We first summarize our general observations:

- In a wireless sensor network of heterogeneous nodes, LEACH goes to unstable operation sooner as it is very sensitive to such heterogeneity.
- Our SEP protocol successfully extends the stable region by being aware of heterogeneity through assigning probabilities of cluster-head election weighted by the relative initial energy of nodes.
- Due to extended stability, the throughput of SEP is also higher than that of current (heterogeneous-oblivious) clustering protocols.
- The performance of SEP is observed to be close to that of an ideal upper bound obtained by distributing the additional energy of advanced nodes uniformly over all nodes in the sensor field.
- SEP is more resilient than LEACH in judiciously consuming the extra energy of advanced nodes—SEP yields longer stability region for higher values of extra energy.

6.1 Results for LEACH

The results of our LEACH simulations are shown in Figure 6(a) for $m = 0.1$ and $\alpha = 2$. We observe that LEACH takes some advantage of the presence of heterogeneity (advanced nodes), as the first node dies after a significantly higher number of rounds (*i.e.* longer stability period) compared to the homogeneous case ($m = \alpha = 0$). The lifetime of the network is increased, but as we will show later this does not mean that the nodes transmit (*i.e.* the throughput is low). The reason is that after the death of a significant number of nodes, the cluster head election process becomes unstable and as a result less nodes become cluster heads. Even worse, during the last rounds, there are only few rounds where more than one cluster head is elected.³

We repeat the same experiment, but now the heterogeneity parameters are set to $m = 0.2$ and $\alpha = 1$, however $m \times \alpha$ remains constant. Our simulation results are shown in Figure 6(bottom). Although the length of the stability region (until the first node dies) is pretty stable, LEACH takes more advantage of the presence of heterogeneity manifested in a higher number of advanced nodes.

In Figure 7, a detailed view of the behavior of LEACH is illustrated, for different distributions of heterogeneity. In Figure

³For both LEACH and SEP, the length of the stable region obtained from independent simulation runs (*i.e.* starting from different random number seeds) is pretty stable for the same values of m and α .




Figure 7: LEACH behavior in the presence of heterogeneity with $m = 0.2$ and $\alpha = 3$: (a) Alive nodes per round; (b) Average number of cluster heads per round per epoch; (c) Normal nodes per round; (d) Advanced nodes per round

7(a), the number of alive nodes is shown for the scenarios ($m = 0.2, \alpha = 1$) and ($m = 0.2, \alpha = 3$). LEACH fails to take full advantage of the heterogeneity (extra energy) as in both scenarios, the first node dies almost at the same round. Moreover, the normal nodes die in both cases very fast (Figure 7(c)) and as a result the sensing field becomes sparse very fast. On the other hand, advanced nodes die in a very slow fashion (Figure 7(d)), because they are not elected very often as cluster heads after the death of the normal nodes (and thus they do not transmit most of the time)—this is because the election process for cluster heads has become unstable and the number of cluster heads elected are less than the optimal number. Furthermore, as shown in Figure 7(b), when a significant number of normal nodes are dead the average number of cluster heads per round per epoch is less than one. This means that in most of the rounds there is no cluster head, so in our model the remaining nodes can not report their values to the sink.

6.2 Results for SEP

In this subsection we compare the performance of our SEP protocol to 1) LEACH in the same heterogeneous setting, and 2) LEACH where the the extra initial energy of advanced nodes is uniformly distributed over all nodes in the sensor field. This latter setting turns

out to provide the highest throughput during the unstable region—we henceforth refer to it as FAIR (for the “fair” distribution of extra energy over existing nodes).

Figure 8(top) shows results for the case of $m = 0.2$ and $\alpha = 1$. It is obvious that the stable region of SEP is extended compared of that of LEACH (by 8%), even though the difference is not very large. Moreover, the unstable region of SEP is shorter than that of LEACH. What is more important to notice is that the stable region of SEP is even greater than FAIR. Furthermore the unstable region of SEP is slightly larger than that of FAIR, and the number of alive nodes per round in SEP is very close to that of FAIR.

Figure 8(bottom) shows results for the case of $m = 0.2$ and $\alpha = 3$. Now SEP takes full advantage of heterogeneity (extra energy of advanced nodes)—the stable region is increased significantly (by 26%) in comparison with that of LEACH. Again the stable region of SEP is greater than that of FAIR. The unstable region of SEP is shorter than that of LEACH, and the number of alive nodes under SEP is close to that of FAIR. This is because the advanced nodes follow the dying process of normal nodes, as the weighted probability of electing cluster heads causes energy of each node to be consumed in proportion to the node’s initial energy.




Figure 8: Comparison between LEACH and SEP in the presence of heterogeneity: (top) $m = 0.2$ and $\alpha = 1$, and (bottom) $m = 0.2$ and $\alpha = 3$.

6.3 Throughput

We assume that the available bandwidth is not tight. Figure 9(top) shows the throughput from cluster heads to the sink. The throughput of SEP is significantly larger than that of LEACH in the stable region and for most of the unstable region. This means that because SEP guarantees cluster heads in more rounds than these cluster heads will report to the sink. It is also worth noticing that the throughput of SEP is greater than that of FAIR during the stable region and very close to that of FAIR at the start of the unstable region. Moreover, the same results are observed in Figure 9(middle) for the throughput of nodes to their cluster heads, as the cluster heads in the case of SEP are elected in a more stable fashion during the unstable period. As a result the overall throughput of SEP is greater than that of LEACH and FAIR during the stable region and close to that of FAIR during the unstable region, as Figure9(bottom) shows.

6.4 Sensitivity of SEP

We study here the sensitivity of our SEP protocol, in terms of the length of the stability period, by varying m and α . Figure 10(top) shows the length of the stability region versus $m \times \alpha$. We found that the performance does not depend on the individual values of m and α but rather on their product, which represents the total amount of




Figure 9: Throughput comparison between LEACH and SEP in the presence of heterogeneity with $m = 0.2$ and $\alpha = 3$: (top) Cluster heads to sink; (middle) Nodes to their cluster heads; and (bottom) Total for the whole network.

extra initial energy brought by advanced nodes. Figure 10(middle) shows the percentage gain in the length of the stability region over the case of $m = 0$ and $\alpha = 0$, i.e. without the added energy of advanced nodes. Figure 10(bottom) shows the percentage gain in the length of the stability region of one protocol over another.

We observe that, as expected, the stability period under FAIR increases linearly with $m \times \alpha$. On the other hand, the stability period under SEP and LEACH increases faster but then more slowly beyond a “knee” point. Moreover, as far as the efficient use of extra energy, the percentage gain in the stability period is maximized under SEP for most values of $m \times \alpha$. In all cases SEP outperforms LEACH.

Interestingly, both SEP and LEACH outperforms FAIR for small amount of heterogeneity (or a small number of advanced nodes)—SEP outperforms FAIR by up to 18% (when $m \times \alpha=0.2$), and LEACH outperforms FAIR by up to 11% (when $m \times \alpha=0.2$). This is because these advanced nodes are uniformly distributed over the sensor field, and when they elect themselves as cluster heads, their “extra” energy is consumed more judiciously than if some of this extra energy was distributed to all nodes (as in FAIR) which are possibly farther away from the sink. This gain over FAIR eventually vanishes when it becomes more beneficial to distribute some extra energy to the fewer normal nodes.

We also notice that the gain of SEP over LEACH increases as $m \times \alpha$ increases—SEP outperforms LEACH by up to 33% when $m \times \alpha=0.9$. The gain of LEACH over FAIR drops much faster than that of SEP after the “knee” point. This indicates that the management of the extra energy of advanced nodes can become difficult, more so for LEACH than our SEP protocol.

7. RELATED WORK

In addition to related work cited throughout the paper, in this section we review specific prior studies that dealt with the heterogeneity in energy of sensor nodes.

The first work that questioned the behavior of clustering protocols in the presence of heterogeneity in clustered wireless sensor networks was [5]. In this work Heinzelman analyzed a method to elect cluster heads according to the energy left in each node. The drawback of this method is that this decision was made per round and assumed that the total energy left in the network was known. The complexity and the assumption of global knowledge of the energy left for the whole network makes this method difficult to implement. Even a centralized approach of this method would be very complicated and very slow, as the feedback should be reliably delivered to each sensor in every round.

In [4], Duarte-Melo and Liu examined the performance and energy consumption of wireless sensor networks, in a field where there are two types of sensors. They consider nodes that are fewer but more powerful that belong to an overlay. All the other nodes have to report to these overlay nodes, and the overlay nodes aggregate the data and send it to the sink. The drawback of this method is that there is no dynamic election of the cluster heads among the two type of nodes, and as a result nodes that are far away from the powerful nodes will die first. The authors estimate the optimal percentage of powerful nodes in the field, but this result is very difficult to use when heterogeneity is a result of operation of the sensor network and not a choice of optimal setting.

In [8], Mhatre and Rosenberg presented a cost based comparative study of homogeneous and heterogeneous clustered wireless sensor networks. They proposed a method to estimate the optimal distribution among different types of sensors, but again this result is hard to use if the heterogeneity is due to the operation of the network. They also studied the case of multi-hop routing within each cluster




Figure 10: Sensitivity of LEACH, SEP, and FAIR to degree of heterogeneity.

(called M-LEACH). Again the drawback of the method is that only powerful nodes can become cluster heads (even though not all of the powerful nodes are used in each round), and that M-LEACH is valid under many assumptions and only when the population of the nodes is very large.

8. CONCLUSIONS AND FUTURE WORK

We proposed SEP (Stable Election Protocol) so every sensor node in a heterogeneous two-level hierarchical network *independently* elects itself as a cluster head based on its initial energy relative to that of other nodes. Unlike [5], we do not require any global knowledge of energy at every election round. Unlike [4, 8], SEP is dynamic in that we do not assume any prior distribution of the different levels of energy in the sensor nodes. Furthermore, our analysis of SEP is not only asymptotic, i.e. the analysis applies equally well to small-sized networks.

We are currently extending SEP to deal with clustered sensor networks with more than two levels of hierarchy and more than two types of nodes. We are also implementing SEP in Berkeley/Crossbow motes and examining deployment issues including dynamic updates of weighted election probabilities based on current heterogeneity conditions.

SEP code and results are publicly available at <http://csr.bu.edu/sep>.

9. REFERENCES

- [1] I. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci. A survey on sensor networks. *IEEE Communications Magazine*, 40(8):102–114, August 2002.
- [2] S. Bandyopadhyay and E. J. Coyle. An energy efficient hierarchical clustering algorithm for wireless sensor networks. In *Proceedings of the 22nd Annual Joint Conference of the IEEE Computer and Communications Societies (Infocom 2003)*, April 2003.
- [3] S. Bandyopadhyay and E. J. Coyle. Minimizing communication costs in hierarchically-clustered networks of wireless sensors. *Comput. Networks*, 44(1):1–16, January 2004.
- [4] E. J. Duarte-Melo and M. Liu. Analysis of energy consumption and lifetime of heterogeneous wireless sensor networks. In *Proceedings of Global Telecommunications Conference (GLOBECOM 2002)*, pages 21–25. IEEE, November 2002.
- [5] W. R. Heinzelman. Application-Specific Protocol Architectures for Wireless Networks, Ph.D. thesis, Massachusetts Institute of Technology, 2000.
- [6] W. R. Heinzelman, A. Chandrakasan, and H. Balakrishnan. Energy-efficient communication protocol for wireless microsensor networks. In *Proceedings of the 33rd Hawaii International Conference on System Sciences (HICSS-33)*, January 2000.
- [7] W. R. Heinzelman, A. P. Chandrakasan, and H. Balakrishnan. An application-specific protocol architecture for wireless microsensor networks. *IEEE Transactions on Wireless Communications*, 1(4):660–670, October 2002.
- [8] V. Mhatre and C. Rosenberg. Homogeneous vs. heterogeneous clustered sensor networks: A comparative study. In *Proceedings of 2004 IEEE International Conference on Communications (ICC 2004)*, June 2004.
- [9] T. J. Shepard. A channel access scheme for large dense packet radio networks. In *Proceedings of ACM SIGCOMM*, pages 219–230, September 1996.
- [10] C.-Y. Wan, A. T. Campbell, and L. Krishnamurthy. PSFQ: a reliable transport protocol for wireless sensor networks. In *2nd IEEE Workshop on Applications and Services in Wireless Networks (WSNA 2002)*, pages 1–11, July 2002.