

August 26, 2002

TR CS-02-13

Efficient Algorithms for Maximum Lifetime Data Gathering and Aggregation in Wireless Sensor Networks

Konstantinos Kalpakis^{1,2}, Koustuv Dasgupta¹, and Parag Namjoshi¹

Computer Science and Electrical Engineering Department
University of Maryland Baltimore County
1000 Hilltop Circle
Baltimore, Maryland 21250

¹Email: {kalpakis,dasgupta,nam1}@csee.umbc.edu

²Supported in part by NASA under Cooperative Agreement NCC5-315 and Contract NAS5-32337, and by NSF under grant IRI-9729495.

Efficient Algorithms for Maximum Lifetime Data Gathering and Aggregation in Wireless Sensor Networks^{1,2}

Konstantinos Kalpakis, Koustuv Dasgupta, and Parag Namjoshi

Abstract

The rapid advances in processor, memory, and radio technology have enabled the development of distributed networks of small, inexpensive nodes that are capable of sensing, computation, and wireless communication. Sensor networks of the future are envisioned to revolutionize the paradigm of collecting and processing information in diverse environments. However, the severe energy constraints and limited computing resources of the sensors, present major challenges for such a vision to become a reality.

We consider a network of energy-constrained sensors that are deployed over a region. Each sensor periodically produces information as it monitors its vicinity. The basic operation in such a network is the systematic gathering and transmission of sensed data to a base station for further processing. During data gathering, sensors have the ability to perform in-network aggregation (fusion) of data packets enroute to the base station. The lifetime of such a sensor system is the time during which we can gather information from all the sensors to the base station. A key challenge in data gathering is to maximize the system lifetime, given the energy constraints of the sensors.

Given the location of n sensors and a base station together with the available energy at each sensor, we are interested in finding an efficient manner in which data should be collected from all the sensors and transmitted to the base station, such that the system lifetime is maximized. This is the maximum lifetime data gathering problem. In this paper, we describe novel algorithms, with worst-case running times polynomial in n , to solve the data gathering problem with aggregation in sensor networks. Our experimental results demonstrate that the proposed algorithms significantly outperform previous methods in terms of system lifetime.

Keywords – sensor networks, data gathering, data aggregation, energy-efficient protocols, lifetime, maximum flow, minimum cut.

I Introduction

The recent advances in micro-sensor technology and low-power analog/digital electronics, have led to the development of distributed, wireless networks of sensor devices [11, 19, 20]. Sensor networks of the future are envisioned to consist of hundreds of inexpensive nodes, that can be readily deployed in physical environments to collect useful information (e.g. seismic, acoustic, medical and surveillance data) in a robust and autonomous manner. However, there are several obstacles that need to be overcome before this vision becomes a reality [9]. Such obstacles arise from the limited energy, computing capabilities and communication resources available to the sensors.

We consider a system of sensor nodes that are homogeneous and highly energy-constrained. Further, replenishing energy via replacing batteries on hundreds of nodes (in possibly harsh terrains) is infeasible. The basic operation in such a system is the systematic gathering of sensed data to be eventually transmitted

¹Supported in part by NASA under Cooperative Agreement NCC5-315 and Contract NAS5-32337, and by NSF under grant IRI-9729495.

²Computer Science and Electrical Engineering Department, University of Maryland Baltimore County, 1000 Hilltop Circle, Baltimore, MD 21250. E-mail: {kalpakis,dasgupta,nam1}@csee.umbc.edu. Phone: 410-455-3143. Fax: 410-455-3969.

Please send all correspondence to Prof. Kalpakis.

to a base station for processing. The key challenge in such data gathering is conserving the sensor energies, so as to maximize their lifetime. To this end, there are several power-aware routing protocols for wireless ad hoc networks discussed in the literature [14, 21]. In the context of sensor networks, LEACH [8] proposes a clustering-based protocol for transmitting data to the base station. The main features include local co-ordination for cluster formation among sensors, randomized rotation of cluster heads for improved energy utilization, and local data compression to reduce global communication. Chang and Tassiulas [3, 4] describe data routing algorithms that maximize the time until the energies of the wireless nodes drain out. In related work, Bhardwaj et al [2] derive upper bounds on the lifetime of a sensor network that collects data from a specified region using some energy-constrained nodes.

Data fusion or aggregation has emerged as a basic tenet in sensor networks. The key idea is to combine data from different sensors to eliminate redundant transmissions, and provide a rich, multi-dimensional view of the environment being monitored. Krishnamachari et al [13] argue that this paradigm shifts the focus from address-centric approaches (finding routes between pairs of end nodes) to a more data-centric approach (finding routes from multiple sources to a destination that allows in-network consolidation of data). Madden et al [18] describe the TinyOS operating system that can be used by an ad-hoc network of sensors to locate each other and route data. The authors discuss the implementation of five basic database aggregates, i.e. COUNT, MIN, MAX, SUM, and AVERAGE, based on the TinyOS platform and demonstrate that such a generic approach for aggregation leads to significant power (energy) savings. The focus of the work in [18] is on a class of aggregation predicates that is particularly well suited to the in-network regime. Such aggregates can be expressed as an aggregate function f over the sets a and b , such that $f(a \cup b) = g(f(a), f(b))$. Other previous works [10, 9, 15, 16] in the related area aim at reducing the energy expended by the sensors during the process of data gathering. Directed diffusion [10] is based on a network of nodes that can co-ordinate to perform distributed sensing of an environmental phenomenon. Such an approach achieves significant energy savings when intermediate nodes aggregate responses to queries. The SPIN protocol [9] uses meta-data negotiations between sensors to eliminate redundant data transmissions through the network. In PEGASIS [15], sensors form chains so that each node transmits and receives from a nearby neighbor. Gathered data moves from node to node, gets aggregated and is eventually transmitted to the base station. Nodes take turns to transmit so that the average energy spent by each node gets reduced. Lindsey et al [16] describe a hierarchical scheme based on PEGASIS that reduces the average energy consumed and delay incurred in gathering the sensed data.

In this paper, we describe novel algorithms for data gathering and aggregation in sensor networks. We define the lifetime of a sensor network to be the time during which we can gather information from all the sensors to the base station. Given the location of sensors and the base station and the available energy at each sensor, we are interested in finding an efficient manner in which the data should be collected and aggregated from all the sensors and transmitted to the base station, such that the system lifetime is maximized. This is the Maximum Lifetime Data Aggregation (MLDA) problem. We first propose a near-optimal polynomial-time algorithm for solving the MLDA problem. The proposed algorithm, while performing significantly better than existing protocols in terms of system lifetime, is computationally expensive for large sensor networks. Consequently, we describe clustering-based heuristics approaches for maximum lifetime data gathering and aggregation in large-scale sensor networks. Finally, we provide experimental results to show that (i) for smaller sensor networks the MLDA algorithm achieves system lifetimes that are 1.15 to 2.32 times better when compared to an existing data gathering protocol, (ii) for larger networks, our clustering-based heuristics achieve as much as a factor of 2.61 increase in the system lifetime when compared to the same protocol.

The rest of the paper is organized as follows. In section II, we describe the system model and the data gathering problem. Section III provides a detailed description of our MLDA algorithm to solve the maximum lifetime data aggregation problem in sensor networks. Next, in section IV, we propose clustering-based heuristics to solve the problem efficiently for larger sensor networks. In section V, we present the experimental results and finally we conclude the paper in section VI.

II The Data Gathering Problem

A. System Model

Consider a network of n sensor nodes $1, 2, \dots, n$ and a base station node t labeled $n + 1$ distributed over a region. The locations of the sensors and the base station are fixed and known apriori. Each sensor produces some information as it monitors its vicinity. We assume that each sensor generates one data packet per time unit to be transmitted to the base station. For simplicity, we refer to each time unit as a *round*. We assume that all data packets have size k bits. The information from all the sensors needs to be gathered at each round and sent to the base station for processing. We assume that each sensor has the ability to transmit its packet to any other sensor in the network or directly to the base station. Further, each sensor i has a battery with finite, non-replenishable energy \mathcal{E}_i . Whenever a sensor transmits or receives a data packet, it consumes some energy from its battery. The base station has an unlimited amount of energy available to it.

Our energy model for the sensors is based on the *first order radio model* described in [8]. A sensor consumes $\epsilon_{elec} = 50nJ/bit$ to run the transmitter or receiver circuitry and $\epsilon_{amp} = 100pJ/bit/m^2$ for the transmitter amplifier. Thus, the energy consumed by a sensor i in receiving a k -bit data packet is given by,

$$RX_i = \epsilon_{elec} \cdot k \tag{1}$$

while the energy consumed in transmitting a data packet to sensor j is given by,

$$TX_{i,j} = \epsilon_{elec} \cdot k + \epsilon_{amp} \cdot d_{i,j}^2 \cdot k \tag{2}$$

where $d_{i,j}$ is the distance between nodes i and j .

B. Problem Statement

We define the *lifetime* T of the system to be the number of rounds until the first sensor is drained of its energy. A *data gathering schedule* specifies, for each round, how the data packets from all the sensors are collected and transmitted to the base station. For brevity, we refer to a data gathering schedule simply as a schedule. Observe that a schedule can be thought of as a collection of T directed trees, each rooted at the base station and spanning all the sensors i.e. a schedule has one tree for each round. The lifetime of a schedule equals the lifetime of the system under that schedule. Clearly, the system lifetime is intrinsically connected to the data gathering schedule. Our objective is to find a schedule that maximizes the system lifetime T .

III MLDA: Maximum Lifetime Data Gathering with Aggregation

Data aggregation performs in-network fusion of data packets, coming from different sensors enroute to the base station, in an attempt to minimize the number and size of data transmissions and thus save sensor energies [10, 8, 13, 15]. Such aggregation can be performed when the data from different sensors are highly correlated. As in previous work [10, 8, 13, 15], we make the simplistic assumption that an intermediate sensor can aggregate multiple incoming packets into a single outgoing packet.

The Maximum Lifetime Data Aggregation (MLDA) problem: Given a collection of sensors and a base station, together with their locations and the energy of each sensor, find a data gathering schedule with maximum lifetime, where sensors are permitted to aggregate incoming data packets.

Consider a schedule \mathcal{S} with lifetime T rounds. Let $f_{i,j}$ be the total number of packets that node i (a sensor) transmits to node j (a sensor or base station) in \mathcal{S} . Since any valid schedule must respect the energy

Objective :

$$\text{maximize } T \tag{4}$$

Constraints :

$$\sum_{j=1}^{n+1} f_{i,j} \cdot TX_{i,j} + \sum_{j=1}^n f_{j,i} \cdot RX_i \leq \mathcal{E}_i. \tag{5}$$

$$\sum_{j=1}^n \pi_{j,i}^{(k)} = \sum_{j=1}^{n+1} \pi_{i,j}^{(k)}, \quad \forall i = 1, 2, \dots, n \text{ and } i \neq k \tag{6}$$

$$T + \sum_{j=1}^n \pi_{j,k}^{(k)} = \sum_{j=1}^{n+1} \pi_{k,j}^{(k)} \tag{7}$$

$$0 \leq \pi_{i,j}^{(k)} \leq f_{i,j}, \quad \forall i = 1, 2, \dots, n \text{ and } \forall j = 1, 2, \dots, n+1 \tag{8}$$

$$\sum_{i=1}^n \pi_{i,n+1}^{(k)} = T \tag{9}$$

where $k = 1, 2, \dots, n$ and all variables T , $f_{i,j}$, and $\pi_{i,j}^{(k)}$ are required to be non-negative integers.

Table 1: Integer program for finding an optimal admissible flow network for the MLDA problem.

constraints at each sensor, it follows that for each sensor $i = 1, 2, \dots, n$,

$$\sum_{j=1}^{n+1} f_{i,j} \cdot TX_{i,j} + \sum_{j=1}^n f_{j,i} \cdot RX_i \leq \mathcal{E}_i. \tag{3}$$

Recall that each sensor, for each one of the T rounds, generates one data packet that needs to be collected, possibly aggregated, and eventually transmitted to the base station.

The schedule \mathcal{S} induces a flow network $G = (V, E)$. The flow network G is a directed graph having as nodes all the sensors and the base station, and having edges (i, j) with capacity $f_{i,j}$ whenever $f_{i,j} > 0$.

Theorem 1 *Let \mathcal{S} be a schedule with lifetime T , and let G be the flow network induced by \mathcal{S} . Then, for each sensor s , the maximum flow from s to the base station t in G is $\geq T$.*

Proof. Each data packet transmitted from a sensor must reach the base station. Observe that, the packets from s could possibly be aggregated with one or more packets from other sensors in the network. Intuitively, we need to guarantee that each of the T values from s *influences* the final value(s) received at the base station. In terms of network flows, this implies that sensor s must have a maximum $s - t$ flow of size $\geq T$ to the base station in the flow network G . ■

Thus, a necessary condition for a schedule to have lifetime T is that each node in the induced flow network can push flow T to the base station t . Stated otherwise, each sensor s must have a minimum $s - t$ cut of capacity (size) $\geq T$ to the base station [5]. Next, we consider the problem of finding a flow network G with maximum T , that allows each sensor to push flow T to the base station, while respecting the energy constraints in (5) at all the sensors. We call such a flow network G an *admissible* flow network with lifetime T . An admissible flow network with maximum lifetime is called an *optimal admissible* flow network. Clearly, what needs to be found are the capacities of the edges in G .

A. Finding a near-optimal admissible flow network

An optimal admissible flow network can be found using an integer program with linear constraints. The integer program, in addition to the variables for the lifetime T and the edge capacities $f_{i,j}$, uses the following

variables: for each sensor $k = 1, 2, \dots, n$, let $\pi_{i,j}^{(k)}$ be a flow variable indicating the flow that k sends to the base station t over the edge (i, j) .

The integer program computes the maximum system lifetime T subject to the energy constraint (5) and the additional linear constraints (6)–(9) for each sensor, as shown in Table 1. For each sensor $k = 1, 2, \dots, n$, constraints (6) and (7) enforce the flow conservation principle at the sensor; constraint (9) ensures that T flow from sensor k reaches the base station; and constraint (8) ensures that the capacity constraints on the edges of the flow network are respected. Moreover, constraint (5) is used to guarantee that the edge capacities of the flow network respect the sensor’s available energy. Finally, for the integer program, all variables are required to take non–negative integer values. The linear relaxation of the above integer program, i.e. when all the variables are allowed to take fractional values, can be computed in polynomial–time. Then, we can obtain a very good approximation for the optimal admissible flow network by first fixing the edge capacities to the floor of their values obtained from the linear relaxation so that the energy constraints are all satisfied; and then solving the linear program (4) subject to constraints (6)–(9) without requiring anymore that the flows are integers (since a solution with integer flows can always be found).³

B. Constructing a schedule from an admissible flow network

Next, we discuss how to get a schedule from an admissible flow network. Recall that a schedule is a collection of directed trees that span all the sensors and the base station, with one such tree for each round. Each such tree specifies how data packets are gathered and transmitted to the base station. We call these trees *aggregation trees*. An aggregation tree may be used for one or more rounds; we indicate the number of rounds f , that an aggregation tree is used, by associating the value f with each one of its edges; we call f to be the lifetime of the aggregation tree.




Figure 1: An admissible flow network G with lifetime 100 rounds, and two aggregation trees A_1 and A_2 with lifetimes 60 and 40 rounds respectively.

Figure 1 provides an example of an admissible flow network G with lifetime $T = 100$ and two aggregation trees A_1 and A_2 , with lifetimes 60 and 40 rounds respectively. By looking at one of these trees, say A_1 , we see that for each one of 60 rounds, sensors 2 and 3 transmit one data packet to sensor 1, which in turn aggregates the incoming packets with its own data packet, and then sends one data packet to the base station. Similarly, for each of the remaining 40 rounds (using A_2), sensors 1 and 2 transmit one data packet to sensor 3, which in aggregates the incoming packets with its own packet, and sends one data packet to the base station. We next describe an algorithm to construct aggregation trees from an admissible flow network G with lifetime T .

Definition 1 : Given an admissible flow network G with lifetime T and a directed tree A rooted at the base station t with lifetime f , we define the (A, f) –*reduction* G' of G to be the flow network that results from G

³The reduction in the system lifetime achieved, w.r.t the fractional optimal lifetime, is at most the maximum cardinality of any min s–t cut.

after reducing by f , the capacities of all of its edges that are also in A . We call G' the (A, f) -reduced G .

Definition 2 : An (A, f) -reduction G' of G is *feasible* if the maximum flow from v to the base station t in G' is $\geq T - f$ for each vertex v in G' .

Note that A does not have to span all the vertices of G , and thus it is not necessarily an aggregation tree. Moreover, if A is an aggregation tree, with lifetime f , for an admissible flow network G with lifetime T , and the (A, f) -reduction of G is feasible, then the (A, f) -reduced flow network G' of G is an admissible flow network with lifetime $T - f$. Therefore, we can devise a simple iterative procedure, to construct a schedule for an admissible flow network G with lifetime T , provided we can find such an aggregation tree A .

We use the GETTREE algorithm in figure 2 to get an aggregation tree A with lifetime $f \leq T$ from an admissible flow network G with lifetime T . Throughout this algorithm, we maintain the invariant that A is a tree rooted at t and the (A, f) -reduction of G is feasible. Tree A is formed as follows. Initially A contains just the base station. While A does not span all the sensors, we find and add to A an edge $e = (i, j)$, where $i \notin A$ and $j \in A$, provided that the (A', f) -reduction of G is feasible – here A' is the tree A together with the edge e , and f is 1 or the minimum of the capacities of the edges in A' . Finally, we can compute a collection of aggregation trees from an admissible flow network G with lifetime T by using the GETSCHEDULE algorithm in figure 3, such that T data packets from each of the sensors are aggregated and transmitted to the base station t . We refer to the algorithm described in this section, for computing a maximum lifetime data gathering schedule, as the MLDA algorithm.⁴

```

GETTREE (Flow Network  $G$ , Lifetime  $T$ , Base Station  $t$ )
1   initialize  $f \leftarrow 1$ 
2   let  $A = (V_o, E_o)$  where  $V_o = \{t\}$  and  $E_o = \emptyset$ 
3   while  $A$  does not span all the nodes of  $G$  do
4     for each edge  $e = (i, j) \in G$  such that  $i \notin V_o$  and  $j \in V_o$  do
5       let  $A'$  be  $A$  together with the edge  $e$ 
6       // check if the  $(A', 1)$ -reduction of  $G$  is feasible
7       let  $G_r$  be the  $(A', 1)$ -reduction of  $G$ 
8       if  $\text{MAXFLOW}(v, t, G_r) \geq T - 1$  for all nodes  $v$  of  $G$ 
9         // replace  $A$  with  $A'$ 
10         $V_o \leftarrow V_o \cup \{i\}$ ,  $E_o \leftarrow E_o \cup \{e\}$ 
11        break
12  let  $c_{min}$  be the minimum capacity of the edges in  $A$ 
13  let  $G_r$  be the  $(A, c_{min})$ -reduction of  $G$ 
14  if  $\text{MAXFLOW}(v, t, G_r) \geq T - c_{min}$  for all nodes  $v$  of  $G$ 
15     $f \leftarrow c_{min}$ 
16  replace  $G$  with the  $(A, f)$ -reduction of  $G$ 
17  return  $f, G, A$ 

```

Figure 2: Constructing aggregation tree A with lifetime f from an admissible flow network G with lifetime T .

Theorem 2 Let $G = (V, E)$ be an admissible flow network with edge capacities $c(e)$, $e \in E$. Let $t \in V$ be the base station and T a positive integer i.e. the lifetime of the system. Then, the GETSCHEDULE algorithm can always find a sequence of aggregation trees that can be used to aggregate and transmit T data packets from each sensor to the base station.

Proof. Refer to the Appendix for proof details.

⁴The MLDA algorithm appears as an extended abstract in [12]

```

GETSCHEDULE(Flow Network  $G$ , Lifetime  $T$ , BaseStation  $t$ )
1    $\mathcal{S} \leftarrow \emptyset$ 
2   while  $T > 0$  do
3      $[f, G, A] \leftarrow \text{GETTREE}(G, T, t)$ 
4      $\mathcal{S} \leftarrow \mathcal{S} \cup \{A\}$ 
5      $T \leftarrow T - f$ 
6   return  $\mathcal{S}$ 

```

Figure 3: Constructing a schedule \mathcal{S} from an admissible flow network G with lifetime T .

C. Worst-case running time of MLDA

The running time analysis of the MLDA algorithm is based on the following three lemmas.

Lemma 1 (Bertsimas and Tsitsiklis [1]) *An ϵ -optimal solution to a linear program with n variables can be found in $O(n^{3.5} \log \frac{1}{\epsilon} + n^5 \log(nU))$ time, where U is the maximum magnitude of the coefficients of the linear program.*

Proof. See the complexity analysis of the potential reduction algorithm in Bertsimas and Tsitsiklis [1, pp. 418]. ■

Lemma 2 (Goldberg and Rao [7]) *Given a flow network $G = (V, E)$ with integral edge capacities bounded by U , a maximum s - t flow can be computed in $O(\min(V^{2/3}, E^{1/2})E \log(V^2/E) \log U)$ time, where s and t are any two vertices of G .*

Proof. See Goldberg and Rao [7]. ■

Lemma 3 *Consider a sensor network with a base station t and n sensors. Each sensor has initial energy no more than \mathcal{E}_{max} and consumes energy for receiving and transmitting k -bit data packets according to the first order radio model [8]. Then, the lifetime of the sensor network is $\leq \mathcal{E}_{max}/(\epsilon_{elec} \cdot k) = O(1)$.*

Proof. Let d_{min} be the minimum distance of a sensor from the base station t . In each round, one data packet needs to be collected from each of the sensors and transmitted to the base station. From the first order radio model, Eqs. (1) and (2), the minimum total energy expended by all the sensors in one round is at least

$$(n - 1) \cdot \epsilon_{elec} \cdot k + (\epsilon_{elec} \cdot k + \epsilon_{amp} \cdot d_{min}^2 \cdot k) = n \cdot \epsilon_{elec} \cdot k + \epsilon_{amp} \cdot d_{min}^2 \cdot k, \quad (10)$$

since there are $n - 1$ packet transmissions over distance at least 0, and one packet transmission over distance at least d_{min} . Therefore, the system lifetime T , i.e. the number of rounds before the first sensor is drained of its energy, is at most

$$\frac{\text{total energy of all the sensors}}{\text{minimum energy consumed in one round}} = \frac{n \cdot \mathcal{E}_{max}}{n \cdot \epsilon_{elec} \cdot k + \epsilon_{amp} \cdot d_{min}^2 \cdot k} \leq \frac{\mathcal{E}_{max}}{\epsilon_{elec} \cdot k} = O(1) \quad (11)$$

Theorem 3 *The worst-case running time of the MLDA algorithm is $O(n^{15} \log n)$, where n is the number of sensors.*

Proof. Consider a sensor network with n sensors and maximum lifetime T .

To compute an admissible flow network, the MLDA algorithm solves the linear program (4) twice. Observe that the linear program (4) has $O(n^3)$ variables and coefficients of magnitude $O(T)$. Using Lemmas 3 and 1, it follows that the time t_{LP} to compute an ϵ -approximate solution to the linear program (4) is

$$t_{LP} = O(n^{10.5} \log \frac{1}{\epsilon} + n^{15} \log n) = O(n^{15} \log n), \quad (12)$$

since ϵ is assumed to be a small constant.

Next, we look at the running time of the GETSCHEDULE procedure. The GETSCHEDULE procedure makes $O(T)$ calls to the GETTREE routine. Each call of the GETTREE routine involves $O(V^2E)$ MAXFLOW computations on the admissible flow network $G = (V, E)$.⁵ Since G has $O(n)$ vertices and $O(n^2)$ edges with capacity $O(T) = O(1)$, using Lemma 2, the running time of a MAXFLOW computation on G is $t_{MAXFLOW} = O(n^{8/3} \log n)$. Thus, the running time of the GETSCHEDULE routine is

$$\leq T \cdot V^2E \cdot t_{MAXFLOW} = O(n^{20/3} \log n). \quad (13)$$

Therefore, the total worst-case running time of MLDA is

$$\leq 2 \cdot t_{LP} + O(n^{20/3} \log n) = O(n^{15} \log n) + O(n^{20/3} \log n) = O(n^{15} \log n). \quad (14)$$

■

We report actual running times of sample runs of the MLDA algorithm in section V.

IV CMLDA: Clustering-based MLDA

Given the location of n sensors and a base station t , we can find a maximum lifetime data gathering schedule using the MLDA algorithm. However, it involves solving a linear program (in Table 1) with $O(n^3)$ variables and constraints. For large sensor networks, i.e. for large values of n , this can be computationally expensive. In this section, we propose a clustering-based approach to solve the data gathering problem efficiently for large networks. In particular, we describe two heuristics – GREEDY CMLDA and INCREMENTAL CMLDA, based on the MLDA algorithm in section III.

Consider the set of n sensors $1, 2, \dots, n$ and a base station t labeled as $n + 1$. We assume, without loss of generality, that each sensor has initial energy \mathcal{E} . Let the sensors be partitioned into m clusters ϕ_1, \dots, ϕ_m each consisting of at most c sensors, i.e. $|\phi_i| \leq c$, for $i = 1, 2, \dots, m$, where c is some constant. We refer to each cluster as a *super-sensor*. Such a partitioning of the sensors can be achieved using an appropriate proximity-based clustering algorithm. Let super-sensor ϕ_{m+1} consist only of the base station t .

A. The greedy CMLDA heuristic

We propose a simple heuristic for finding a maximum lifetime data gathering schedule in sensor networks. Our approach is to compute a maximum lifetime schedule for the super-sensors ϕ_1, \dots, ϕ_m with the base station ϕ_{m+1} , and then use this schedule to construct aggregation trees for the sensors.

⁵Note that since each edge of G needs to be checked at most once for inclusion in the aggregation tree [17], the GETTREE routine can be implemented to use $O(VE)$ MAXFLOW computations.

```

BUILD-TREE(AS-Tree  $\mathcal{T}_k$ , Super-Sensor  $\phi$ , Aggregation Tree  $A$ , Base Station  $t$ )
1   while  $A$  does not contain all the sensors in  $\phi$  do
2       find a pair  $(i, j)$ , where  $i \in \phi - A$  and  $j \in A$ , with maximum residual energy
3       add the edge  $(i, j)$  to  $A$ 
4       update the residual energy of sensor  $i$  as  $\mathcal{E}^r[i] \leftarrow \mathcal{E}^r[i] - Tx_{i,j}$ 
5       if  $j \neq t$  then update the residual energy of  $j$  as  $\mathcal{E}^r[j] \leftarrow \mathcal{E}^r[j] - Rx_j$ 
6   foreach child  $\phi'$  of  $\phi$  in  $\mathcal{T}_k$  do
7        $A \leftarrow$  BUILD-TREE( $\mathcal{T}_k, \phi', A, t$ )
8   return  $A$ 

```

Figure 4: Constructing an aggregation tree A for the sensors from an AS -tree \mathcal{T} .

Figure 6 gives a high level view of the greedy clustering-based MLDA (GREEDY CMLDA) heuristic. In the first phase, we assign the initial energy of each super-sensor ϕ_i ($i = 1, 2, \dots, m$) to be the sum of the initial energies of the sensors within it, i.e. $\mathcal{E}_{\phi_i} = \mathcal{E} \cdot |\phi_i|$. The distance between two super-sensors ϕ_i and ϕ_j is assigned to be the maximum distance between any two nodes (sensor or base station) u and v , such that $u \in \phi_i$ and $v \in \phi_j$. Having set up the initial energies and the distances between the super-sensors, we can find a maximum lifetime schedule for the super-sensors ϕ_1, \dots, ϕ_m with the base station as ϕ_{m+1} , using the MLDA algorithm. Recall that such a schedule consists of a collection of directed trees $\mathcal{T}_1, \dots, \mathcal{T}_k$, each rooted at ϕ_{m+1} and spanning over all the super-sensors. To distinguish it from an aggregation tree for the sensors, we refer to each such tree as an aggregation super-tree (or simply an AS -tree).

Next, we use the BUILD-TREE procedure (in figure 4) to construct an aggregation tree A for the sensors from an AS -tree \mathcal{T}_k . Observe that A is a directed tree rooted at t that is used to aggregate one data packet from each sensor. We denote $\mathcal{E}^r[i]$ to be the residual energy at sensor i . Initially, $\mathcal{E}^r[i] = \mathcal{E}$ for each sensor i in the network. Our objective is to construct (one or more) aggregation trees such that the *minimum residual energy* among the n sensors is maximized, thereby maximizing the lifetime of the corresponding data gathering schedule.

Initially, aggregation tree A contains only the base station t . We perform a (pre-order) traversal of the AS -tree \mathcal{T}_k . For each visited super-sensor ϕ , we add the sensors in ϕ to the current aggregation tree A . Let $\phi - A$ denote the set of sensors in ϕ that are not included in A . We define the *residual energy* of a pair (i, j) as $\min\{\mathcal{E}^r[i] - Tx_{i,j}, \mathcal{E}^r[j] - Rx_j\}$, where $i \in \phi - A$ and $j \in A$. Intuitively, on adding a directed edge (i, j) to A , the residual energy at sensor i is reduced by the energy consumed in transmitting a data packet from i to j . Moreover, if j is not the base station, its residual energy is reduced by the energy consumed in receiving a data packet. Among all pairs (i, j) , such that $i \in \phi - A$ and $j \in A$, the BUILD-TREE procedure chooses one with the maximum residual energy and includes the edge (i, j) in A . The process is repeated until all sensors in ϕ are included in A , upon which it continues with the next super-sensor in \mathcal{T}_k . Figure 5 gives an illustration of the BUILD-TREE procedure. The running time of the procedure is polynomial in the number of sensors. Finally, observe that a maximum lifetime schedule for the super-sensors could possibly consist of one or more AS -trees. In this case, we choose (in step 9 of figure 6) the AS -trees in decreasing order of their respective lifetimes; while constructing no more than f_k aggregation trees from a particular AS -tree \mathcal{T}_k , where f_k is the lifetime of the AS -tree \mathcal{T}_k .

Worst-case running time of GREEDY CMLDA:

Consider the GREEDY CMLDA heuristic in figure 6. Steps 1 and 2 (of Phase I) use a greedy proximity-based algorithm to cluster the n sensors into $m + 1$ super-sensors, which can be done in $O(n^2)$ time. In steps 3 and 4, we initialize the energies of and distances between the super-sensors in $O(n^2)$ time. Then, we compute an admissible flow network and a schedule for the super-sensors in $O(m^{15} \log m)$ worst-case running time,

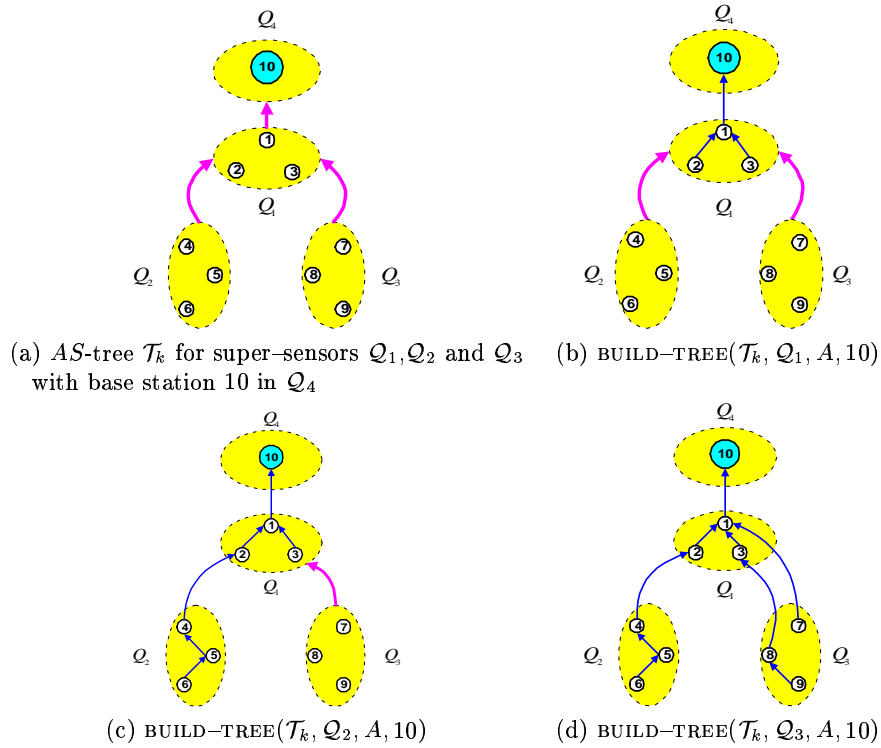


Figure 5: Illustration of the BUILD-TREE procedure for an AS-tree \mathcal{T}_k . Sensors 1, 2, \dots , 9 are partitioned among super-sensors $\mathcal{Q}_1, \mathcal{Q}_2$, and \mathcal{Q}_3 . Super-sensor \mathcal{Q}_4 contains the base station 10. The aggregation tree A is used for collecting one data packet from each sensor and transmitting to the base station.

by solving the linear program (4) with $O(m^3)$ variables (see also Lemma 1 and Theorem 3).

In Phase II of the heuristic, we use this flow network to construct aggregation trees for the sensors. In steps 6 and 7, we initialize the lifetime and the residual energy of each sensor in $O(n)$ time. The *while* loop in step 8 is executed at most $T = O(1)$ times, where T is the lifetime found in Phase I. The time for each iteration of this *while* loop is dominated by the running time of the BUILD-TREE routine. Observe that, the worst-case running time of the BUILD-TREE routine is $O(n^3)$. Thus, the worst-case running time of Phase II is $O(n^3)$, and consequently the worst-case running time of the GREEDY CMLDA heuristic is $O(m^{15} \log m + n^3)$.

As discussed in section V, by appropriately choosing the number of super-sensors m , we can achieve a significant reduction in the actual running time of GREEDY CMLDA with respect to MLDA. For example, for $m = n^{3/17}$, the worst-case running time of GREEDY CMLDA is $O(n^3)$.

B. The Incremental CMLDA Heuristic

In this section, we describe an improved heuristic for the MLDA problem in large-scale sensor networks. Recall that, in solving the MLDA problem, we are essentially interested in *provisioning* the (edge) capacities of an admissible flow network G – such that the s - t flow from each sensor s to the base station t in G is maximized. The proposed heuristic builds such a flow network by incrementally provisioning capacities on its edges.

Consider the set of n sensor nodes and a base station node t . Let the initial energy of each sensor be

INPUT:

Location of n sensors $1, 2, \dots, n$ and a base station t .
Initial energy \mathcal{E} in each sensor.

OUTPUT:

A data gathering schedule \mathcal{S} , i.e. a collection of aggregation trees, with lifetime T .

ALGORITHM:

PHASE I:

1. partition the sensors into m super-sensors ϕ_1, \dots, ϕ_m
2. let super-sensor ϕ_{m+1} consist only of the base station t
3. let the energy of each super-sensor ϕ_i , $i = 1, 2, \dots, m$, be $\mathcal{E}_{\phi_i} \leftarrow \mathcal{E} \cdot |\phi_i|$
4. let the distance between any two super-sensors ϕ_i and ϕ_j be $d_{\phi_i, \phi_j} \leftarrow \max\{d_{u,v} : u \in \phi_i, v \in \phi_j\}$
5. find an admissible flow network G for the super-sensors ϕ_1, \dots, ϕ_m with base station ϕ_{m+1} , and compute a schedule $\mathcal{T} \leftarrow \{\mathcal{T}_1, \dots, \mathcal{T}_k\}$ from G

PHASE II:

6. initialize the schedule $\mathcal{S} \leftarrow \emptyset$ and lifetime $T \leftarrow 0$
 7. let the residual energy of each sensor $i = 1, 2, \dots, n$ be $\mathcal{E}^r[i] = \mathcal{E}$
 8. while $\min\{\mathcal{E}^r[i] : i = 1, 2, \dots, n\} > 0$ do
 9. choose an *AS*-tree \mathcal{T}_k from \mathcal{T}
 10. initialize A to contain only the base station t
 11. compute an aggregation tree $A \leftarrow \text{BUILD-TREE}(\mathcal{T}_k, \phi_{m+1}, A, t)$
 12. update the schedule $\mathcal{S} \leftarrow \mathcal{S} \cup A$ and lifetime $T \leftarrow T + 1$
-

Figure 6: A high level description of the GREEDY CMLDA heuristic.

\mathcal{E} . Let the sensors be partitioned into m super-sensors ϕ_1, \dots, ϕ_m , each having at most c sensors for some constant c . Let super-sensor ϕ_{m+1} consist only of the base station t . The linear relaxation of the integer program in Table 1 can be used to find an optimal admissible flow network G for the super-sensors. Our basic approach is to use flow network G in order to construct an admissible flow network for the sensors.

The INCREMENTAL MLDA heuristic consists of four phases and works as follows.

Phase I:

In the first phase, we assign the initial energy of each super-sensor ϕ_i ($i = 1, 2, \dots, m$) to be the sum of the initial energies of the sensors within it, i.e. $\mathcal{E}_{\phi_i} = \mathcal{E} \cdot |\phi_i|$. The distance between two super-sensors ϕ_i and ϕ_j is assigned to be the maximum distance between any two nodes (sensor or base station) u and v , such that $u \in \phi_i$ and $v \in \phi_j$. Having set up the initial energies and the distances between the super-sensors, we compute an optimal admissible flow network G of maximum lifetime T for the super-sensors ϕ_1, \dots, ϕ_m (with the base station ϕ_{m+1})⁶, using the linear relaxation of the integer program in Table 1. Specifically, we provision the capacity f_{ϕ_i, ϕ_j} between every pair of super-sensors ϕ_i and ϕ_j , such that the system of super-sensors has a lifetime T .

Phase II:

Given the flow network G for the super-sensors and the target lifetime T computed in Phase I, we next determine the capacity provisions that are required for each of the sensors within every super-sensor.

⁶For brevity, we refer to both t and ϕ_{m+1} as the base station.

Objective:

$$\text{minimize } \mathcal{E}_{max} \tag{15}$$

Energy and Capacity Constraints:

$$\sum_{j \in \phi_i} f_{s,j} \cdot TX_{s,j} + \sum_{j \in \phi_i} f_{j,s} \cdot RX_s + \sum_{j \neq i} f_{s,\phi_j} \cdot TX_{s,j} + \sum_{j \neq i} f_{\phi_j,s} \cdot RX_s \leq \mathcal{E}_{max}, \quad \forall s \in \phi_i \tag{16}$$

$$\sum_{s \in \phi_i} f_{s,\phi_j} = f_{\phi_i,\phi_j}, \quad \forall j \neq i \tag{17}$$

$$\sum_{s \in \phi_i} f_{\phi_j,s} = f_{\phi_j,\phi_i}, \quad \forall j \neq i \tag{18}$$

Flow Constraints:

$$\sum_{j \in \phi_i} \pi_{j,k}^{(s)} = \sum_{j \in \phi_i} \pi_{k,j}^{(s)} + \sum_{j \neq i} \pi_{k,\phi_j}^{(s)}, \quad \forall k \in \phi_i \text{ and } k \neq s \tag{19}$$

$$\sum_{j \in \phi_i} \pi_{j,s}^{(s)} + T = \sum_{j \in \phi_i} \pi_{s,j}^{(s)} + \sum_{j \neq i} \pi_{s,\phi_j}^{(s)} \tag{20}$$

$$0 \leq \pi_{k,j}^{(s)} \leq f_{k,j}, \quad \forall j \in \phi_i \text{ and } \forall k \in \phi_i \tag{21}$$

$$0 \leq \pi_{k,\phi_j}^{(s)} \leq f_{k,\phi_j}, \quad \forall j \neq i \tag{22}$$

$$\sum_{j \neq i, k \in \phi_i} \pi_{k,\phi_j}^{(s)} = T, \tag{23}$$

where $s \in \phi_i$ and all variables are required to be non-negative.

Table 2: Linear program for Phase II of INCREMENTAL CMLDA.

Consider a super-sensor ϕ_i . In order to live for T rounds, each sensor s in ϕ_i needs to transmit T packets to the base station t (ϕ_{m+1}). Observe that, each super-sensor ϕ_i has a maximum flow of size $\geq T$ to ϕ_{m+1} in G . Stated otherwise, each super-sensor ϕ_i has a minimum ϕ_i - ϕ_{m+1} cut of size $\geq T$ in G . Therefore, for each sensor s in ϕ_i , we need to determine the capacity provisions between s and the remaining sensors in ϕ_i , as well as between s and each of the super-sensors $\phi_1, \dots, \phi_{i-1}, \phi_{i+1}, \dots, \phi_{m+1}$, such that:

- (i) the sum of the provisioned capacities from (to) all the sensors in ϕ_i to (from) each super-sensor ϕ_j equals the provisioning f_{ϕ_i,ϕ_j} (f_{ϕ_j,ϕ_i}) obtained from Phase I, and
- (ii) each sensor s in ϕ_i can push T packets (flow) to the remaining super-sensors. Note that this guarantees that each s - ϕ_{m+1} (s - t) minimum cut in the resulting flow network has capacity $\geq T$, since each minimum ϕ_i - ϕ_{m+1} cut in G has capacity $\geq T$.

In doing so, our objective is to *minimize* the maximum energy consumed by any sensor within the super-sensor ϕ_i , thereby extending the lifetime of the sensors. The necessary provisions can be computed in polynomial-time using the linear program in Table 2.

For each super-sensor ϕ_1, \dots, ϕ_m , the linear program minimizes the maximum energy consumed by any sensor within the super-sensor, subject to the energy constraint (16), the capacity constraints (17) and (18), and the flow constraints (19)–(23). For each sensor $s \in \phi_i$, constraint (16) gives the energy consumed at the sensor. Constraints (17) and (18) ensure that the sum of the provisioned capacities from (to) all the sensors in ϕ_i to (from) each super-sensor ϕ_j equals the provisioning f_{ϕ_i,ϕ_j} (f_{ϕ_j,ϕ_i}). Constraints (19) and (20) enforce the flow conservation principle at a sensor; constraints (21) and (22) ensure that the capacity constraints on the edges of the flow network are respected; and constraint (23) ensures that each sensor $s \in \phi_i$ has a minimum cut of size T to the remaining super-sensors.

Objective:

$$\text{minimize } \mathcal{E}_{max} \tag{24}$$

Energy and Capacity Constraints for ϕ_l :

$$\sum_{j \in \phi_r} f_{i,j} \cdot TX_{i,j} + \sum_{j \in \phi_r} f_{j,i} \cdot RX_i \leq \mathcal{E}_{max}, \quad \forall i \in \phi_l \tag{25}$$

$$\sum_{j \in \phi_r} f_{i,j} = f_{i,\phi_r}, \quad \forall i \in \phi_l \tag{26}$$

$$\sum_{j \in \phi_r} f_{j,i} = f_{\phi_r,i}, \quad \forall i \in \phi_l \tag{27}$$

Energy and Capacity Constraints for ϕ_r :

$$\sum_{j \in \phi_l} f_{i,j} \cdot TX_{i,j} + \sum_{j \in \phi_l} f_{j,i} \cdot RX_i \leq \mathcal{E}_{max}, \quad \forall i \in \phi_r \tag{28}$$

$$\sum_{j \in \phi_l} f_{i,j} = f_{i,\phi_l}, \quad \forall i \in \phi_r \tag{29}$$

$$\sum_{j \in \phi_l} f_{j,i} = f_{\phi_l,i}, \quad \forall i \in \phi_r \tag{30}$$

where all variables are required to be non-negative integers.

Table 3: Linear program for Phase III of INCREMENTAL CMLDA.

Phase III:

From Phase II, we obtain the capacity provisions between any sensor s and all other sensors in the same super-sensor. In addition, we know the provisions that are required between s and each of the remaining super-sensors. However, in order to construct the admissible flow network for the sensors, we need to determine the capacities that need to be provisioned between individual sensors in different super-sensors.

In particular, consider any two distinct super-sensors ϕ_l and ϕ_r ($l, r = 1, 2, \dots, m + 1$).⁷ For each sensor i in ϕ_l , we know the capacity f_{i,ϕ_r} ($f_{\phi_r,i}$) provisioned from (to) sensor i to (from) super-sensor ϕ_r . Similarly, for each sensor j in ϕ_r , we know the capacity f_{j,ϕ_l} ($f_{\phi_l,j}$) provisioned from (to) sensor j to (from) ϕ_l . To complete the construction, we provision capacities between pairs of sensors from ϕ_l and ϕ_r , while ensuring that :

- (i) the total capacity provisioned from (to) each sensor $i \in \phi_l$ to (from) all the sensors in ϕ_r equals the provisioning f_{i,ϕ_r} ($f_{\phi_r,i}$) obtained from Phase II, and
- (ii) the total capacity provisioned from (to) each sensor $j \in \phi_r$ to (from) all the sensors in ϕ_l equals the provisioning f_{j,ϕ_l} ($f_{\phi_l,j}$).

In doing so, we again maintain the objective of minimizing the maximum energy consumed by each sensor in any of the super-sensors. This is a *matching* problem which can be solved in polynomial-time using the linear program in Table 3.

For each pair of super-sensors ϕ_l and ϕ_r , the linear program minimizes the maximum energy consumed by any sensor within the super-sensors, subject to the constraints (25)–(30). For each sensor $i \in \phi_l$ and each sensor $j \in \phi_r$, constraints (25) and (28) give the energy consumed at the respective sensors. Constraints (26) and (27) guarantee that the total capacity provisioned from (to) each sensor $i \in \phi_l$ to (from) all the sensors

⁷We do not provision capacities from the base station t to any sensor.

in ϕ_r equals the provisioning f_{i,ϕ_r} ($f_{\phi_r,i}$). Similarly, constraints (29) and (30) ensure that the total capacity provisioned from (to) each sensor $j \in \phi_r$ to (from) all the sensors in ϕ_l equals the provisioning f_{j,ϕ_l} ($f_{\phi_l,j}$).

Phase IV:

At the end of Phase III, we obtain an admissible flow network with lifetime T for all n sensors and the base station t with capacities provisioned between them. Note that, these capacities are fractional non-negative numbers. In order to obtain a data gathering schedule for the sensors, we first compute the the maximum energy \mathcal{E}_{max} consumed by any sensor. Then, using Lemma 4 below, we scale the provisioned capacities by a factor of $\alpha = \mathcal{E}/\mathcal{E}_{max}$, and thus obtain an admissible flow network G_o with lifetime $T_o = \alpha T$ that respects the initial energy of the sensors.

Lemma 4 *Given the location of n sensors and a base station t , the lifetime of an admissible flow network scales linearly with the energy of the sensors.*

Proof. Follows directly from the energy constraint (5) for the sensors in the formulation of the MLDA problem in Table 1. ■

Next, we floor all the capacities of G_o to obtain an admissible flow network with integer capacities (flows). We call this flow network the INCREMENTAL CMLDA flow network. Using this flow network, we finally compute the integral system lifetime (as in the MLDA algorithm), and a data gathering schedule \mathcal{S} using the GETSCHEDULE algorithm from section III.

Figure 7 provides an illustration of the INCREMENTAL CMLDA heuristic. We consider a set of 4 sensors and a base station labeled 5, clustered into super-sensors \mathcal{Q}_1 , \mathcal{Q}_2 and \mathcal{Q}_3 . Figure 7(a) shows an optimal admissible flow network for the super-sensors with lifetime 100. We use this flow network to provision capacities for the sensors within each super-sensor, during Phase II of the heuristic. The resulting network is shown in figure 7(b). Observe that, sensors 1 and 2 in \mathcal{Q}_1 , each have a min-cut of size 100 to super-sensors \mathcal{Q}_2 and \mathcal{Q}_3 . Similarly, sensors 3 and 4 in \mathcal{Q}_2 , each have a min-cut of size 100 to super-sensors \mathcal{Q}_1 and \mathcal{Q}_3 . Finally, we solve the matching problem in Phase III of the heuristics, to determine the capacity provisioning between pairs of sensors from \mathcal{Q}_1 and \mathcal{Q}_2 . The INCREMENTAL CMLDA flow network for the sensors with lifetime of 100 rounds is shown in figure 7(c). Note that each sensor has a min-cut of size 100 to the base station 5.


Worst-case running time of INCREMENTAL CMLDA:

The INCREMENTAL CMLDA heuristic consists of four phases. Phase I involves solving the linear program (4), which has $O(m^3)$ variables, and thus the worst-case running time of Phase I is $O(m^{15} \log m)$ (see Lemma 1 and Theorem 3). In Phase II, we solve m instances of the linear program (4), each with $O(cm^2)$ variables. Each instance of the linear program (4) can be solved in worst-case running time $O(c^5 m^{10} \log(cm))$, using Lemma 1. Hence, the worst-case running time of Phase II is $O(c^5 m^{11} \log(cm))$. In Phase III, we solve the linear program (4) with $O(c^2)$ variables in worst-case running time $O(c^{10} \log c)$. Finally, Phase IV has worst-case running time $O(n^2)$, since we need to find the maximum energy consumed by any sensor.


Therefore, the total worst-case running time of INCREMENTAL CMLDA is

$$O(m^{15} \log m) + O(c^5 m^{11} \log(cm)) + O(c^{10} \log c) + O(n^2). \tag{31}$$


As discussed in section V, by appropriately choosing the number of super-sensors m and the maximum number c of sensors within any super-sensor, we can achieve a significant reduction in the actual running time of INCREMENTAL CMLDA with respect to the MLDA. For example, for $m = n^{5/10.5}$, the worst-case running time of the INCREMENTAL CMLDA is $O(n^{82.5/10.5} \log n) = O(n^{7.86} \log n)$.



(a) An admissible flow network from Phase I for super-sensors Q_1 and Q_2 with base station in Q_3 .



(b) Capacity provisioning from Phase II for sensors 1, 2 in super-sensor Q_1 and sensors 3, 4 in super-sensor Q_2 .



(c) The INCREMENTAL CMLDA flow network for the sensors with lifetime of 100 rounds.

Figure 7: Illustrative example of the INCREMENTAL CMLDA heuristic.

V Experiments

In this section, we compare the data gathering schedule given by the near-optimal MLDA algorithm, the GREEDY CMLDA heuristic and the INCREMENTAL CMLDA heuristic, with that obtained from a chain-based 3-level hierarchical protocol proposed by Lindsey, Raghavendra and Sivalingam [16]. For brevity, we refer to this protocol as the LRS protocol. We choose this protocol since it outperforms other competitive protocols (e.g. LEACH [8]) in terms of system lifetime.

LRS protocol for constructing a data gathering schedule: In this protocol, sensor nodes are initially grouped into clusters based on their distances from the base station. A chain is formed among the sensor nodes in a cluster at the lowest level of the hierarchy. Gathered data, moves from node to node, gets aggregated, and reaches a designated leader in the chain i.e. the cluster head. At the next level of the

hierarchy, the leaders from the previous level are clustered into one or more chains, and the data is collected and aggregated in each chain in a similar manner. Thus, for gathering data in each round, each sensor transmits to a close neighbor in a given level of the hierarchy. This occurs at every level, the only difference being that nodes that are receiving at each level are the only nodes that rise to the next level in the hierarchy. Finally at the top level, there is a single leader node transmitting to the base station. To increase the lifetime of the system, the leader in each chain is chosen in a round-robin manner in each round. Observe that, the manner in which chain leaders are selected in each level of the hierarchy, naturally defines an aggregation tree, for each round of data gathering.

For the initial set of experimental results, we consider a network of sensors randomly distributed in a $50\text{m} \times 50\text{m}$ field. The number of sensors in the network, i.e. the network size n , is varied to be 40, 50, 60, 80 and 100 respectively. Each sensor has an initial energy of 1J and the base station is located at (25m, 150m). Each sensor generates packets of size 1000 bits. The energy model for the sensors is based on the first order radio model described in section II.

Each experiment corresponds to a random placement of the sensors, for a particular network size. In each experiment, we measure the lifetime T , i.e. the number of rounds before the first sensor is drained of its energy, for the data gathering schedule given by the LRS protocol. For the same placement of sensors, we measure the lifetime of the data gathering schedules obtained from MLDA, GREEDY CMLDA (G-CMLDA) and INCREMENTAL CMLDA (I-CMLDA). We define the *performance ratio* R_M as the ratio of the system lifetime achieved using MLDA to the lifetime given by the LRS protocol. Similarly, we define R_G (R_I) to be the ratio of the system lifetime achieved using the GREEDY CMLDA (INCREMENTAL MLDA) heuristic to the lifetime given by the LRS protocol. Recall that, the (integral) solution given by MLDA is an approximation of the optimal fractional solution. We denote OPT to be the optimal system lifetime for any particular experiment.

For a data gathering schedule \mathcal{S} , we define the *depth* of a sensor v to be its average number of hops from the base station in the schedule, i.e. the average of its depths in each of the aggregation trees in \mathcal{S} . The depth of the schedule is defined as $\max\{\text{depth}(i)\}$, among all sensors i in the network. We measure the depth D of a schedule constructed using each of the MLDA, GREEDY (INCREMENTAL) CMLDA and LRS algorithms. Note that, the depth of a data gathering schedule is an interesting metric since it gives an estimate of the (maximum) average delay⁸ that is incurred in sending data packets from any sensor to the base station.

Finally, for the GREEDY (INCREMENTAL) CMLDA heuristics, we denote c to be the number of sensors in a cluster (super-sensor). Given the location of the sensors and the base station, we employ a greedy clustering algorithm similar to the chain-forming algorithm used by the LRS protocol [16] – pick a sensor i farthest from the base station and form a cluster that includes i and its $c - 1$ nearest neighbors; continue the process with the remaining sensors until all sensors have been included in some cluster. For a particular network size, we assign the size of a cluster in GREEDY (INCREMENTAL) CMLDA to be identical to the size of a chain in the LRS protocol. By clustering the sensors in the above manner, we can efficiently compute maximum lifetime schedules for the super-sensors in large networks. Further, with a proper choice of the number of super-sensors and the size of each super-sensor, we can solve the linear programs in Phases II and III of the INCREMENTAL CMLDA heuristic in a fast and efficient manner. Observe that, the MLDA algorithm and the GREEDY (INCREMENTAL) CMLDA heuristics presented in this paper are essentially centralized in nature. This implies that the clustering of the sensors need to be pre-computed at the base station. Similarly, an appropriate data gathering schedule is pre-computed at the base station (which is less likely to be resource-constrained) and transmitted to the individual sensors. We take advantage of the fact that the base-station is aware of the locations of the sensors and have sufficient processing capabilities to compute efficient data-gathering schedule(s) for the sensors.

Tables 4 and 5 summarize our main results. Note that the presented values for lifetime and depth are averaged across 20 different experiments for each network size. Further, the MIN and MAX columns for

⁸On a 1Mbps link, a 1000 bit message can incur a delay of 1 ms on each hop to the base station.

R_M , R_G and R_I indicate the corresponding minimum and maximum performance ratios observed from those experiments. We make the following key observations for the smaller sensor networks:

INPUT		MLDA			G-CMLDA		I-CMLDA		LRS	
n	c	OPT	T	D	T	D	T	D	T	D
40	5	6611.8	6610	4.9	6442	4.6	6512	4.8	5592	4.4
50	5	6809.0	6808	5.8	6747	5.6	6786	5.5	5466	5.1
60	5	7176.2	7174	6.2	6896	6.0	7084	6.3	5872	5.2
80	10	7946.9	7945	7.5	7509	6.6	7809	6.9	6002	6.1
100	10	8292.6	8290	8.2	8011	7.2	8121	7.8	5526	6.6

Table 4: Experimental results for 50m \times 50m sensor network.

INPUT		R_M		R_G		R_I	
n	c	MIN	MAX	MIN	MAX	MIN	MAX
40	5	1.15	1.48	1.10	1.42	1.15	1.45
50	5	1.20	1.90	1.20	1.65	1.20	1.72
60	5	1.18	1.66	1.18	1.62	1.16	1.64
80	10	1.27	2.05	1.21	2.01	1.27	1.96
100	10	1.42	2.32	1.33	2.16	1.38	2.24

Table 5: Performance ratios for 50m \times 50m sensor network.

- the lifetime of a schedule obtained using the INCREMENTAL CMLDA heuristic is always within 3% of the optimal fractional solution, while the lifetime of a schedule given by the GREEDY CMLDA heuristic is always within 9% of the optimal solution.
- the lifetime of a schedule given by the MLDA algorithm is near-optimal. In fact, the approximation scheme in section III leads to a reduction in the (integral) system lifetime by no more than 3 rounds, when compared to the optimal solution.
- the MLDA algorithm significantly outperforms the LRS protocol in terms of system lifetime. Further, the GREEDY CMLDA heuristic performs 1.10 to 2.16 times better than LRS, while the INCREMENTAL CMLDA heuristic performs 1.15 to 2.24 times better than LRS.
- the average depth of a data gathering schedule attained by the GREEDY (INCREMENTAL) CMLDA heuristic is slightly higher than that of the LRS protocol. Note that the 3 level protocol in LRS is specifically devised to reduce the average depth of each sensor [16]. To this end, the GREEDY (INCREMENTAL) CMLDA heuristic does quite well in attaining comparable sensor depths, while delivering significant improvements in system lifetime.

INPUT	MLDA	G-CMLDA	I-CMLDA
40	6.52	11.8	12.5
50	48.0	16.8	22.9
60	108.5	17.2	48.1
80	252.6	24.1	80.8
100	335.6	29.2	106.5

Table 6: Comparison of CPU times (in minutes) for MLDA, GREEDY CMLDA and INCREMENTAL CMLDA.

For each of the above experiments, we also measured its total running time (in terms of CPU time) when using MLDA, GREEDY CMLDA and INCREMENTAL CMLDA respectively. Table 6 shows the average CPU

times for different network sizes. Observe that, each of these experiments were conducted using MATLAB implementations of the algorithms on typical machine configurations of 866 MHz (Intel) processor/ 1GB physical memory. As shown in Table 6, the INCREMENTAL CMLDA heuristic has an average running time that is approximately 33% of that for MLDA, while the GREEDY CMLDA heuristic has an average running time that is about 10% of that for MLDA.

For our next set of experiments, we consider larger networks of sensors randomly distributed in a 100m \times 100m field. The number of sensors in the network, i.e. the network size n , is varied to be 100, 200, 300 and 400 respectively. Each sensor has an initial energy of 1J and the base station is located at (50m, 300m). Once again, the presented values for lifetime and depth are averaged across 20 different experiments for each network size. Due to the high complexity of the algorithm, we do not include any results regarding the performance of MLDA for the large-scale networks. The clustering in CMLDA (chain formation in LRS) is done in the manner described above. We summarize our results in Tables 7 and 8.

We make the following observations:

- the GREEDY CMLDA and INCREMENTAL CMLDA heuristics significantly outperform the LRS protocol in terms of system lifetime. In particular, the GREEDY CMLDA heuristic delivers system lifetimes that are 1.20 to 2.27 times larger than LRS; while the INCREMENTAL CMLDA heuristic attains lifetimes that are 1.22 to 2.61 times larger than LRS.
- the INCREMENTAL CMLDA heuristic obtains systems lifetimes that are 10–20% larger than the GREEDY CMLDA heuristic. Observe that, this improvement comes at the cost of increased complexity involved with INCREMENTAL CMLDA.
- the average depth of a data gathering schedule attained by the GREEDY (INCREMENTAL) CMLDA heuristic is only slightly higher than that of the LRS protocol.

INPUT		G-CMLDA		I-CMLDA		LRS	
n	c	T	D	T	D	T	D
100	10	3200	7.5	3611	7.8	2458	6.9
200	10	4086	10.6	4512	10.6	2854	9.6
300	15	4858	13.2	5560	13.1	3212	12.1
400	20	5202	20.6	6142	19.8	3654	18.6
500	25	5533	26.1	6577	25.8	3596	24.8

Table 7: Experimental results for 100m \times 100m sensor network.

INPUT		R _G		R _I	
n	c	MIN	MAX	MIN	MAX
100	10	1.32	1.91	1.22	2.61
200	10	1.20	2.27	1.30	2.43
300	15	1.29	2.13	1.50	2.41
400	20	1.31	1.88	1.40	2.02
500	25	1.33	2.09	1.36	2.33

Table 8: Performance ratios for 100m \times 100m sensor network.

In conclusion, our experimental results demonstrate that the clustering-based heuristics can achieve as much as a factor of 2.61 increase in the lifetime of large-scale sensor networks, when compared to the LRS protocol, and incurs a small increase in the delay experienced by individual sensors. Thus, while it is true that our algorithms are centralized, the payback is significant in terms of improvement in the system lifetime.

VI Conclusions

In this paper, we proposed a polynomial-time near-optimal algorithm (MLDA) for solving the maximum lifetime data gathering problem for sensor networks, when the sensors are allowed to perform in-network aggregation of data packets. Given the complexity of the MLDA algorithm, we next described efficient clustering-based heuristics to solve the maximum lifetime data aggregation problem in large sensor networks. Further, we presented experimental results demonstrating that the proposed methods attain significant improvements in system lifetime, when compared to existing protocols.

There are a number of important issues related to the maximum lifetime data gathering problem that need to be investigated in the future. In the work presented in this paper, we make the simplistic assumption that a sensor can always aggregate its own data packets with those of any other sensor in the network. As part of our current research, we are exploring a more complex scenario where a sensor is permitted to aggregate its own packets with only certain sensors, while acting as a router for other incoming packets. In the future, we plan to investigate modifications to the MLDA algorithm that would allow sensors to be added to (removed from) the network, without having to re-compute the entire schedule. Further, we plan to study the data gathering problem with depth (delay) constraints for individual sensors, in order to attain desired tradeoffs between the delay experienced by the sensors and the lifetime achieved by the system.

References

- [1] D. Bertsimas and J. N. Tsitsiklis. *Introduction to Linear Optimization*. Athena Scientific, Massachusetts, 1997.
- [2] M. Bhardwaj, T. Garnett, and A.P. Chandrakasan. Upper Bounds on the Lifetime of Sensor Networks. In *Proceedings of International Conference on Communications*, 2001.
- [3] J.H. Chang and L. Tassiulas. Energy Conserving Routing in Wireless Ad-hoc Networks. In *Proceedings of IEEE INFOCOM*, 2000.
- [4] J.H. Chang and L. Tassiulas. Maximum Lifetime Routing in Wireless Sensor Networks. In *Proceedings of Advanced Telecommunications and Information Distribution Research Program, College Park, MD*, 2000.
- [5] T. H. Cormen, C. E. Leiserson, and R. L. Rivest. Max-flow Min-cut Theorem. In *Introduction to Algorithms, MIT Press*, 1998.
- [6] J. Edmonds. Edge-disjoint branchings. In *Combinatorial Algorithms, Academic Press*, 1973.
- [7] A. V. Goldberg and S. Rao. Beyond the Flow Decomposition Barrier. In *Proceedings of 38th IEEE Annual Symposium on Foundation of Computer Science*, 1997.
- [8] W. Heinzelman, A.P. Chandrakasan, and H. Balakrishnan. Energy-Efficient Communication Protocols for Wireless Microsensor Networks. In *Proceedings of Hawaiian International Conference on Systems Science*, 2000.
- [9] W. Heinzelman, J. Kulik, and H. Balakrishnan. Adaptive Protocols for Information Dissemination in Wireless Sensor Networks. In *Proceedings of 5th ACM/IEEE Mobicom Conference*, 1999.
- [10] C. Intanagonwiwat, R. Govindan, and D. Estrin. Directed diffusion: A scalable and robust communication paradigm for sensor networks. In *Proceedings of 6th ACM/IEEE Mobicom Conference*, 2000.
- [11] J. M. Kahn, R. H. Katz, and K. S. J. Pister. Mobile Networking for Smart Dust. In *Proceedings of 5th ACM/IEEE Mobicom Conference*, 1999.

- [12] K. Kalpakis, K. Dasgupta, and P. Namjoshi. Maximum Lifetime Data Gathering and Aggregation in Wireless Sensor Networks. In *To Appear in Proceedings of IEEE Networks'02 Conference*, 2002.
- [13] B. Krishnamachari, D. Estrin, and S. Wicker. The Impact of Data Aggregation in Wireless Sensor Networks. In *Proceedings of International Workshop on Distributed Event-Based Systems*, 2002.
- [14] X. Lin and I. Stojmenovic. Power-aware Routing in Ad Hoc Wireless Networks. In *University of Ottawa, TR-98-11*, 1998.
- [15] S. Lindsey and C. S. Raghavendra. PEGASIS: Power Efficient GAttering in Sensor Information Systems. In *Proceedings of IEEE Aerospace Conference*, 2002.
- [16] S. Lindsey, C. S. Raghavendra, and K. Sivalingam. Data Gathering in Sensor Networks using the Energy*Delay Metric. In *Proceedings of the IPDPS Workshop on Issues in Wireless Networks and Mobile Computing*, 2001.
- [17] L. Lovász. On two minimax theorems in graph theory. In *Journal of Combinatorial Theory Series B, Vol. 21*, 1976.
- [18] S. Madden, R. Szewczyk, M. J. Franklin, and D. Culler. Supporting Aggregate Queries Over Ad-Hoc Wireless Sensor Networks. In *Proceedings of 4th IEEE Workshop on Mobile Computing and Systems Applications*, 2002.
- [19] R. Min, M. Bhardwaj, S.H. Cho, A. Sinha, E. Shih, A. Wang, and A.P. Chandrakasan. Low-Power Wireless Sensor Networks. In *VLSI Design*, 2001.
- [20] J. Rabaey, J. Ammer, J.L. da Silva Jr, and D. Patel. PicoRadio: Ad-hoc Wireless Networking of Ubiquitous Low-Energy Sensor/Monitor Nodes. In *Proceedings of the IEEE Computer Society Annual Workshop on VLSI*, 2000.
- [21] S. Singh, M. Woo, and C. Raghavendra. Power-aware Routing in Mobile Ad Hoc Networks. In *Proceedings of 4th ACM/IEEE Mobicom Conference*, 1998.

Appendix: Proof of Theorem 2

The proof is based on a powerful theorem in graph theory ([6, 17]).

Preliminaries:

A *multi-graph* is a graph which may have parallel edges, i.e. it may have multiple edges with the same end-vertices. An *arborescence* \mathcal{A} rooted at a vertex t is a directed graph such that there is a unique directed path from t to each vertex in \mathcal{A} . A *spanning arborescence* of a multi-graph is an arborescence that spans all its vertices. The *transpose* G^T of a directed graph G is the graph that results from G by reversing the directions of all of its edges. Given a graph $G = (V, E)$ and an edge $e \in V \times V$, let $G + e$ denote the graph resulting from G by adding e to it.

Consider a directed graph $G = (V, E)$ with positive integer capacities $c(e)$ assigned to its edges $e \in E$. Let \hat{G} be the graph that results from G by replacing each one of its edges $e = (u, v)$ with $c(e)$ parallel edges (u, v) ; for the multi-graph $c(e) = c(u, v)$ will denote the number of parallel edges (u, v) .

For any subset $S \subset V$, let (S, \bar{S}) , called a *cut*, denote the set of edges in G from S to $\bar{S} = V - S$. The *capacity (size)* of a cut (S, \bar{S}) , denoted with $\omega_G(S)$, is sum of the capacities of the edges in (\bar{S}, S) ; the capacity $\omega_{\hat{G}}(S)$ of the cut (S, \bar{S}) with respect to \hat{G} equals the total number of edges in (S, \bar{S}) in \hat{G} . A $u-v$ cut is a cut (S, \bar{S}) such that $u \in S$ and $v \in \bar{S}$. A $u-v$ min-cut is a $u-v$ cut with minimum capacity.

For any subgraph H of G let $G - H$ denote the graph resulting from G by deleting the edges in H . Similarly, let $\hat{G} - \lambda H$ denote the graph resulting from the multi-graph \hat{G} by deleting λ copies of H from it, $0 \leq \lambda \leq \min\{c(e) : e \in H\}$. Note that the operation $\hat{G} - \lambda H$ corresponds to a reduction of the capacities of the edges of G , which are also in H , by λ ; thus, for convenience, we also denote $\hat{G} - \lambda H$ by $G - \lambda H$.

Definition 3: Consider a multi-graph $\hat{G} = (V, \hat{E})$. Let t be a vertex in V , and k a positive integer. Suppose that, for all $v \in V$, each t - v cut has size $\geq k$. An arborescence $\mathcal{A} = (\mathcal{V}_{\mathcal{A}}, \mathcal{E}_{\mathcal{A}})$ of the multi-graph \hat{G} is called *admissible* if for every $S \subset V$, with $t \in S$, $\omega_{\hat{G}-\mathcal{A}}(S) \geq k - 1$.

Note that an arborescence with no edges is admissible.

Lemma 5 Consider a graph $G = (V, E)$ with capacities associated with its edges. Let (S, \bar{S}) be any cut in G . Then the capacity of (S, \bar{S}) in G is equal to the capacity of (\bar{S}, S) in G^T and vice-versa.

Proof. Follows directly from the definition of cut capacities, and the definition of the transpose of a graph. ■

Theorem 4 (Edmonds (1973 [6]), Lovász(1979 [17])) Let $\hat{G} = (V, \hat{E})$ be a multi-graph, t be a vertex in V , and k be a positive integer. If, for all $v \in V$, each t - v cut has size $\geq k$, then \hat{G} has k edge-disjoint spanning arborescences rooted at t .

Moreover, the following corollary follows from the construction in the proof of Theorem 4.

Corollary 1 Let $\hat{G} = (V, \hat{E})$ be a multi-graph, t be a vertex in V , and k a positive integer. Suppose that, for all $v \in V$, each t - v cut has size $\geq k$. Let $\mathcal{A} = (\mathcal{V}_{\mathcal{A}}, \mathcal{E}_{\mathcal{A}})$ be an admissible arborescence of \hat{G} . If \mathcal{A} is not spanning V , there there exists an edge $e \in (V_{\mathcal{A}}, \bar{V}_{\mathcal{A}})$ such that \mathcal{A} together with e is also an admissible arborescence.

Proof sketch. Follows directly from the construction in the proof of Theorem 4. ■

Proof of Theorem 2:

Since G is admissible, each sensor node $v \in V$ has a v - t cut of capacity $\geq T$ to the base station t .

Observe that, the transpose \mathcal{A}^T of a spanning admissible arborescence \mathcal{A} of the multi-graph that corresponds to G^T , provides us with an aggregation tree for G for aggregating 1 value from each sensor. Moreover, from Corollary 1 we know that a spanning admissible arborescence for the multi-graph \hat{G}^T can be constructed using a simple greedy algorithm. Specifically, start with an arborescence \mathcal{A} that contains just t , and while \mathcal{A} is not spanning V , add an edge $e \in (V_{\mathcal{A}}, \bar{V}_{\mathcal{A}})$ such that $\mathcal{A} + e$ is admissible for \hat{G}^T . Thereby, it is easy to see that the GETTREE algorithm can always find an aggregation tree A for collecting one data packet from each sensor, such that the $(A, 1)$ -reduction of G is feasible. Further, if the (A, f) -reduction of G is feasible for some $f > 1$, we can use the aggregation tree A to collect f data packets from each sensor. This results in fewer aggregation trees and hence an improvement in the running time of our algorithm. Finally, we can iteratively invoke the GETTREE procedure to find a sequence of aggregation trees for collecting T data packets from each of the sensors. ■

Contents

I	Introduction	2
II	The Data Gathering Problem	4
	<i>A. System Model</i>	4
	<i>B. Problem Statement</i>	4
III	MLDA: Maximum Lifetime Data Gathering with Aggregation	4
	<i>A. Finding a near-optimal admissible flow network</i>	5
	<i>B. Constructing a schedule from an admissible flow network</i>	6
	<i>C. Worst-case running time of MLDA</i>	8
IV	CMLDA: Clustering-based MLDA	9
	<i>A. The greedy CMLDA heuristic</i>	9
	<i>B. The Incremental CMLDA Heuristic</i>	11
V	Experiments	16
VI	Conclusions	20